Limits of Global Games

Latest Version

Rafael Veiel*

November 10, 2024

Abstract

Games with strategic complementarities often exhibit multiple equilibria. In a global game, players privately observe a noisy signal of the underlying payoff matrix. As the noise diminishes, a unique equilibrium is selected in almost all binary-action games with strategic complementarities - a property known as "limit uniqueness." This paper describes the limits of that approach in two-player games, as we move beyond two actions. Unlike binary-action games, limit uniqueness is not an intrinsic feature of all games with strategic complementarities. When the noise is symmetric, we demonstrate that limit uniqueness holds if and only if the payoffs exhibit a generalized ordinal potential property. Moreover, we provide an example illustrating how this condition can be easily violated.

*veiel@mit.edu, Massachusetts Institute of Technology. I would like to thank my advisors, Stephen Morris, Rob Townsend and Muhamet Yildiz for invaluable comments and suggestions. I would also like to thank Ian Ball, Roberto Corrao, Drew Fudenberg, Olivier Gossner, Nima Haghpanah, Moritz Poll, Iván Werning, Alex Wolitzky, the seminar participants at SAET and the participants of MIT theory lunch for very valuable and thoughtful feedback.

1 Introduction

Economic models with a coordination component, like investment games and bank runs, often have multiple equilibria. One way to resolve this indeterminacy is to relax the common knowledge assumption on payoffs, by letting players privately observe a noisy signal of the underlying payoff structure. Players are nudged into taking a certain action because their signal makes them believe others will do the same. In particular, if some signals make an action dominant for some player, then best-response behavior can unravel and select the equilibrium in which that action is played.

Carlsson and Van Damme (1993) have shown that for two-player, binary-action, coordination games, having players privately observe noisy signals of payoffs almost always selects a unique equilibrium as the noise vanishes. This property was later called "limit uniqueness." In their setup, players may observe a signal about more than one parameter of the payoff structure. They call the induced game of incomplete information a "global game." Frankel et al. (2003) obtain a similar result for games of strategic complementarities with many actions, but where players see a noisy signal of a one-dimensional parameter of the game.

Global games have been used extensively in macroeconomics and finance to model, among other phenomena: Liquidity crises (Morris and Shin, 2004), currency attacks (Morris and Shin, 1998), and bank runs (Goldstein and Pauzner, 2005). They offer an equilibrium selection device with an economic story: A small grain of doubt about the commitment of some players to playing the actions of one equilibrium and playing another action instead can "panic" rational players into playing a best response to that action. All of these applications introduce uncertainty about a one-dimensional parameter of a binary action game and obtain sharp predictions as the noise vanishes.

However, when we examine models with more than two actions, limiting players' uncertainty to a one-dimensional parameter may no longer be a reasonable assumption. Oury (2013) provides a first example of a supermodular game where two-dimensional noise breaks the limit uniqueness result that would hold under one-dimensional noise. This example indicates that in some cases, the effectiveness of the global game approach relies on one-dimensional noise. It is then important to understand the limitations of the global game approach: under what conditions does a small amount of noisy, private information lead to coordination on a single, unique equilibrium?

This paper answers that question in a two-player, many-actions generalization of the global game model in Carlsson and Van Damme (1993) with symmetric noise. In the spirit of the original global game model, we let players observe a noisy signal of every aspect of the game, moving beyond the one-dimensional set-up in Frankel et al. (2003). Signals are thus naturally multi-dimensional. To state our results, we generalize the condition of risk-dominance to action sets and show that every game admits such a set. In contrast to the binary-action case, limit uniqueness is no longer an intrinsic property of coordination games: For concave games with strategic complementarities, limit uniqueness holds if and only if payoffs admit a generalized ordinal potential on all risk-dominant sets of actions. The "only if" part of the result means that the global game approach actually relies on a lot of structure to serve as a reliable selection device when we move beyond two actions and one-dimensional noise.

Indeed, in games with a generalized ordinal potential, the incentive of all players to switch actions can be expressed using a common function. Moreover, Monderer and Shapley (1996) show that games with a generalized ordinal potential are precisely the games where all sequences of better-replies converge to an equilibrium. We show in an example that better-response cycles, and thus the nonexistence of a generalized ordinal potential, can arise in concave supermodular investment games when there is a simple asymmetry in the payoffs: Under-investment hurts one player more than over-investment while over-investment hurts the other player more than under-investment.

For general games, we provide an upper bound for the selection in terms of risk-dominant, best-reply closed sets and provide a sufficient condition for an action set to survive iterative deletion of dominated strategies as the noise vanishes: Limit multiplicity arises for games that have risk-dominant best-response cycles.

Global Game Model We study two-player games and introduce a novel framework for studying global games with many actions and high-dimensional noise. Our model departs from the original global game framework of Carlsson and Van Damme (1993) in several key ways.

First, we represent all two-player games with a fixed set of actions on a sphere by scaling payoffs. We introduce noise by rotating the payoffs by a small random angle, contrasting with the original model, where noise was added in an additively separable way. In a spherical model with rotational noise the amount of uncertainty is constant everywhere on the sphere. In a global game model that comprises payoff matrices of every scale, the amount of uncertainty depends on the scale of the payoffs contained in a signal. Indeed, in such a model payoffs with large entries carry less uncertainty for any fixed size of noise. We show that in the limit where the noise vanishes, both models deliver identical predictions.

Second, we restrict our analysis to symmetric noise distributions. We impose two types of symmetries: the distribution of random rotations (i.e., the noise) remains invariant under any orthogonal change of basis, and both players have the same noise distribution. By applying rotational noise on a sphere, we can more effectively leverage the symmetry properties of the noise.[1]

Third, our global game model can be viewed as a private values model. This means that a player's private signal perfectly reveals her own payoffs, allowing us to concentrate on the primary source of uncertainty faced by a player: the signal observed by her opponent. A player's signal reveals her own payoffs and also contains information about the payoffs faced by her opponent.

Finally, we focus on interim correlated rationalizability (ICR), an incomplete information version of rationalizability introduced by Dekel et al. (2007). This solution concept coincides with Bayes Nash equilibrium in supermodular games but may be more permissive in other games.

Methodological Idea This paper focuses on studying regions within the signal space where certain actions are rationalizable in the presence of noise, as well as the boundaries that separate these regions. Signals that lie on the boundary make a player just indifferent between two actions. Boundaries are thus characterized by a system of indifference constraints. As the noise diminishes, some boundaries may converge and ultimately collide. The regions they enclose can contain multiple rationalizable actions and may disappear in the limit. This must happen in the region of the sphere where limit uniqueness holds: regions, where a unique action profile is rationalizable, must collide with other such regions and leave no space for regions with multiplicity that may lie in between.

We characterize the points that must exist at the collision of any set of boundaries. By leveraging the symmetry of the noise, we demonstrate

[1] In Carlsson and Van Damme (1993) symmetry is not required for limit uniqueness but we were not able to establish a similar result in the multi-dimensional model.

that these collision points are contained within the zero-set of a symmetric multilinear form. This form is linear in the payoffs of each of the indifference constraints corresponding to the colliding boundaries. It thus represents the limit of all relevant indifference constraints involved in the collision. In the case of binary actions, this form becomes bilinear and describes the set of games where two action profiles are risk-dominant. The multilinear form thus provides a necessary condition for points to reside at the intersection of the boundaries.

To establish a necessary and sufficient condition, we require the stability of the zeros of the multilinear form: a small perturbation of players' signals in a direction that makes an action profile more profitable should introduce the appropriate slack with the correct sign in both players' indifference constraints. We derive a condition on payoffs that ensures this stability, which we term "aligned incentives." We demonstrate that aligned incentives are both necessary and sufficient for achieving limit uniqueness. For strictly concave and strictly supermodular games, we can restate this result in terms of the existence of a generalized ordinal potential. Furthermore, we utilize results of Monderer and Shapley (1996) to express this condition in terms of the nonexistence of better-response cycles.

Related Literature Our paper relates to Carlsson and Van Damme (1993), who were the first to introduce the global game framework for two-player, binary action games. There is a long list of applied theory papers using those techniques, see Morris and Shin (2003) for a survey. Ui (2001) first establishes a connection between potential games, as introduced in Monderer and Shapley (1996) and robustness to incomplete information as introduced in Kajii and Morris (1997). That paper shows that maximizing a potential is sufficient for an equilibrium to be robust to incomplete information. This notion of robustness allows for a richer set of perturbations compared to the global game model. In our, more restrictive global game model, we are able to prove the necessity of a potential condition.

Global games with strategic complementarities and many actions have also been studied in Frankel et al. (2003). In their set-up, players receive noisy signals about a one-dimensional parameter which affects payoffs monotonically. In that case, limit uniqueness holds for all games with strategic complementarities. The selected outcome may however depend on the fine details of the noise distribution. Going from one-dimensional noise to manydimensional noise breaks the limit uniqueness result. While limit uniqueness is harder to obtain in the multi-dimensional case, we show that the selection no longer depends on the details of the noise within the class of symmetric noise distributions considered in this paper.

Oury (2013) also studies global games with multi-dimensional noise and provides a sufficient condition for limit uniqueness: If an equilibrium is selected in every one-dimensional global game considered in Frankel et al. (2003) independently of the structure of the noise, then it is selected in the global game with multi-dimensional noise. Frankel et al. (2003) provide a sufficient condition for noise-independent limit uniqueness in terms of a local potential property. This condition is a special case of the generalized ordinal potential property used in our paper.

Combining Oury (2013) and Frankel et al. (2003) thus provides a sufficient condition for limit uniqueness in global games with multi-dimensional noise that is consistent with our characterization. In light of this result, the main contribution of this paper is threefold: First, we provide a full characterization of limit uniqueness for concave, supermodular games. Second, we introduce the spherical global game model as a way to analyze multi-dimensional global games. Our set-up and solution techniques are very different from Oury (2013) and allow for the study of multi-dimensional global games without reference to one-dimensional global game techniques. Third, we generalize risk-dominance to action sets and provide a sufficient condition for limit multiplicity in all games. The condition is based on an upper bound for the limit selection that we obtain from arguments in Kajii and Morris (1997) and from a generalization of risk-dominance in two action games.

Organization of the Paper The rest of the paper is organized as follows: Section 2 introduces the spherical global game model. Section 3 uses the binary action case as an illustrative example for the techniques used in the general analysis. The section recovers the results in Carlsson and Van Damme (1993) expressed in the language of this paper. Section 4 provides a sufficient condition for limit multiplicity in terms of risk-dominance. Section 5 derives key topological and algebraic properties of collisions. Section 6 introduces the property of aligned incentives and characterizes limit uniqueness. Section 7 applies this characterization to strictly concave, strictly supermodular games and characterizes limit uniqueness in terms of generalized ordinal potentials. Section 8 illustrates the failure of limit uniqueness in examples. Section

9 discusses the connection to the literature. All proofs can be found in Appendix A.

2 Global Game Model

We study two-player games and introduce a new framework for analyzing global games with many actions, when there is noise about every entry in the payoff matrix. The main distinctive feature of this model is the way in which noise is applied to payoffs. In the original global game model noise is applied in an additively separable way, while this paper considers rotational noise. Figure 1 provides a simple picture that illustrates both approaches. On the left, we sketch the elements of the basic global game model in Carlsson and Van Damme (1993). A payoff matrix y is randomly drawn from the space U of all games of a fixed size (i.e. for fixed action sets). Players don't observe y but are privately informed of a noisy signal of y, which is of the form $s = y + \sigma\epsilon$, where σ is a small positive number and ϵ a random payoff matrix in U. On the right, we illustrate the approach taken in this paper. We restrict attention to the unit sphere $S \subseteq U$ of scaled payoffs. A player's signal s is of the form $s = e^{\sigma E}(y/||y||_2)$, where $e^{\sigma E}$ represents a random rotation matrix, where the rotation is of magnitude σ.

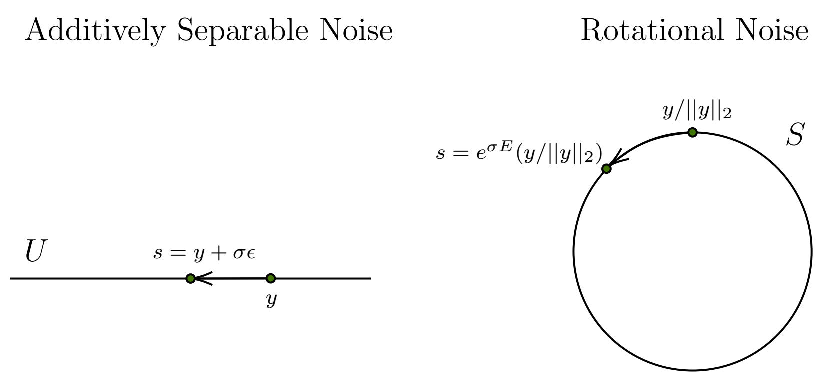

Figure 1: Global game with additively separable noise and spherical global game with rotational noise.

In Section 9.2 we provide the mapping between both models and show that in the limit where the noise vanishes, both models make identical predictions on the unit sphere S.

Notation We let $\mathbb{R}_+ := (0, \infty)$ denote the positive real numbers and $\mathbb{N}_+ := \{1, 2, \dots\}$ the set of positive integers. For any $n \in \mathbb{N}$, we let the Euclidean norm on $\mathbb{R}^n$ be denoted by $||\cdot||_2$. For any $n \in \mathbb{N}$, let $\mathrm{O}^n \subseteq \mathbb{R}^{n\times n}$ denote the set of orthogonal[2] $n \times n$ matrices and let $\mathrm{SO}^n \subseteq \mathrm{O}^n$ denote the set of orthogonal matrices with determinant equal to one. Note that these are rotation matrices. Every rotation matrix $R \in \mathrm{SO}^n$ can be written as a matrix exponential[3] $R = e^E$ for some skew-symmetric[4] matrix E.

2.1 Model

We now introduce the global game model, which is spherical. We consider a pair of players $\{1, 2\}$, where each player[5] i has an action set given by A_i, with $|A_i| < \infty$. We model all two-player games with action sets A_1, A_2 on a sphere, scaling payoffs accordingly. A game is represented by a vector of payoffs $u = (u_1, u_2) \in \mathbb{R}^{A_1\times A_2} \times \mathbb{R}^{A_2\times A_1}$. For any vector of payoffs s, we use superscripts to refer to its action-pair-index and subscripts to refer to the player-index, e.g. $s_i^{a_1,a_2}$ refers to player i's payoff when (a_1, a_2) is played. We define the global game state space S as the unit-sphere in $U := \mathbb{R}^{A_1\times A_2} \times \mathbb{R}^{A_2\times A_1}$,

$$S := \{u = (u_1, u_2) \in U : ||u||_2 = 1\}. \tag{2.1}$$

We introduce noise by applying a small random rotation to the scaled payoffs in S. This means that players will each privately observe a noisy signal of an underlying point in S, which we represent itself as a point in S. That is, signals will be draws in S. Let $n = |A_1 \times A_2|$, so that $2n = \dim(U)$. In order to describe random rotations on S we introduce the set of $2n \times 2n$-skew-symmetric matrices, which can be used to represent rotation matrices on the unit sphere in $\mathbb{R}^{2n}$. Indeed, every rotation matrix can be written as the matrix exponential e^T for some skew-symmetric matrix T. The advantage of representing rotations via matrix exponentials is that we can easily control the magnitude of the rotation by scaling the exponent: $e^{\sigma T}, \sigma \in \mathbb{R}_+$. This provides a useful modeling device to describe vanishing noise using matrix algebra.

[2]A matrix $T \in \mathbb{R}^{n\times n}$ is orthogonal if $T^{-1} = T^\top$, where $T^\top$ represents the transpose.

[3]The matrix exponential is defined as the series $e^E = \sum_{k=0}^{\infty} \frac{1}{k!} E^k$, where $E^0 = \mathrm{id}$.

[4]A matrix $E \in \mathbb{R}^{n\times n}$ is skew-symmetric if $E = -E^\top$.

[5]We will use i as generic notation for a player. If we use player indices i and $-i$, then $-i$ refers to the other player, i.e. $-i \neq i$.

Let $\mathcal{E}$ denote the collection of bounded $2n \times 2n$ skew-symmetric matrices:

$$\mathcal{E} := \{E \in [-1,1]^{2n\times 2n} : E^\top = -E\}. \tag{2.2}$$

A noisy signal consists of a random rotation on S. Since rotations are matrix exponentials of skew symmetric matrices, a random rotation can be identified with a random draw in $\mathcal{E}$.

A *spherical global game* is defined as a tuple of distributions $(\nu_0, (\nu_1, \nu_2)) \in \Delta(S) \times \Delta(\mathcal{E})^2$ with continuous, bounded densities. For every $\sigma \in \mathbb{R}_+$, a spherical global game gives rise to a Bayesian game, where each player i privately observes a signal

$$s = e^{\sigma E_i} y, \tag{2.3}$$

where the two random variables $(E_i, y) \in \mathcal{E} \times S$ are drawn independently: The random matrix $E_i \in \mathcal{E}$ is called the *noise term* and is drawn with distribution ν_i, and the random vector $y \in S$ is called the *latent common state* and is drawn with distribution ν_0.

The matrix exponential $e^{\sigma E_i}$ of σE_i represents a random rotation of the latent common state y. Let $\nu = \nu_1 \times \nu_2$ denote the product distribution of the noise terms. For each $\sigma > 0$, $\nu_\sigma \in \Delta(S^3)$ represents the induced joint distribution on the latent common state and signal pairs. This distribution ν_σ defines the common prior in the Bayesian game, where each player i receives a private signal s.

Each signal, a point in S, describes payoffs for both players at every action profile. While a player's own signal reveals her payoffs at all action profiles, it does not fully reveal the payoffs of her opponent. Thus, the payoffs of player $-i$, as described by player i's private signal, are not necessarily the true payoffs faced by $-i$. This distinction becomes negligible as $\sigma \to 0$, where the analysis simplifies.

Figure 2 below illustrates the key elements of a spherical global game with two actions per player, denoted $A_i = \{a, b\}$, for every player i. The payoff matrix in the middle represents the latent common state. With two actions, it is an element of $\mathbb{R}^8$, scaled appropriately. Players do not observe nor do they care about this state directly. Instead, each player i observes another payoff matrix, her private signal, that is obtained by applying a random rotation to the latent state. The private signals of each of the two players are depicted by the two payoff matrices at the bottom of Figure 2. The left one represents what player 1 sees and the right one represents what player 2 sees. Player 1's payoffs from choosing action a, when player 2 chooses action b is then given by the corresponding entry in her private signal: $s_1^{a,b}$ (highlighted in red). Similarly, player 2's payoffs from choosing action b, when player 1 chooses action a is then given by the corresponding entry in her private signal: $\hat{s}_2^{a,b}$ (highlighted in blue). Neither player observes the signal that the other player receives and so no player knows the payoffs that their opponent is facing.

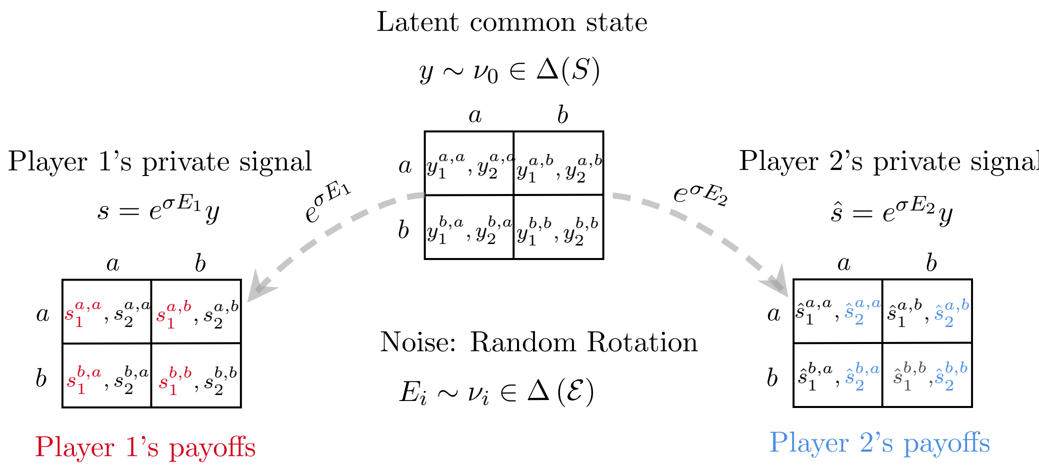

Figure 2: Illustration of spherical global game with two actions.

Say that ν is *symmetric* if $\nu_1 = \nu_2$ and for every $E \in \mathcal{E}$, every player i and every $X \in \mathrm{O}^n$,

$$f_{\nu_i}(X^\top E X) = f_{\nu_i}(E), \tag{2.4}$$

where f_{ν_i} is the density of ν_i. We will assume that ν is symmetric throughout. Examples include the uniform distribution or the normal distribution on $\mathbb{R}^{2n \times 2n}$ restricted to $\mathcal{E}$.

2.2 Rationalizability and Limit Selection

We now state the definition of Interim Correlated Rationalizability (ICR) as defined in Dekel et al. (2007). This solution concept builds on the foundations laid by rationalizability for complete information games introduced in (Bernheim, 1984; Pearce, 1984). Under ICR, for each signal they could receive, players iteratively delete actions which are strictly dominated when considering their expectation over state-contingent conjectures on their opponent's play at each of their signals.[6] ICR is a more permissive solution concept than Bayes-Nash equilibrium but coincides with it on supermodular games. Our results on multiplicity outside of supermodular games rely heavily on rationalizability being the solution concept.

Best Response Let the collection of non-empty action sets of player i be denoted by $\mathcal{A}_i := 2^{A_i} \setminus \varnothing$. For every player i, let $S_i := \{s_i : s \in S\}$ denote the projection of S onto i's payoffs. For any payoff of player i, $s_i \in S_i$, let the best-response to a belief $p \in \Delta(A_{-i})$ be given by

$$\mathrm{br}_i(p|s_i) := \arg\max_{a_i} \sum_{a_{-i} \in A_{-i}} p(a_{-i}) s_i^{a_i, a_{-i}}. \tag{2.5}$$

A correlated conjecture of player i is a ν_σ-measurable stochastic map, $\kappa_{-i} \colon S \times S \to \Delta(A_{-i})$. Every signal $s \in S$ defines a probability on A_{-i},

$$\nu_\sigma \circ \kappa_{-i}(a_{-i}|s) := \int_{S \times S} \kappa_{-i}(a_{-i}|y, s') \mathrm{d}\nu_\sigma(y, s'|s), \ \forall \ a_{-i} \in A_{-i}. \tag{2.6}$$

Player i's best-reply to a correlated conjecture κ_{-i}, given signal s, is then given by

$$\mathrm{BR}_i^\sigma(\kappa_{-i}|s) := \mathrm{br}_i(\nu_\sigma \circ \kappa_{-i}(\cdot|s)|s). \tag{2.7}$$

Rationalizability We now define ICR: Let $\mathrm{ICR}_i^{0,\sigma}(s) = A_i$ for each player i and signal $s_i \in S$. Given $\mathrm{ICR}_{-i}^{m-1,\sigma}(s') \subseteq A_{-i}$ for every signal $s' \in S$ of player $-i$, define

$$\mathrm{ICR}_i^{m,\sigma}(s) := \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^{m,\sigma}} \mathrm{BR}_i^\sigma(\kappa_i|s), \tag{2.8}$$

where $\mathcal{S}_{-i}^{m,\sigma} := \{\kappa_{-i} \colon S \times S \to \Delta(A_{-i}) : \forall\, y, s',\ \mathrm{supp}(\kappa_{-i}(y, s')) \subseteq \mathrm{ICR}_{-i}^{m-1,\sigma}(s')\}$.

Definition 2.1 (Rationalizability)**.** ICR *is given by*

$$\mathrm{ICR}_i^\sigma(s) := \bigcap_{m=0}^{\infty} \mathrm{ICR}_i^{m,\sigma}(s). \tag{2.9}$$

[6] Another natural solution concept would be interim independent rationalizability (IIR), where players' conjectures are not allowed to be correlated with the underlying state. We have not explored if there would be a difference between IIR and ICR in our global game model.

We end this section by introducing key terminology: Limit selection and limit uniqueness. The limit selection is defined as the limit of the ICR correspondence as the noise vanishes.

Definition 2.2 (Limit Selection)**.** *The* limit selection *is the map* ICR: $S \to \mathcal{A}_1 \times \mathcal{A}_2$ *so that for every* $s \in S$ *and player* i*,*

$$\liminf_{\sigma \downarrow 0} \mathrm{ICR}_i^\sigma(s) = \mathrm{ICR}_i(s). \tag{2.10}$$

We define limit uniqueness as a property of subsets of the sphere: A set satisfies limit uniqueness if, as the noise goes to zero, the probability of signals where ICR contains more than one action for at least one player also vanishes. We can easily extend this definition to points rather than sets, by requiring that every small enough neighborhood of the point satisfy the set-based definition of limit uniqueness below.

Definition 2.3 (Limit Uniqueness)**.** *Limit uniqueness holds on a set* $O \subseteq S$ *if*

$$\lim_{\sigma \downarrow 0} \nu_\sigma(\{s \in O : \exists\ i \text{ s.t. } |\mathrm{ICR}_i^\sigma(s)| > 1\}) = 0. \tag{2.11}$$

Boundaries and Collisions The main approach in this paper is to study regions within the signal space S where certain actions are rationalizable in the presence of noise, as well as the boundaries that separate these regions. The set of points in S with a given profile $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ of rationalizable actions is called the *Rationalizable Set of* B and is defined

$$\mathcal{R}^\sigma(B) \coloneqq \{s \in S : \mathrm{ICR}^\sigma(s) = B\}. \tag{2.12}$$

The limit set is denoted $\mathcal{R}(B) \coloneqq \liminf_{\sigma \downarrow 0} \mathcal{R}^\sigma(B)$. The boundary between B and B' is called the *boundary* and is defined as

$$\partial \mathcal{R}^\sigma(B, B') \coloneqq \overline{\mathcal{R}^\sigma(B)} \cap \overline{\mathcal{R}^\sigma(B')}, \tag{2.13}$$

where $\overline{\mathcal{R}^\sigma(\cdot)}$ denotes the topological closure of $\mathcal{R}^\sigma(\cdot)$ in S. Figure 3 provides a sketch of the rationalizable sets and boundaries. The rationalizable sets partition the sphere for every fixed choice of σ. Understanding properties of all the boundaries between all rationalizable sets allows us to provide a complete picture of the limit selection and thus of regions where limit uniqueness holds.

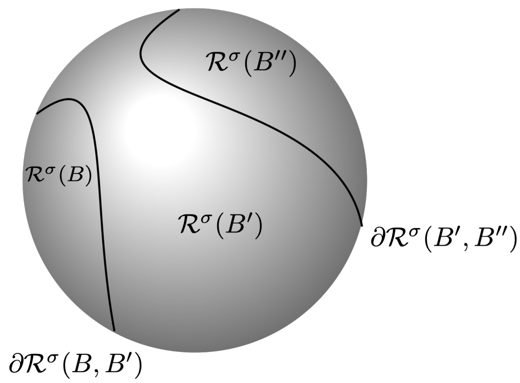

Figure 3: Rationalizable sets and their boundaries.

To illustrate these concepts in a more concrete setting, consider again the binary action environment. Figure 4 below illustrates three rationalizable regions for $\sigma > 0$. These regions partition S for every $\sigma > 0$. The figure shows three regions. The left most region (in light gray) is the region where both players rationalize the singleton $\{a\}$. In the middle region, (in white) player 2 rationalizes both actions while player 1 still rationalizes the singleton $\{a\}$. Finally, on the right most region (in dark gray) both players rationalize the full action set. The dashed curves that separate the regions are the boundaries.[7]

[7]As we show later, upper-hemi continuity properties of ICR imply that rationalizable sets contain their boundary with rationalizable sets in which fewer actions are rationalizable. This is the reason we define boundaries using the topological closure of the rationalizable sets.

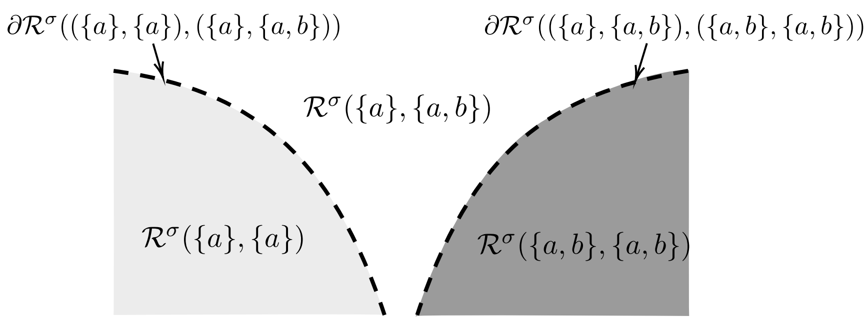

Figure 4: Illustration of rationalizable sets and their boundaries in binary actions case.

As the noise vanishes, some boundaries may converge and ultimately collide. The regions they enclose can contain multiple rationalizable actions and may disappear in the limit. We will study the limit selection through the lens of collisions of boundaries. A collision is the intersection of two or more limit boundaries.

Definition 2.4 (Collisions)**.** *For any collection of action set pairs* $Z = \{(B_1^1, B_2'^1), \ldots, (B_1^m, B_1'^m)\}$, *define the* ICR-collision *of* Z,

$$\mathscr{C}(Z) \coloneqq \bigcap_{(B,B')\in Z} \partial\mathcal{R}(B, B'), \tag{2.14}$$

where $\partial\mathcal{R}(B,B') \coloneqq \lim_{\sigma\downarrow 0} \partial\mathcal{R}^\sigma(B,B')$.

Figure 5 illustrates a collision between the boundary separating $(\{a\},\{a\})$ from $(\{a\},\{a,b\})$ and the boundary separating $(\{a\},\{a,b\})$ from $(\{a,b\},\{a,b\})$.

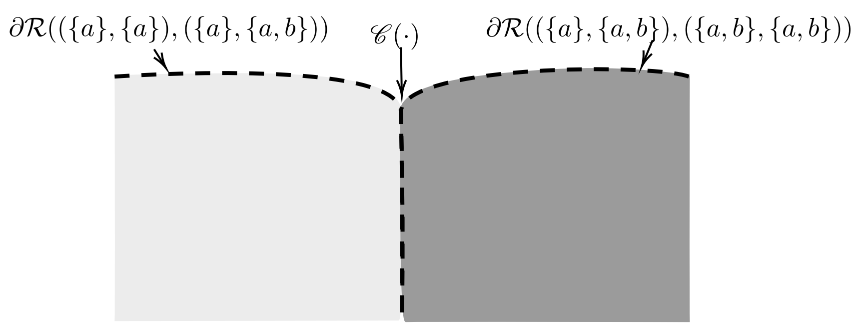

Figure 5: Illustration of collisions.

3 Illustrative Example: Binary-Action

In this section we again consider the special case where actions are binary, i.e. $A_i = \{a, b\}$, for every i. In this case, $U = \mathbb{R}^8$ and S is the unit sphere in $\mathbb{R}^8$,

$$S = \{u \in \mathbb{R}^8 : ||u||_2 = 1\}. \tag{3.1}$$

With additively separable noise rather than rotational noise, this case[8] was studied in Carlsson and Van Damme (1993). We now study binary action games using the tools introduced in this paper. We give a quick recap of our set-up: After a common state $y \in S$ is drawn with distribution $\nu_0 \in \Delta(S)$, every player i is privately informed of a game $s = e^{\sigma E_i} y$, where $E_i \sim \nu_i \in \Delta(\mathcal{E})$ gives rise to a random rotation of magnitude σ. Recall that we have assumed $\nu_1 = \nu_2$ with symmetric density. Given a signal $s \in S$, each player i can thus compute her interim correlated rationalizable actions $\text{ICR}_i^\sigma(s)$. The partition of rationalizable sets induced by the correspondence $\text{ICR}^\sigma : S \to \mathcal{A}_1 \times \mathcal{A}_2$ allows each player i to compute her beliefs on the rationalizable action sets of her opponent:

$$P_i^\sigma(B_{-i}|s) = \nu_0 \times \nu \left(\left\{ (y, E_i, E_{-i}) : B_{-i} = \text{ICR}_{-i}(e^{\sigma E_{-i}} y) \right\} \middle| e^{\sigma E_i} y = s \right). \tag{3.2}$$

We now provide a preview of the results and analysis presented in the main sections of this paper for the binary action case.

Indifference Constraints (IC) We will focus on the collision involving singletons $\{a\}, \{a\}$ and $\{b\}, \{b\}$. We show in Lemma 5.1 that the indifference constraint

$$g_i^{1,2,\sigma}(s) = P_i^\sigma(\{a\}|s)(s_i^{1,1} - s_i^{2,1}) + P_i^\sigma(\{b\}|s)(s_i^{1,2} - s_i^{2,2}) = 0, \tag{3.3}$$

must hold at the boundary where player i switches from $\{a\}$ to $\{b\}$. Condition (3.3) is linear in payoffs/signals when beliefs are held fixed, however beliefs also depend on payoffs.

Symmetry Constraints We use the symmetry of the noise distribution to constrain the limiting behavior of beliefs at points of collision. In Lemma 5.2 we show that at any point where boundaries collide, rationalizable sets are

[8] We discuss the connection between additively separable noise and rotational noise in Section 9.2.

symmetric about an axis of symmetry. In particular, we show that colliding beliefs are equal up to a permutation of the action labels. We illustrate this symmetry in Figure 6. In the figure we illustrate three rationalizable sets, for some positive $\sigma > 0$: 1) The light gray region on the left which corresponds to the region where both players rationalize $\{a\}$; 2) The white region in the middle, where player 1 rationalizes action $\{a\}$ but player 2 rationalizes $\{b\}$; 3) The dark gray region on the right, where both players rationalize $\{b\}$.

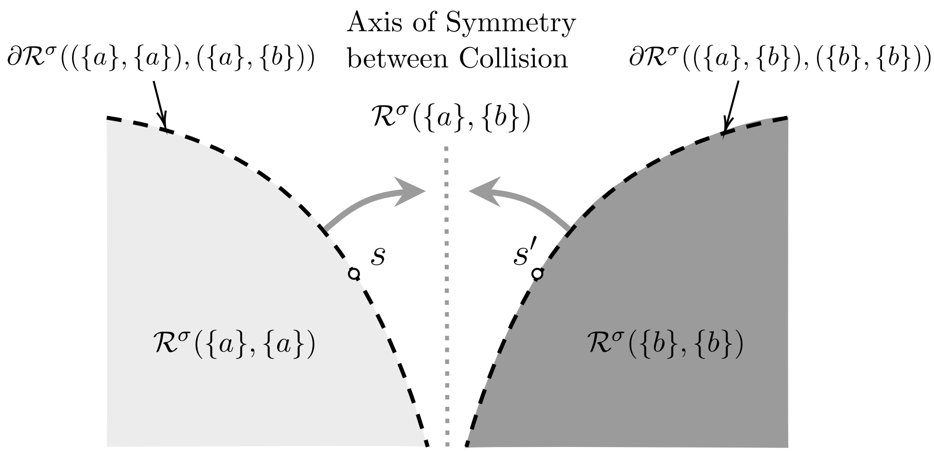

Figure 6: Symmetry of boundaries.

Points s and s' lie on their respective boundaries, where player 2 switches rationalizable actions at s and player 1 switches at s'. Both points are reflections of each other about the dashed axis of symmetry. We show that symmetry of the noise distribution implies that the beliefs at s and the beliefs at s' are equal up to a label permutation. In particular, we have that

$$P_2^\sigma(\{a\}|s) = P_1^\sigma(\{b\}|s'), \tag{3.4}$$

that is, the probability of the light gray and white regions, conditional on s, is equal to the probability of its mirror image: the dark gray and white region conditional on s'. Assuming both boundaries collide as $\sigma \downarrow 0$, thus making the white region vanish, imposes strong symmetry constraints on the limit beliefs at the collision. Assuming s and s' collide to a common point s^*, the symmetry constraint (3.5) becomes,

$$\lim_{\sigma\downarrow 0} P_2^\sigma(\{a\}|s^*) = \lim_{\sigma\downarrow 0} P_1^\sigma(\{b\}|s^*). \tag{3.5}$$

Risk-Dominance as Multilinear Form Combining the symmetry constraint (3.5) with the necessary indifference constraints (3.3), we obtain the binary action version of Proposition 5.4: Points on a collision are contained in the zeros of a multilinear form given by

$$\delta(s) = \sum_{\alpha=a,b} \prod_{i=1,2} (s_i^{a,\alpha} - s_i^{b,\alpha}) = 0. \tag{3.6}$$

In the case where two boundaries collide, the form is in fact bilinear: It is linear in the payoff coefficients of each of the colliding indifference constraints. By the symmetry constraints, beliefs of player i are implicitly described by the payoff-differences of the other player. Expression (3.6) is equivalent to both actions being risk-dominant[9]: The product of the deviation losses is exactly equal for both actions. In Proposition 5.4 we derive the generalization of this expression to the many-action environment and show that it is necessary for a point to lie on a collision. Carlsson and Van Damme (1993) have shown that condition (3.6) is also sufficient to characterize the selection in the region of S where payoffs are supermodular.

Aligned Incentives To obtain a sufficient condition for collisions we study the local stability properties of the zeros of the multilinear form (3.6). We consider the derivative at a point satisfying (3.6),

$$\frac{\partial \delta}{\partial (s_i^{a,a} - s_i^{b,a})} = (s_{-i}^{a,a} - s_{-i}^{b,a}). \tag{3.7}$$

We say that s is *critical* if it satisfies (3.6) and *incentives are aligned* if

$$\frac{\partial \delta}{\partial (s_i^{a,a} - s_i^{b,a})} \frac{\partial \delta}{\partial (s_{-i}^{a,a} - s_{-i}^{b,a})} > 0. \tag{3.8}$$

Aligned incentives thus requires that a perturbation in a common direction for both players (here the dominance region of (a, a)) has the same effect on their ICs. Hence s has *aligned incentives* on $\{(a, a), (b, b)\}$ if $(s_i^{a,\alpha} - s_i^{b,\alpha})$ and $(s_{-i}^{a,\alpha} - s_{-i}^{b,\alpha})$ have the same sign of every $\alpha \in \{a, b\}$. In Proposition 6.1, we show that aligned incentives and criticality are necessary and sufficient for points to lie on a collision.

[9] See for instance Harsanyi and Selten (1988).

When s is given by a "matching pennies" payoff function, incentives are not aligned:

	1	2
1	1,-1	-1,1
2	-1,1	1,-1

Table 1: Matching Pennies Game.

Then we have that $\frac{\partial \delta}{\partial (s_i^{a,a} - s_i^{b,a})} \frac{\partial \delta}{\partial (s_{-i}^{a,a} - s_{-i}^{b,a})} < 0$. We conclude that for all payoffs where incentives are misaligned, $(\{a\}, \{a\})$ does not collide with either $(\{a\}, B_{-i})$ for $B_{-i} \in \{\{b\}, \{a, b\}\}$. In Proposition 4.1 we show that the presence of risk dominant[10] best-response cycles are in fact enough for limit multiplicity.

For a coordination game, incentives are aligned and we can recover the result in Carlsson and Van Damme (1993).

	1	2
1	1,1	-1,-1
2	-1,-1	1,1

Table 2: Coordination Game.

In Theorem 7.1 we state our characterization of limit uniqueness for strictly supermodular games with strictly concave payoffs.

Limit Selection In the case of binary action, we can thus provide an explicit expression for ICR_i. For every player i, the dominance regions $D_i(a), D_i(b) \subseteq S$ are given by

$$\begin{aligned} D_i(a) &= \{s \in S : s_i^{a,\alpha} - s_i^{b,\alpha} > 0, \ \forall \ \alpha \in \{a, b\}\}, \\ D_i(b) &= \{s \in S : s_i^{a,\alpha} - s_i^{b,\alpha} < 0, \ \forall \ \alpha \in \{a, b\}\}. \end{aligned} \tag{3.9}$$

The selection when a dominance region is involved can be described easily: For every $s \in D_i(a_i)$ there is $\sigma > 0$ small enough so that

$$\text{ICR}_i^\sigma(s) = \{a_i\}, \ \text{ICR}_{-i}^\sigma(s) = \arg \max_{a_{-i} \in A_{-i}} s_{-i}^{a_{-i}, a_i}. \tag{3.10}$$

[10]We generalize risk-dominance to collections of action profiles in the next section.

Define the set of payoffs where no player has a dominant action,

$$Z := S \setminus \cup_i (D_i(a) \cup D_i(b)). \tag{3.11}$$

The selection when no dominance region is involved takes the following form: For every player i and $s \in \text{int}(Z)$,

$$\text{ICR}_i(s) = \begin{cases} \{a\}, & \text{if } \delta(s) > 0, \\ \{b\}, & \text{if } \delta(s) < 0, \\ A_i, & \text{otherwise.} \end{cases} \tag{3.12}$$

Moreover, ICR_i extends continuously to the closure $\overline{\text{int}(Z)} = Z$.

4 Risk-Dominance and Limit Multiplicity

In Subsection 4.2, we provide an upper-bound for the limit selection using risk-dominance. A pair of action sets (B_1, B_2) - one action set for each player - satisfies *risk-dominance* if the minimal probabilities that each player i needs to assign to their opponent $-i$ playing an action in B_{-i}, so that i's best-reply is contained in B_i, sum to at most one. We apply existing results from Kajii and Morris (1997) to show that the limit selection is always contained in a risk-dominant set that is closed under best-replies (BRC). In Subsection 4.3 we provide a sufficient condition for the limit selection to exhibit multiple actions for a player: Any risk-dominant BRC set consisting of a best-response cycle involving two or more actions for a player will be contained in the limit selection.

4.1 Example: Coordination on Surplus Splitting

We illustrate the results in this subsection by a simple example. We consider a coordination game with a surplus splitting component that gives rise to a best-response cycle. In this game, each player makes a production capacity choice (H, high and L, low) and a surplus splitting choice (a or b).

<table>
<tr><td></td><td></td><td colspan="2">H</td><td colspan="2">L</td></tr>
<tr><td></td><td></td><td>a</td><td>b</td><td>a</td><td>b</td></tr>
<tr><td rowspan="2">H</td><td>a</td><td>1,-1</td><td>-1,1</td><td colspan="2" rowspan="2">0</td></tr>
<tr><td>b</td><td>-1,1</td><td>1,-1</td></tr>
<tr><td rowspan="2">L</td><td>a</td><td colspan="2" rowspan="2">0</td><td>ϵ,-ϵ</td><td>-ϵ,ϵ</td></tr>
<tr><td>b</td><td>-ϵ,ϵ</td><td>ϵ,-ϵ</td></tr>
</table>

Table 3: Coordination Game with Surplus-Splitting Component.

Conditional on picking either H or L, players play a matching pennies game with payoffs ± 1 or $\pm\epsilon$. With their choice of H or L players coordinate on which of the two matching pennies games to play. Conditional on choosing H, best-replies are cyclic. When we consider the solution concept of ICR, the limit selection must always contain a pair of actions of the form $(\alpha, a), (\alpha, b)$, for $\alpha \in \{H, L\}$. If we allow players to make their most optimistic conjecture over the choice a, b of their opponent, choosing H has payoff 1 and choosing L has payoff ϵ. The reduced binary action game is a simple coordination game:

	H	L
H	1,1	0,0
L	0,0	ϵ, ϵ

Table 4: Coordination Game induced by optimistic conjectures over a, b.

For ϵ small enough, players don't need to be too certain that their opponent is playing H for them to also play H. We conclude that in the binary choice of H and L, action H *risk-dominates* L in the sense of Harsanyi and Selten (1988). If players receive a scaled version of the payoff matrix in Table 3, s, we would expect that for ϵ small, the limit selection is of the form

$$\mathrm{ICR}_i(s) = \{(H, a), (H, b)\}, \ \forall\ i = 1, 2. \tag{4.1}$$

In the section below we confirm this conjecture and show how to generalize this observation.

4.2 Risk-Dominance

A pair of action sets $(B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ is *best-reply closed* (BRC) at $s \in S$ if all best-replies to beliefs that are supported on B are themselves equal to B,

$$\bigcup_{p \in \Delta(B_{-i})} \mathrm{br}_i(p_i|s_i) = B_i, \ \forall\ i. \tag{4.2}$$

For every pair $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$, define the minimal probability player i needs to attach to B_{-i} so that she has a best-reply in B_i,

$$p_i^B(s) \coloneqq \min \left\{ p_i^B \in [0,1] : \begin{array}{c} \exists\ p \in \Delta(A_{-i}) \text{ s.t. } p(B_{-i}) \leq p_i^B \\ \mathrm{br}_i(p|s_i) \cap B_i \neq \varnothing \end{array} \right\}. \tag{4.3}$$

We generalize the concept of risk-dominance for binary action games from Harsanyi and Selten (1988) to BRC sets. A BRC set is *risk-dominant* if

$$p_1^B(s) + p_2^B(s) \leq 1. \tag{4.4}$$

The result below shows that all payoffs in S whose action sets can be partitioned into at least two BRC sets admit a risk-dominant BRC set:

Lemma 4.1. *Fix $s \in S$ and let $B = (B_1, B_2), \hat{B} = (\hat{B}_1, \hat{B}_2)$ be disjoint BRC sets satisfying $A_i = B_i \cup \hat{B}_i$, for every player i. Then $\{B, \hat{B}\}$ contains a risk-dominant BRC set.*

Lemma 4.1 follows from the following observation: If a BRC set does not satisfy risk-dominance, then its complement does. Indeed, if $B = (B_1, B_2)$ is a BRC set violating risk-dominance, then $p_1^B(s) + p_2^B(s) > 1$. But then the complement $\hat{B} = (\hat{B}_1, \hat{B}_2)$ must satisfy $p_i^{\hat{B}}(s) \leq 1 - p_i^B(s)$, which establishes risk-dominance of $\hat{B}$.

For every $s \in S$, let $\mathscr{R}(s)$ denote collection of risk-dominant BRC sets (if s has such a BRC set) and let it be equal to the full action set otherwise. The result below establishes an upper bound for the limit selection: The limit selection is always contained in the union of all risk-dominant BRC sets.

Lemma 4.2 (Upper Bound on Limit Selection)**.** *For every $s \in S$ and every player i,*

$$\mathrm{ICR}_i(s) \subseteq \bigcup_{(B_1,B_2) \in \mathscr{R}(s)} B_i. \tag{4.5}$$

The result is a consequence of the critical path result in Kajii and Morris (1997). Any action that is excluded from every risk dominant BRC set cannot be rationalizable on the set of signals where there is common (p_1, p_2)-certainty of players playing actions in a risk-dominant BRC set, for $p_1 + p_2 \leq 1$. The critical path result ensures that, as the noise vanishes, the set of signals satisfying common (p_1, p_2)-certainty of players playing actions in a risk-dominant BRC set have ex-ante probability one. Hence the result follows.

4.3 Limit Multiplicity

We now provide a sufficient condition for limit multiplicity. We show that if all risk-dominant BRC sets are contained in one large best-response cycle then they are all selected.

Best-Response Cycles For every $s \in S$ and any player i a reaction function $\beta_i \colon A_{-i} \to A_i$ satisfies for every $a_{-i} \in A_{-i}$,

$$\beta_i(a_{-i}) \in \arg\max_{a_i \in A_i} s_i^{a_i, a_{-i}}. \tag{4.6}$$

Note that generically, points in S admit a unique reaction function. A *best-response cycle* of $s \in S$ is an ordered list of distinct actions for each player, $c = (c_1, c_2)$, where for every player i, c_i takes the form $c_i = (a_i^1, \dots, a_i^{n_c})$ and satisfies for any $l < n_c$,

$$\beta_1(a_1^l) = a_2^l, \;\; \beta_2(a_2^l) = a_1^{l+1}, \tag{4.7}$$

with $a_2^{n_c}$ satisfying $\beta_2(a_2^{n_c}) = a_1^1$, for some reaction functions β_1, β_2. Call $n_c \in \mathbb{N}$ the length of cycle $c = (c_1, c_2)$. For every $s \in S$, let $C_i(s) \subseteq A_i$ denote the collection actions of player i that are contained in cycles of any length $n_c > 1$, and $C_i^*(s) \subseteq A_i$ the collection of actions contained in cycles of length $n_c = 1$. Note that if s admits a unique reaction function, then $C_i(s)$ and $C_i^*(s)$ are disjoint. The cyclic decomposition of such a game consists of the pair of disjoint action sets $\bar{C}(s) = (C_i(s), C_i^*(s))_{i=1,2}$.

Multiplicity Let $C(s) := (C_1(s), C_2(s))$ be the profile of cycles. Consider $p^{C(s)}(s) = (p_1^{C(s)}(s), p_2^{C(s)}(s)) \in [0, 1]^2$, the minimal probabilities that each player i needs to assign to $C_{-i}(s)$ so that her best-reply is in $C_i(s)$ as defined in (4.3). Say that s has a *risk-dominant cyclic component* if the cycle is the union of all risk-dominant BRC sets, i.e.

$$C_i(s) = \bigcup_{(B_1,B_2)\in\mathscr{R}(s)} B_i. \tag{4.8}$$

Proposition 4.1 below shows that a risk-dominant cyclic component is always contained in the selection.

Proposition 4.1 (Sufficient Condition for Limit Multiplicity)**.** *Let $s \in S$ have a risk-dominant cyclic component $C(s)$ then for every player i,*

$$C_i(s) \subseteq \mathrm{ICR}_i(s). \tag{4.9}$$

Proof. Since $\mathrm{ICR}_i(s) \neq \varnothing$, we conclude from Lemma 4.2 that $\mathrm{ICR}_i(s) \subseteq C_i(s)$. Since $\mathrm{ICR}(s)$ is itself a BRC set it cannot be smaller than $C(s)$. □

Coming back to the example game in Table 3: There are two symmetric BRC sets in this game: $B_i^H = \{(H,a),(H,b)\}$ and $B_i^L = \{(L,a),(L,b)\}$. Under the most optimistic conjecture, playing H gives payoffs of 1 and under the most optimistic conjecture, playing L gives payoffs ϵ. We then have that

$$p_i^{B^H} = \frac{\epsilon}{1+\epsilon}. \tag{4.10}$$

So if $\epsilon < 1$, $B^H = (B_1^H, B_2^H)$ becomes the risk-dominant cyclic component and so we have that the limit selection must contain B^H.

5 Properties of the Limit Selection

In this section we derive basic properties of the limit selection. In Subsection 5.1 we establish basic topological properties of the limit selection. We prove a basic upper-hemi continuity property of rationalizable sets and show that for every pair of action sets, there is an open region in S, where this pair is selected. In Subsection 5.2 we establish an algebraic property of collisions. We show that every collision is contained in the zero-set of a symmetric multilinear form. The coefficients of the multilinear form are derived from the indifference conditions that need to hold on each of the colliding boundaries. They are pinned down by the symmetry of the noise distribution.

5.1 Topological Properties

Proposition 5.1 below shows that everything is selected in some open set. The proof is constructive: For any pair of action sets we construct a game in S with a risk-dominant best-response cycle involving all actions in that set. The result then follows from an application of Proposition 4.1.

Proposition 5.1 (All Rationalizable Sets)**.** *For every pair of action sets* $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ *so that* $|B_1| = |B_2|$*, there is an open*[11] *set* $O \subseteq S$ *so that*

$$O \subseteq \lim_{\sigma \downarrow 0} \mathcal{R}^\sigma(B). \tag{5.1}$$

Proposition 5.2 below establishes upper-hemi continuity of ICR^σ. Upper-hemi continuity was established in Dekel et al. (2007) with respect to the product topology on hierarchies of beliefs. In our case, the type space is given by the sphere S and the property follows from continuity of players' beliefs as a function of their signal with respect to the standard Euclidean topology.

Proposition 5.2 (Upper-Hemi Continuity)**.** *For every* $\sigma > 0$ *and any pair* $(B, B') \in (\mathcal{A}_1 \times \mathcal{A}_2)^2$ *so that* $\partial\mathcal{R}^\sigma(B, B') \neq \varnothing$,

$$\partial\mathcal{R}^\sigma(B, B') = \partial\mathcal{R}^\sigma(B \cup B', B'). \tag{5.2}$$

Define the collection of nested action set profiles,

$$\mathscr{A} := \{(B, B') \in (\mathcal{A}_1 \times \mathcal{A}_2)^2 : \forall\, i, B'_i \subseteq B_i\}. \tag{5.3}$$

We now study the structure of collisions. By Proposition 5.2 it is without loss of generality to restrict attention to collisions of action set-pairs in $\mathscr{A}$.

5.2 Algebraic Properties

Signals on the boundary between two rationalizable sets make a player indifferent between two actions. These boundaries are defined by indifference constraints. We will now examine how these boundaries interact as the noise vanishes, focusing on the limits of the indifference constraints at each colliding boundary.

[11] A subset $O \subseteq S$ is open if it is the intersection of an open set in U and S.

First, we will establish an expression for these indifference constraints. Then, we'll use the symmetry of the noise distribution to determine the beliefs that players hold at the points where the boundaries collide. We show that the rationalizable sets with colliding boundaries are mirror images of each other. As a result, the beliefs associated with one indifference constraint can be linked to those of the other colliding boundaries via a set of symmetry constraints.

For every $s \in S$, every $\sigma > 0$, every player i and every action set $B_{-i} \in \mathcal{A}_{-i}$, define the beliefs at s,

$$P_i^\sigma(B_{-i}|s) := \nu \times \nu_0 \left(\left\{ (E_i, E_{-i}, y) : B_{-i} = \mathrm{ICR}_{-i}^\sigma(e^{\sigma E_{-i}} y) \right\} \middle| e^{\sigma E_i} y = s \right). \tag{5.4}$$

$P_i^\sigma(B_{-i}|s)$ represents the probability that player i attaches to her opponent rationalizing B_{-i} after i has observed signal s. We will start by deriving the Indifference Constraints that must hold at each point on a boundary.

Indifference Constraints Fix $\sigma > 0$. For any player i and any pair of actions $a_i, a_i' \in A_i$, define the expected deviation loss of switching from a_i to a_i' under the most optimistic selection rule,

$$g_i^{a_i, a_i', \sigma}(s) := \max_{\kappa_{-i} \in \Sigma_{-i}} \sum_{B_{-i} \in \mathcal{A}_{-i}} P_i^\sigma(B_{-i}|s) \sum_{a_{-i} \in B_{-i}} \kappa_{-i}(a_{-i}|B_{-i})(s_i^{a_i, a_{-i}} - s_i^{a_i', a_{-i}}), \tag{5.5}$$

where $\Sigma_{-i} := \{\kappa \colon \mathcal{A}_{-i} \to \Delta(A_{-i}) : \forall\ a_{-i} \in A_{-i},\ \forall\ B_{-i} \in \mathcal{A}_{-i}, \kappa(a_{-i}|B_{-i}) > 0 \implies a_{-i} \in B_{-i}\}$ is the set of random selections. For every player i and action pair $(a_i, a_i') \in A_i \times A_i$ define the set of signals where this deviation loss is zero

$$\mathcal{G}_i^\sigma(a_i, a_i') := \left\{ s \in S : g_i^{a_i, a_i', \sigma}(s) = 0 \right\}. \tag{5.6}$$

The result below shows that for any nested pair of action sets, the boundary between the two associated rationalizable sets satisfies indifference constraints given by (5.6).

Lemma 5.1 (Indifference Constraints)**.** *For every $(B, B') \in \mathscr{A}$, every $\sigma > 0$ and every $s \in \partial \mathcal{R}^\sigma(B, B')$, there exist $(a_i, a_i') \in B_i \times B_i'$ for every i so that*

$$s \in \bigcap_{i=1,2} \mathcal{G}_i^\sigma(a_i, a_i'). \tag{5.7}$$

The indifference constraints in (5.6) are linear in payoffs when beliefs are fixed. However, beliefs are also a function of payoffs. We exploit the symmetry properties of the noise distribution to constrain the behavior of limit beliefs and establish the multilinearity of collisions. We start by writing every collision as an *IC-collision*, a collision of indifference constraints for action pairs. For any $(a_i, a_i') \in A_i \times A_i$, define the limit set

$$\mathcal{G}_i(a_i, a_i') \coloneqq \lim_{\sigma \downarrow 0} \mathcal{G}_i^\sigma(a_i, a_i'). \tag{5.8}$$

For every player i, let $C_i = \{(a_i^1, a_i'^1), \dots, (a_i^{m_i}, a_i'^{m_i})\} \subseteq A_i \times A_i$ be a collection of action pairs. Define the $C = (C_1, C_2)$*-IC-collision*

$$\mathcal{C}(C) \coloneqq \bigcap_{i=1,2} \bigcap_{n=1}^{m_i} \mathcal{G}_i(a_i^n, a_i'^n). \tag{5.9}$$

It follows from Lemma 5.1 that every collision is contained in an IC-collision. Indeed, for every Z there exists

$$C_{Z,i} \subseteq \bigcup_{(B,B') \in Z} (B_i \times B_i'). \tag{5.10}$$

so that $\mathscr{C}(Z) \subseteq \mathcal{C}(C_{Z,1}, C_{Z,2})$.

Corollary 5.1. *For every $Z \subseteq \mathscr{A}$ there exists $C = (C_1 \subseteq A_1 \times A_1, C_2 \subseteq A_2 \times A_2)$ so that*

$$\mathscr{C}(Z) \subseteq \mathcal{C}(C). \tag{5.11}$$

We will now study IC-collisions.

Symmetry of Colliding Beliefs Fix a collection of action pairs $C_i = \{(a_i^1, a_i'^1), \dots, (a_i^{m_i}, a_i'^{m_i})\} \subseteq A_i \times A_i$, for every player i. We show that symmetry of the noise distribution (i.e. condition (5.18)) implies that for any $n, n' \leq m_i$, the beliefs required to make the nth IC hold can be obtained by permuting the beliefs required to make the n'th IC hold. We prove this property by induction on the rounds of deletion of ICR. Define the joint beliefs $\bar{P}^\sigma(\cdot|s) \in \Delta(\mathcal{A}_1 \times \mathcal{A}_2)$, which for every pair of action sets $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$ are defined as

$$\bar{P}^\sigma(B|s) \coloneqq \nu \times \nu_0 \left(\{(E, E', y) : B = \mathrm{ICR}^\sigma(e^{\sigma E'} y)\} | e^{\sigma E} y = s \right). \tag{5.12}$$

By symmetry of the noise distribution across players, we have that for every player i and action set $B_{-i} \in \mathcal{A}_{-i}$,

$$P_i^\sigma(B_{-i}|s) = \sum_{\tilde{B} \in \mathcal{A}_1 \times \mathcal{A}_2 : \tilde{B}_{-i} = B_{-i}} \bar{P}^\sigma(\tilde{B}|s). \tag{5.13}$$

Let $\mathrm{P}^{2n} \subseteq \{-1, 0, 1\}^{2n \times 2n}$ denote the collection of signed permutation matrices. An *invariance* is a pair (η, X), where $\eta\colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathcal{A}_1 \times \mathcal{A}_2$ is a permutation of the action set labels and $X \in \mathrm{P}^{2n}$ so that for all $s \in S$ and every $\sigma > 0$,

$$\eta(\mathrm{ICR}^\sigma(s)) = \mathrm{ICR}^\sigma(Xs). \tag{5.14}$$

Condition (5.14) means the following: A rationalizable set $\mathcal{R}^\sigma(B) \subseteq S$ can be transformed into another rationalizable set $\mathcal{R}^\sigma(\eta(B)) \subseteq S$ by applying a permutation matrix X to all elements in $\mathcal{R}^\sigma(B)$. This means that rationalizable sets which are mapped to each-other by η are in fact mirror images of each-other.

A pair (η, X) preserves best-replies if for every $\tilde{s} \in S$ and every $p \in \Delta(\mathcal{A}_1 \times \mathcal{A}_2)$

$$\eta(\widetilde{\mathrm{BR}}(p|\tilde{s})) = \widetilde{\mathrm{BR}}(p \circ \eta | X\tilde{s}), \tag{5.15}$$

where for every i,

$$\widetilde{\mathrm{BR}}_i(p|s) = \bigcup_{\substack{\kappa_{-i}\colon \mathcal{A}_1 \times \mathcal{A}_2 \to \Delta(A_{-i}) \text{ s.t.} \\ \kappa_{-i}(B_{-i}|B_1, B_2) = 1, \ \forall\ (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2}} \mathrm{br}_i(p \circ \kappa_{-i}|s_i), \tag{5.16}$$

and $p \circ \kappa_{-i}(a_{-i}) \coloneqq \sum_{(B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2} \kappa_{-i}(a_{-i}|B_1, B_2) p(B_1, B_2)$. Condition (5.15) is analogous to (5.14): It requires that the set of beliefs with a given best-reply can be transformed into the set of beliefs with another best-reply, by permuting payoffs and the support of each belief. Below, we show that preserving best-replies is a sufficient condition for a pair (η, X) to also be an invariance. Moreover, we establish a symmetry property of beliefs associated to each invariance, given by expression (5.17).

Lemma 5.2 (Symmetry of Colliding Beliefs)**.** *Every (η, X) that preserves best-replies is an invariance. Moreover, for every $\sigma > 0$, every $\hat{s} \in S$ and every $B \in \mathcal{A}_1 \times \mathcal{A}_2$,*

$$\bar{P}^\sigma(B|\hat{s}) = \bar{P}^\sigma(\eta(B)|X\hat{s}). \tag{5.17}$$

The result in Lemma 5.2 follows from symmetry of the noise distribution. Indeed, if (η, X) is an invariance, then for any $\sigma > 0$, players' beliefs at s and at Xs will be symmetric up to the permutation of action labels induced by η:

$$\begin{aligned}\bar{P}^{\sigma}(B|s) &= \nu\left(\left\{(E,E') : B = \mathrm{ICR}^{\sigma}(e^{\sigma(E'-E)}s)\right\}\right)\\ &= \nu\left(\left\{(E,E') : \eta(B) = \mathrm{ICR}^{\sigma}(Xe^{\sigma(E'-E)}s)\right\}\right)\\ &= \nu\left(\left\{(E,E') : \eta(B) = \mathrm{ICR}^{\sigma}(Xe^{\sigma(E'-E)}X^{\top}Xs)\right\}\right)\\ &= \nu\left(\left\{(E,E') : \eta(B) = \mathrm{ICR}^{\sigma}(e^{\sigma(E'-E)}Xs)\right\}\right)\\ &= \bar{P}^{\sigma}(\eta(B)|Xs).\end{aligned} \tag{5.18}$$

The first line follows from the definition of $\bar{P}^{\sigma}$ in (5.13). The second line, follows from the fact that (η, X) is an invariance. The third line uses the fact that X is an orthogonal matrix and so $X^{\top}X$ is the identity. Finally, the fourth line uses a property of the matrix exponential: $Xe^{\sigma(E'-E)}X^{\top} = e^{\sigma X(E'-E)X^{\top}}$ and symmetry of the noise distribution. When (η, X) preserves best-replies we show inductively on the rounds of elimination of ICR, that (η, X) is an invariance and so (5.18) holds.

For every invariance (η, X) an *invariant point* is $s^* \in S$ so that,

$$Xs^* = s^*. \tag{5.19}$$

Fix $Z \subseteq \mathscr{A}$. Say that a permutation $\eta\colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathcal{A}_1 \times \mathcal{A}_2$ *acts on* Z if for every $(B, B') \in Z$, we have that

$$(\eta(B), \eta(B')) \in Z. \tag{5.20}$$

Say that a collision Z is *full* at $s \in \mathscr{C}(Z)$ if for every permutation η that acts on Z there exists a permutation matrix $X \in \mathrm{P}^{2n}$ so that (η, X) is an invariance with invariant point s.

The symmetry properties of beliefs established in Lemma 5.2 only have bite for collisions if collisions consisted of invariant points, or, if collisions are full. Lemma 5.3 establishes that fact.

Lemma 5.3 (Every Collision is full)**.** *For every $Z \subseteq \mathscr{A}$ and $s \in \mathscr{C}(Z)$, the collision Z is full at s.*

Combining Lemmas 5.2 and 5.3 establishes that any label permutation η that acts on Z defines an invariance whose invariant points contain the collision of Z.

Proposition 5.3 (Symmetry of Colliding Beliefs). *For every collision $\mathscr{C}(Z)$, every $s \in \mathscr{C}(Z)$ and every permutation η that acts on Z, there exists a permutation matrix $X \in \mathrm{P}^{2n}$ so that (η, X) is an invariance with invariant point s.*

Proof. This is an immediate consequence of Lemmas 5.2 and 5.3. □

The symmetry of colliding beliefs established in Proposition 5.3 implies that the sum of probabilities attached to any action set pair by points on the colliding boundaries must sum to one: In particular, we conclude that for every collision $\mathscr{C}(Z)$, every $s \in \mathscr{C}(Z)$, every $\sigma > 0$ and every $(B, B') \in Z$, there exist $s^{\sigma}_{B,B'} \in \partial\mathcal{R}^{\sigma}(B, B')$ so that

(i) $||s - s^{\sigma}_{B,B'}||_2 < \sigma$,

(ii) $||s^{\sigma}_{\hat{B},\hat{B}'} - s^{\sigma}_{B,B'}||_2 < \sigma$, $\forall\ (\hat{B}, \hat{B}') \in Z$, and

(iii) by varying the permutation η we also have that

$$\lim_{\sigma \downarrow 0} \sum_{(\hat{B},\hat{B}') \in Z} \bar{P}^{\sigma}(B|s^{\sigma}_{\hat{B},\hat{B}'}) = 1. \tag{5.21}$$

Property (iii) is the generalization of the symmetry property (3.5) in the binary action case.

Multilinearity We are now ready to state the main result of this subsection. It states that every collision is in the zero-set of a symmetric multilinear form. For every $a_i, a'_i \in A_i, a_{-i} \in A_{-i}$, write the payoff-difference as

$$d^{a_i,a'_i,a_{-i}}_{s_i} := s^{a_i,a_{-i}}_i - s^{a'_i,a_{-i}}_i. \tag{5.22}$$

We prove Lemma 5.4 by exploiting the symmetry property of the noise distribution and the resulting relation (5.21). This relation across boundaries pins down the coefficients of a multilinear form, which we specify in the Lemma below.

Lemma 5.4 (Multilinearity of IC-Collision)**.** *For every collection of action pairs* $C_i = \{(a_i^1, a_i'^1), \ldots, (a_i^{m_i}, a_i'^{m_i})\}$*, and for every player* i*, there exists a multilinear form* δ_C *so that*

$$\mathcal{C}(C) = \{s \in S : \delta_C(d_s^1, \ldots, d_s^K) = 0\}, \tag{5.23}$$

where for all $s \in S$*,*

$$\delta_C(d_s^1, \ldots, d_s^K) = \sum_{a \in A_1 \times A_2} \chi_{a,C} \prod_{i=1,2} \prod_{k=1}^{m_i} d_{s_i}^{a_i^k, a_i'^k, a_{-i}}, \tag{5.24}$$

for some family of binary coefficients $(\chi_{a,C})_a \in \{0,1\}^{A_1 \times A_2}$ *and for every* $k \le m_i$*,*

$$d_s^k := \left(d_{s_i}^{a_i^k, a_i'^k, a_{-i}}\right)_{a_{-i} \in A_{-i}}. \tag{5.25}$$

We thus obtain a necessary condition for points to lie on a collision:

Proposition 5.4 (Multilinearity)**.** *For every* $Z \subseteq \mathscr{A}$ *there exist a corresponding set* $C = (C_1, C_2)$*, of the form* $C_i = \{(a_i^1, a_i'^1), \ldots, (a_i^{m_i}, a_i'^{m_i})\}$ *for every* i*, and a family of binary coefficients* $(\chi_{a,C})_a \in \{0,1\}^{A_1 \times A_2}$*, so that*

$$\mathscr{C}(Z) \subseteq \{s \in S : \delta_C(d_s^1, \ldots, d_s^K) = 0\}, \tag{5.26}$$

where δ_C *is given by (5.24).*

Proof. This follows immediately from Lemma 5.4 and Corollary 5.1. □

Proposition 5.4 thus provides the generalization of the bilinear form (3.6) derived in the binary action case.

6 Aligned Incentives and Limit Uniqueness

Fix a collection of action pairs for each player $C = (\{(a_i^1, a_i'^1), \ldots, (a_i^{m_i}, a_i'^{m_i})\})_{i=1,2}$. Say that $\mathcal{C}(C)$ is a *singleton collision* if there exists $\epsilon > 0$ so that limit uniqueness holds on

$$\{s \in S : \exists\ \tilde{s} \in \mathcal{C}(C) \text{ s.t. } ||s - \tilde{s}||_2 < \epsilon\}. \tag{6.1}$$

In this section, we characterize limit uniqueness based on the local properties of the zero set of the multi-linear form discussed earlier. We demonstrate that a point on a collision is in a region of limit uniqueness if and only if it lies in the zero set of the corresponding multi-linear form and a condition we call "aligned incentives" is satisfied.

The aligned incentives condition is checked by perturbing the zero set of the multi-linear form in various directions. Each perturbation breaks the ties associated with the collision's indifference constraints, favoring a specific action. Additionally, when we perturb player i's signal in a way that favors an action for their opponent, it should increase the probability that player i assigns to that action. This shift in beliefs may then break a tie in player i's indifference constraints, leading to a preferred action for player i. However, this change also affects player $-i$'s beliefs after the initial perturbation, creating a ripple effect.

We establish that for limit uniqueness to hold, any perturbation must not create a cycle of tie breaks involving more than one action per player. This condition is both necessary and sufficient for limit uniqueness.

Aligned Incentives Fix a player i and let $(a_i, a_i') \in C_i$ and $a_{-i} \in A_{-i}$. Define the derivative

$$\gamma^{a_i,a_i',a_{-i}}(d_s) \coloneqq \frac{\partial \delta_C}{\partial d_{s_i}^{a_i,a_i',a_{-i}}}(d_s). \tag{6.2}$$

Define the sets of actions

$$I^{a_i,a_i'}(d_s) \coloneqq \left\{a_{-i} \in A_{-i} : \gamma^{a_i,a_i',a_{-i}}(d_s) > 0\right\}, \tag{6.3}$$

$I^{a_i,a_i'}(d_s)$ corresponds to the set of actions of player $-i$ so that an increase in i's payoff increment breaks the tie in favor of action a_i. In the binary action example from Section 3, we perturbed the zero-set of the bilinear form in equation (3.6) in the region of matching pennies games. When we perturbed the game in the direction where playing action (a, a) is dominant, we also raise the probability that the player who prefers to mismatch the other player's action assigns to a, thus breaking the tie in favor of action b. This occurs because the matching pennies game has a cycle of best responses, leading to a recurring pattern of tie breaks. To generalize this observation, let

$$G_{C,i} \coloneqq \bigcup_m \{(a_i^1, \dots, a_i^m) : (a_i^{k \mod m}, a_i^{(k+1) \mod m}) \in C,\ \forall\, k < m\}. \tag{6.4}$$

A sequence of cyclic perturbations in C consists of distinct actions for each player $(a_1^1, \dots, a_1^m) \in G_{C,1}$, $(a_2^1, \dots, a_2^m) \in G_{C,2}$ for $m \geq 2$, and a player i so that for all $k < m$,

$$a_{-i}^{k+1} \in I^{a_i^{k+1}, a_i^k}, \; a_i^k \in I^{a_{-i}^{k+1}, a_{-i}^k} \tag{6.5}$$

and

$$a_{-i}^1 \in I^{a_i^1, a_i^m}, \; a_i^m \in I^{a_{-i}^1, a_{-i}^m}. \tag{6.6}$$

A sequence of cyclic perturbations in C, $((a_1^1, \dots, a_1^m), (a_2^1, \dots, a_2^m)) \in G_{C,1} \times G_{C,2}$, is *un-dominated* if for every player i, every $k_i, k_{-i} \leq m$ and every action $a_i \in A_i$ so that $(a_i, a_i^{k_i}) \in C_i$, and

$$a_{-i}^{k_{-i}} \in I^{a_i, a_i^{k_i}}, \tag{6.7}$$

there a sequence of actions for every player $(\tilde{a}_1^1, \dots, \tilde{a}_1^h)$, $(\tilde{a}_2^1, \dots, \tilde{a}_2^h)$ so that

(i) for every $l < h$, $(\tilde{a}_1^l, \tilde{a}_1^{l+1}) \in C_1$, $(\tilde{a}_2^l, \tilde{a}_2^{l+1}) \in C_2$, $\tilde{a}_{-i}^1 = a_{-i}^{k_{-i}}$,

(ii) for every $l < h$, $\tilde{a}_{-i}^l \in I^{\tilde{a}_i^{l+1}, \tilde{a}_i^l}$, $\tilde{a}_i^{l+1} \in I^{\tilde{a}_{-i}^{l+1}, \tilde{a}_{-i}^l}$, and

(iii) there exists $k_i^* \leq m$ so that $\tilde{a}_i^h = a_i^{k_i^*}$.

A sequence of cyclic perturbations in C is thus un-dominated if every tie break in favor of an action outside of the cyclic sequence can be extended to break a tie in favor of some other action in the cyclic sequence.

Say that *incentives are aligned on C* if there is no un-dominated sequence of cyclic perturbations in C.

Lemma 6.1 (AI $\iff$ Singleton Collisions)**.** *Let $C = (C_1 \subseteq A_1 \times A_1, C_2 \subseteq A_2 \times A_2)$. $\mathcal{C}(C)$ is a singleton collision if and only if incentives are aligned on C.*

We conclude with our main result of this section, which provides a characterization of limit uniqueness in terms of aligned incentives.

Proposition 6.1 (Characterization of Limit Uniqueness)**.** *$O \subseteq S$ satisfies limit uniqueness if and only if for every collision C, every $s \in \mathcal{C}(C) \cap O$, incentives are aligned on C.*

Proof. This result follows from Lemma 6.1. □

Aligned incentives allows us to characterize limit uniqueness without any restrictions on the game. However, this condition is hard to interpret and check in practice. In order to obtain more interpretable results, we will restrict attention to a subclass of payoffs: Concave supermodular games.

7 Limit Uniqueness for Supermodular Games

Consider the special case where action sets are an interval of the form $A_i = \{1, \dots, N_i\}$, for every player i, where $N_i \in \mathbb{N}$. Using the total order on integers, let $\preceq$ denote the product order on $A_1 \times A_2$. Payoffs $s \in S$ are *strictly supermodular* if for all $(a_i, a_{-i}), (a'_i, a'_{-i}) \in A_1 \times A_2$ so that $(a_i, a_{-i}) \preceq (a'_i, a'_{-i})$ but not $(a'_i, a'_{-i}) \preceq (a_i, a_{-i})$,

$$d_{s_i}^{a'_i, a_i, a_{-i}} < d_{s_i}^{a'_i, a_i, a'_{-i}}, \ \forall\, i = 1, 2. \tag{7.1}$$

That is, switching to higher actions becomes more profitable when the other player also plays higher actions. Say that $s \in S$ is concave if for every player i and every $a_{-i} \in A_{-i}$, the mapping

$$a_i \mapsto s_i^{a_i, a_{-i}}, \tag{7.2}$$

is strictly concave.[12] Let $S^* \subseteq S$ denote the collection of strictly concave, strictly supermodular payoffs in S.

Generalized Ordinal Potentials Let $C_i = \{(a_i^1, a_i'^1), \dots, (a_i^{m_i}, a_i'^{m_i})\}$ be a collection of action pairs for player $i = 1, 2$. Say that s admits an *generalized ordinal potential on* (C_1, C_2) if there exists a function $F_s \colon A_1 \times A_2 \to \mathbb{R}$ so that for every player i, every action pair $(a_i, a'_i) \in C_i$, and any $a_{-i} \in A_{-i}$,

$$d_{s_i}^{a_i, a'_i, a_{-i}} > 0 \implies F_s(a_i, a_{-i}) - F_s(a'_i, a_{-i}) > 0. \tag{7.3}$$

Say that $s \in S$ has a *better-response cycle* on (C_1, C_2) if there are sequences $(a_1^1, \dots, a_1^m) \in G_{C,1}$ and $(a_2^1, \dots, a_2^m) \in G_{C,2}$ and a player i so that for every $n < m$,

$$d_{s_i}^{a_i^{l+1}, a_i^l, a_{-i}^l} > 0, \ d_{s_{-i}}^{a_{-i}^{l+1}, a_{-i}^l, a_i^l} > 0, \tag{7.4}$$

and

$$d_{s_i}^{a_i^1, a_i^m, a_{-i}^m} > 0 \text{ and } d_{s_{-i}}^{a_{-i}^1, a_{-i}^m, a_i^m} > 0. \tag{7.5}$$

The following result is from Monderer and Shapley (1996):

[12] We define concavity on finitely supported functions on A_i in the usual way: For any interval $a_i \preceq a'_i \preceq a''_i$ with distinct actions $a_i, a'_i, a''_i \in A_i$, we have that $s_i^{a'_i, a_{-i}} > (s_i^{a_i, a_{-i}} + s_i^{a''_i, a_{-i}})/2$.

Proposition 7.1 (Monderer and Shapley (1996)). *A game s admits a generalized ordinal potential on (C_1, C_2) if and only if it admits no better-response cycle on (C_1, C_2).*

We use this characterization to relate the generalized ordinal potential property to aligned incentives for strictly concave, strictly supermodular games.

Lemma 7.1 (GOP $\implies$ AI). *Let s be strictly supermodular and concave. If s admits a generalized ordinal potential on (C_1, C_2) then incentives are aligned on (C_1, C_2).*

7.1 Limit Uniqueness

We conclude with our main result of this subsection, which is a characterization of limit uniqueness for strictly concave, strictly supermodular games:

Theorem 7.1 (Limit Uniqueness). *Let $O \subseteq S^*$ be open. Limit uniqueness holds on O if and only if every $s \in O$ admits a generalized, ordinal potential on every risk-dominant BRC set.*

We prove Theorem 7.1 as follows: Suppose every $s \in O$ admits a generalized, ordinal potential on all risk-dominant BRC sets. Then by Lemma 7.1 incentives are aligned. Consider the shortest path from any point in O to a dominance region. We show that all games along this path also admit a generalized ordinal potential on all risk-dominant BRC sets and so incentives are aligned everywhere on the path. By Proposition 6.1 every collision along the path is a singleton collision. Since the starting point in O is arbitrary, we conclude that limit uniqueness holds on O.

For the converse, suppose that limit uniqueness holds on O and let $s \in O$ fail to have a generalized ordinal potential on some risk-dominant BRC set $B = (B_1, B_2)$. By Proposition 7.1, the failure of the generalized potential property is equivalent to the existence of a better-response cycle in some risk-dominant BRC set. We show that there is a path from s to some $\bar{s} \in S$, where $\bar{s}$ has a risk-dominant best-response cycle. By Proposition 4.1 we conclude that $\bar{s}$ satisfies limit multiplicity. We then show that every collision on the path fails aligned incentives. By Proposition 6.1 all points along the path have multiplicity. This last part of the argument is illustrated with a three-action cycle $\{a, b, c\}$ in Figure 7 below. The figure depicts the path from $\bar{s}$ and s and three hypothetical boundaries separating three hypothetical rationalizable sets of a player for some small $\sigma > 0$. Suppose that close to $\bar{s}$, the set of rationalizable actions of player i is given by $\{a, b, c\}$ because these actions are involved in a best-response cycle. The first boundary in Figure 7 eliminates c (shaded in gray) for player i. The next boundary eliminates b on top of c (shaded in gray). We will argue that this is impossible when all games along the path have a better-response cycle involving actions $\{a, b, c\}$ for player i. Indeed: at the first boundary, player i is indifferent between, say actions a and c. A perturbation that increases payoffs at a for player i breaks ties among actions in $\{a, b, c\}$ in a cyclic fashion, eventually favoring all actions in the better-response cycle. Hence c must also be rationalizable in the middle region. Similarly, b cannot be eliminated after the second boundary.

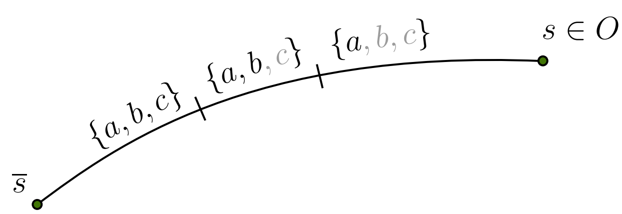

Figure 7: Path between game with best-response cycle and game with better response cycle.

8 Examples

In this section we show that the failure of admitting a generalized ordinal potential, i.e. the existence of better response cycles, and thus the failure of limit uniqueness can result from economically interpretable conditions. We illustrate in a four-action game that a small asymmetry in payoffs is enough to generate a better response cycle. We show that this asymmetry can be described in smooth, continuum action games in terms of the third-order derivative of a player's payoff function. We then use this condition to argue in an example that budget balanced policy interventions, i.e. redistributing payoffs, in games where limit uniqueness holds preserves the limit uniqueness property.

8.1 Example: Investment with Asymmetric Returns

We consider two agents faced with an investment choice. Agent i's investment choice is denoted a_i and is taken from the set $A_i := \{a = 1, b = 2, c = 3, d = 4\}$ with $a \prec b \prec c \prec d$. The return to a pair of investment choices $a = (a_1, a_2)$ is given by a strictly supermodular and strictly concave payoff matrix depicted in Table 5. The game has three pure strategy Nash equilibria: $\{(a, a), (c, b), (d, d)\}$ indicated in bold. The game contains a better-response cycle, highlighted in gray:

$$(b, a) \to (b, b) \to (d, b) \to (d, c) \to (c, c) \to (c, a) \to (b, a). \tag{8.1}$$

	a		b		c		d	
a	**9**	**7**	4	4	3	-1	-1	-6
b	8	6	7	7	10	3	7	-1
c	4	8	**10**	**10**	16	7	14	4
d	0	4	8	10	15	15	**15**	**19**

Table 5: Concave Supermodular Game with Better Response Cycle: $(0, 2), (0, 1), (3, 1), (3, 5), (4, 5), (4, 2), (0, 2)$.

The cycle arises because of an asymmetry in the payoffs: For the column player, under-investment (investing below the best-response) is better than over-investment (investing above the best-response) while for the row player over-investment is better than under-investment. Note that $(\{a, b, c, d\}, \{a, b, c, d\})$ is a BRC set.

By Monderer and Shapley (1996) we conclude that the game does not admit a generalized ordinal potential $(\{(b, c), (c, d), (b, d)\}, \{(a, b), (a, c), (b, c)\})$ and so limit uniqueness fails.

8.2 Example: Large Asymmetric Investment Games

We now consider a parametric class of payoff functions and show that almost every game in this class has a better response cycle arbitrarily close to every interior equilibrium.

Consider the following parametrized payoff function for action choices $(a_1, a_2) \in [0,1]^2$,

$$g_i(a_1, a_2) = \alpha_i a_i - a_i^{\alpha_i} a_{-i}^{1-\alpha_i}, \tag{8.2}$$

where $(\alpha_1, \alpha_2) \in (1, \infty)^2$ satisfy $\alpha_1 + \alpha_2 = 4$. There is a continuum of pure strategy, symmetric Nash equilibria given by $\{(a^*, a^*) : a^* \in [0,1]\}$. Player i's payoff from playing (a^*, a^*) is given by

$$g_i(a^*, a^*) = \frac{\alpha_i - 1}{2} 2a^*, \tag{8.3}$$

and similarly, player $-i$'s payoff at (a^*, a^*) can be expressed in terms of α_i as follows

$$g_{-i}(a^*, a^*) = \frac{3 - \alpha_i}{2} 2a^*. \tag{8.4}$$

In equilibrium, the total amount of investment is distributed among the players with shares determined by α_i: $g_i(a^*, a^*) + g_{-i}(a^*, a^*) = \frac{\alpha_i - 1 + 3 - \alpha_i}{2} 2a^* = 2a^*$. For any player i, and any fixed action a_{-i}, $\alpha_i \in (1,2)$ means that i's payoffs, as a function of her own action, are skewed to the right, while $\alpha_i > 2$ means that payoffs are skewed to the left. Figure 8 below illustrates this. The left most graph shows the shape of the function $a_i \mapsto g_i(a_i, a_{-i})$ when $\alpha_i > 2$. The middle graph considers the case where $\alpha_i = 2$ and the right most graph considers the case where $\alpha_i \in (1,2)$.

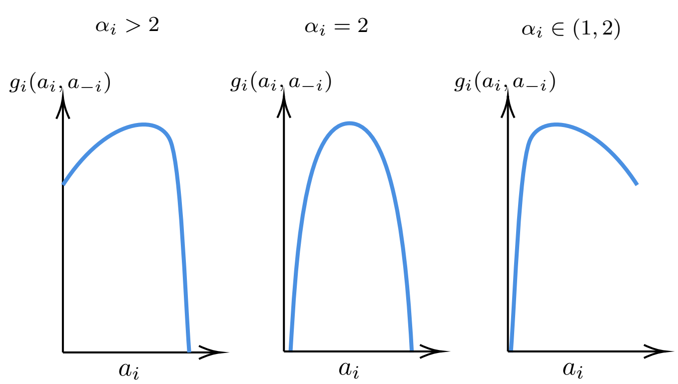

Figure 8: Payoffs for different choices of α_i.

When α_i is high and player i increases her investment, her payoffs decrease strongly after the peak. In this case, over-investment is a lot worse than under-investment. Conversely, when α_i is low, her payoff decreases a lot less after the peak. In this case, under-investment is a lot worse than over-investment. It can readily be verified that these payoffs give rise to a better-response cycle around every interior Nash equilibrium for a sufficiently fine grid of discrete actions for all $\alpha_i \neq 2$.

The key property that ensures the opposite skewness of players' payoff functions is the third-order derivative with respect to a player's own actions at each Nash equilibrium (a^*, a^*):

$$\prod_{i=1,2} \frac{\partial^3 g_i}{\partial a_i^3}(a^*, a^*) < 0. \tag{8.5}$$

When player i's third order derivative is negative, her payoffs are skewed to the left while a positive third-order derivative means that player i's payoffs are skewed to the right.

8.3 Example: Is Limit Uniqueness Robust to Redistribution?

Do budget balanced policy interventions, i.e. redistribution of payoffs, ever move a game outside of the region satisfying limit uniqueness?

Note that in Example 8.2, the parameter α_i only affected the distribution of equilibrium payoffs, not necessarily respecting budget balance outside of equilibrium play. We now provide an example which suggests that limit uniqueness is in fact robust to redistributive policies, at least within the class of concave supermodular games.

Consider a smooth, strictly concave, strictly supermodular game $g = (g_1, g_2)$ on action set $A_i = [0, 1]$, for every i. Specifically, for every i we consider a smooth payoff function $g_i \colon A_i \times A_{-i} \to \mathbb{R}$ satisfying $\frac{\partial^2 g_i}{\partial a_i \partial a_{-i}} > 0$, and $\frac{\partial^2 g_i}{\partial a_i^2} < 0$, $\forall a \in A_1 \times A_2$. Suppose, as it was the case in Example 8.2, that there is a continuum of symmetric, pure strategy Nash equilibria $\{(a^*, a^*) : a^* \in [0, 1]\}$ and that for every such Nash equilibrium (a^*, a^*), players' payoff functions are skewed in the same direction. That is,

$$\frac{\partial^3 g_i}{\partial a_i^3}(a_1, a_2) \geq 0, \ \forall \, i \in \{1, 2\}, \ \forall \, (a_1, a_2) \in A_1 \times A_2. \tag{8.6}$$

In that case, the game admits a generalized ordinal potential everywhere and so, as we argued in Example 8.2, there is a finite grid of $[0, 1]$, so that the restriction of g to that grid lies in a region satisfying limit uniqueness.

A *budget balanced intervention* is a player-specific smooth function $f_i\colon A_i \times A_{-i} \to \mathbb{R}$ for every player i, so that

$$f_1(a_1, a_2) + f_2(a_1, a_2) = 0, \ \forall \ (a_1, a_2) \in A_1 \times A_2. \tag{8.7}$$

Can we find a budget balanced intervention so that the induced game $\hat{g} = (g_1 + f_1, g_2 + f_2)$ satisfies (8.5) at some Nash equilibrium? We argue that, as long as $\hat{g}$ is also a concave supermodular game, the restriction of $\hat{g}$ to a finite grid will remain in the region of limit uniqueness.

In order to satisfy (8.5), redistribution must introduce skewness in opposing directions into players' payoff functions.

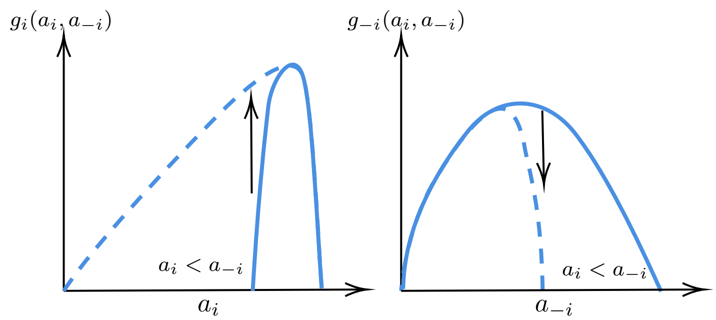

Figure 9: Budget Balanced Intervention cannot generate "opposite skewness" property.

Figure 9 illustrates the incompatibility between budget balance and generating opposite skewness: To skew player i's payoffs to the left when player $-i$ plays a_{-i}, $f_i(\cdot, a_{-i})$ would either have to increase i's payoffs when $a_i < a_{-i}$ (as can be seen in Figure 9) or decrease i's payoffs when $a_i > a_{-i}$. But in order to skew player $-i$'s payoffs to the right when player i plays an action a_i, $f_{-i}(a_i, \cdot)$ would either have to increase i's payoffs when $a_i < a_{-i}$ or decrease i's payoffs when $a_i > a_{-i}$. But that is impossible under budget balance.

9 Discussion

9.1 Spherical State Space and Rotational Noise

The sphere is a coarse representation of payoffs that is invariant to symmetric transformations. It is an appropriate state space to exhibit the symmetry properties of ICR. Admitting a compact state space, spherical global games have the advantage of admitting uniform priors. With classical global games we would have to resort to improper priors. An important difference to the classic set-up in Carlsson and Van Damme (1993) is the way in which the noise enters players' private signals. In this paper, the noise is linear (applied via a random rotation) rather than additively separable. This simplifies the analysis greatly as it allows us to use basic matrix algebra when deriving the symmetry properties of the limit selection in subsection 5.2. In subsection 9.2 below, we show that the affine set-up in Carlsson and Van Damme (1993) and our set-up are outcome-equivalent in the limit: every global game with additively separable noise induces an limit-outcome equivalent global game with rotational noise when the state space is restricted to the unit sphere.

9.2 Additively Separable Noise

A *classical global game* is a tuple of distributions $(\mu_0, (\mu_1, \mu_2))$ on the space of payoff function-pairs $U := \mathbb{R}^A \times \mathbb{R}^A$ with continuous, symmetric, bounded densities centered around the origin. For every $\sigma \in \mathbb{R}_+$, each player draws a private signal

$$x = w + \sigma\varepsilon_i, \tag{9.1}$$

where $w \in U$ is drawn according to μ_0 and $\varepsilon_i \in U$ is drawn (independently of w and ε_{-i}) according to distribution μ_i. Payoffs of player i when receiving private signal x are given by her payoff component in her own signal, $(x_i^{a_i,a_{-i}})_{a_i,a_{-i}} \in \mathbb{R}^A$. We let player i's ICR correspondence for the induced Bayesian game be denoted $\overline{\text{ICR}}_i^\sigma : U \to \mathcal{A}_i$ and let the induced limit selection be denoted $\overline{\text{ICR}}\colon U \to \mathcal{A}_1 \times \mathcal{A}_2$, so that for every $x \in U$ and player i,

$$\liminf_{\sigma \downarrow 0} \overline{\text{ICR}}_i^\sigma(x) = \overline{\text{ICR}}_i(x). \tag{9.2}$$

Let $\varphi\colon U \to S$ denote the spherical projection of payoffs onto the unit sphere of payoffs,

$$\varphi(u) := \frac{u}{||u||_2}. \tag{9.3}$$

Proposition 9.1 below establishes that every classical global game can be replicated by an outcome equivalent spherical global game.

Proposition 9.1 (Spherical Global Game Representation)**.** *For every global game $(\mu_\sigma)_{\sigma>0}$ there is a spherical global game $(\nu_\sigma)_{\sigma>0}$ so that for every i and* *every $u \in U$,*

$$\overline{\mathrm{ICR}}(\varphi(u)) = \mathrm{ICR}(\varphi(u)). \tag{9.4}$$

9.3 Related Literature

Frankel et al. (2003) consider supermodular games with many actions and players but restrict attention to one-dimensional noise, where the parameter over which there is noise enters monotonically into the payoffs. In their set-up, limit uniqueness holds for all supermodular games. They show that the equilibrium that gets selected may depend on the details of the noise. We made a substantive symmetry assumption on the structure of the noise. Under this assumption, the exact shape of the noise does not affect the coefficients of the multilinear form derived in Lemma 5.4. This contrasts our global games model with the one-dimensional model in Frankel et al. (2003). They also provide a sufficient condition for noise independent limit uniqueness in terms of local potentials, which is a special case of the generalized ordinal potential property holding on a set of actions. They say that an action profile $a^* \in A_1 \times A_2$ is a *local potential maximizer at s* if there exists a function $F\colon A_1 \times A_2 \to \mathbb{R}$ so that for every player i, for every $a_i \prec a_i^*$ there is $\overline{a}_i \succ a_i$; and for every $a_i \succ a_i^*$ there is $\underline{a}_i \prec a_i$, so that for all

$$(a_i, a_i') \in H_i(a_i^*) := \left\{ (a_i, a_i') : a_i \in A_i, \begin{matrix} a_i' \in [a_i, \overline{a}_i], \text{ if } a_i \prec a_i^* \\ a_i' \in [\underline{a}_i, a_i], \text{ if } a_i \succ a_i^* \end{matrix} \right\}, \tag{9.5}$$

there exists $\mu_i(a_i) > 0$ so that for all $a_{-i} \in A_{-i}$,

$$d_{s_i}^{a_i, a_i', a_{-i}} \mu_i(a_i) \geq F(a_i, a_j) - F(a_i', a_j). \tag{9.6}$$

Note that if a^* is a local potential maximizer at s for some potential function F then s admits a generalized ordinal potential F on $(H_1(a_1^*), H_2(a_2^*))$.

Oury (2013) considers the multidimensional noise set-up similar to this paper and finds that if a supermodular game satisfies noise independent limit uniqueness in every one-dimensional global game considered in Frankel et al. (2003), then limit uniqueness holds in the multi-dimensional case. Combining Frankel et al. (2003) and Oury (2013), we thus obtain a sufficient condition for limit uniqueness which is consistent with our characterization: If the game admits a local potential maximizer (and thus a generalized ordinal potential on some set of actions) then limit uniqueness holds. Oury (2013) considers multi-dimensional global games with additively separable noise and does not require symmetry of the noise. Under the additional structure we introduce, our paper is able to provide a converse to the sufficient condition in Oury (2013).

A more general notion of robustness to incomplete information is studied in Kajii and Morris (1997). Potential conditions have been found to be sufficient for this more general notion of robustness to incomplete information in Ui (2001).

9.4 Conclusion

This paper characterizes limit uniqueness in concave supermodular games. While achieving limit uniqueness in global games with multi-dimensional noise requires stronger conditions on payoffs compared to one-dimensional global games, the selection remains unaffected by the specific details of the noise. Theorem 6.1 provides a full characterization of limit uniqueness for all games, but is stated in terms of conditions which are hard to interpret.

An intriguing question for future research is how to characterize limit uniqueness on the sphere in terms of economically interpretable properties. Specifically, exploring the properties of the region where limit uniqueness holds could yield valuable insights. What types of interventions on payoffs either preserve or break limit uniqueness? We provide an example in which budget-balanced transfer schemes preserve limit uniqueness. We conjecture that this result may hold more broadly, but we will leave a formal proof for future research.

References

Bernheim, B. D. (1984): "Rationalizable strategic behavior," *Econometrica: Journal of the Econometric Society*, 1007–1028.

Carlsson, H. and E. Van Damme (1993): "Global games and equilibrium selection," *Econometrica: Journal of the Econometric Society*, 989–1018.

Dekel, E., D. Fudenberg, and S. Morris (2007): "Interim correlated rationalizability," *Theoretical Economics*.

Frankel, D. M., S. Morris, and A. Pauzner (2003): "Equilibrium selection in global games with strategic complementarities," *Journal of Economic Theory*, 108, 1–44.

Goldstein, I. and A. Pauzner (2005): "Demand–deposit contracts and the probability of bank runs," *the Journal of Finance*, 60, 1293–1327.

Harsanyi, J. C. and R. Selten (1988): "A general theory of equilibrium selection in games," *MIT Press Books*, 1.

Kajii, A. and S. Morris (1997): "The robustness of equilibria to incomplete information," *Econometrica: Journal of the Econometric Society*, 1283–1309.

Monderer, D. and L. S. Shapley (1996): "Potential games," *Games and economic behavior*, 14, 124–143.

Morris, S. and H. S. Shin (1998): "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, 587–597.

——— (2003): *Global Games: Theory and Applications*, Cambridge University Press, 56114, Econometric Society Monographs.

——— (2004): "Coordination risk and the price of debt," *European Economic Review*, 48, 133–153.

Oury, M. (2013): "Noise-independent selection in multidimensional global games," *Journal of Economic Theory*, 148, 2638–2665.

Pearce, D. G. (1984): "Rationalizable strategic behavior and the problem of perfection," *Econometrica: Journal of the Econometric Society*, 1029–1050.

Ui, T. (2001): "Robust equilibria of potential games," *Econometrica*, 69, 1373–1380.

A Appendix: Omitted Proofs

Proof of Lemma 4.1

Proof. Let $B = (B_1, B_2)$ and $\hat{B} = (\hat{B}_1, \hat{B}_2)$ be two disjoint BRC sets satisfying $A_i = B_i \cup \hat{B}_i$, for every player i. Suppose for the sake of a contradiction that both BRC seta fail the risk-dominance criterion (4.4), i.e.

$$p_1^B(s) + p_2^B(s) > 1. \tag{A.1}$$

Then for every player i we must have that $p_i^B(s) > 0$. Fix an arbitrary player i and so

$$\mathrm{br}_i(p_i^*|s_i) \cap B_i = \varnothing, \tag{A.2}$$

for all $p_i^* \in \Delta(A_{-i})$ so that $p_i^*(B_{-i}) < p_i^B(s)$. But then, for every player i, there is a probability $p_i^{**} \in \Delta(A_{-i})$ satisfying $p_i^{**}(\hat{B}_{-i}) \leq 1 - p_i^B(s)$ so that

$$\mathrm{br}_i(p_i^{**}|s_i) \subseteq \hat{B}_i, \tag{A.3}$$

where $p_i^{**}(\hat{B}_{-i}) + p_{-i}^{**}(\hat{B}_i) \leq 1$, a contradiction. □

Proof of Lemma 4.2

Proof. Fix $\sigma > 0$ and suppose that for all $B \in \mathscr{R}(s)$, $a_i \notin B_i$. We show that $a_i \notin \mathrm{ICR}(s)$. Let $B^* = (B_1^*, B_2^*)$ denote the union of all risk-dominant BRC sets (or the full action set if no such BRC set exists), i.e. for every player i

$$B_i^* := \bigcup_{(B_1,B_2)\in\mathscr{R}(s)} B_i. \tag{A.4}$$

Since $a_i \notin B_i^*$, we have that for all $(p_1, p_2) \in \Delta(A_2) \times \Delta(A_1)$ so that $a_i \in \mathrm{br}_i(p_i)$ for each $i = 1, 2$,

$$p_1(B_2) + p_2(B_1) > 1, \ \forall\ (B_1, B_2) \in \mathscr{R}(s). \tag{A.5}$$

Put differently for any $(B_1, B_2) \in \mathscr{R}(s)$, for all $p_1, p_2 \in [0, 1]$ so that $p_1 + p_2 \leq 1$, every player i and every $q_i \in \Delta(A_{-i})$ so that $q_i(B_{-i}) < p_i$, we have that

$$a_i \notin \mathrm{br}_i(q_i|s). \tag{A.6}$$

Fix $(B_1, B_2) \in \mathscr{R}(s)$. For any $p_1, p_2 \in [0,1]$ define the sequence of sets $(R^{m,\sigma}_{p,i})_{m\in\mathbb{N}}$ for every player i,

$$
\begin{aligned}
R^{1,\sigma}_{p,i} &:= \bigcup_{\hat{a}_i \in B_i} \{s \in S : s_i^{\hat{a}_i, a_{-i}} - s_i^{a'_i, a_{-i}} > 0,\ \forall\, a'_i \neq \hat{a}_i, \forall\, a_{-i} \in A_{-i}\} \\
R^{m,\sigma}_{p,i} &:= \{\hat{s} \in S : \nu^\sigma(R^{m-1,\sigma}_{p,j} | \hat{s}) \geq p_i\} \\
&\vdots \qquad\qquad \vdots
\end{aligned}
\tag{A.7}
$$

First note that

$$R^{1,\sigma}_{p,i} \supseteq \{s \in S : \mathrm{ICR}^{\sigma,1}_i(s) \subseteq B_i\}. \tag{A.8}$$

Moreover,

$$R^{1,\sigma}_{p,i} \cap \{s \in S : a_i \in \mathrm{ICR}^{\sigma,1}_i(s)\} = \varnothing. \tag{A.9}$$

But then we must have that for any numbers $p_1, p_2 \in [0,1]$ so that $p_1 + p_2 \leq 1$, and any $m \in \mathbb{N}$,

$$R^{m,\sigma}_{p,i} \cap \{s \in S : a_i \in \mathrm{ICR}^{\sigma,m}_i(s)\} = \varnothing. \tag{A.10}$$

For every m, define the product $R^{m,\sigma}_p := R^{m,\sigma}_{p,1} \times R^{m,\sigma}_{p,2}$ and let $R^{\infty,\sigma}_p := \bigcup_{m=1}^\infty R^{m,\sigma}_p$. Recall that ν_σ is the joint probability on the signal profiles and the latent state y. Applying the critical path theorem in Kajii and Morris (1997) we have that

$$\nu_\sigma\left(R^{\infty,\sigma}_p | y \notin \cup_i R^{1,\sigma}_{p,i}\right) \geq 1 - \frac{1 - \min_i p_i}{1 - p_1 - p_2}\left(1 - \nu_\sigma\left(R^{1,\sigma}_p | y \notin \cup_i R^{1,\sigma}_{p,i}\right)\right) \tag{A.11}$$

and so in particular, we must have that $p_1 + p_2 < 1$. Moreover,

$$\lim_{\sigma \downarrow 0} \nu_\sigma\left(R^{1,\sigma}_p | y \notin \cup_i R^{1,\sigma}_{p,i}\right) = 0. \tag{A.12}$$

Then for all $\epsilon > 0$ there must exist σ so that

$$\nu_\sigma\left(\{\tilde{s} \in S : ||\tilde{s} - s||_2 < \epsilon\} \setminus R^{\infty,\sigma}_p \mid y \notin \cup_i R^{1,\sigma}_{p,i}\right) < \epsilon. \tag{A.13}$$

But then $a_i \notin \mathrm{ICR}_i(s)$. □

Proof of Proposition 5.1

Proof. For any pair of action sets $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$, we construct a game $s_B \in S$ with a risk dominant best-response cycle involving all actions in B. For the remaining actions, we scale the payoffs down to ϵ to ensure that all BRC sets involving actions outside of B are risk-dominated. Consider any order on the actions $B_i = \{a_i^1, \dots, a_i^{|B_i|}\}$. For any $k < |B_i|$, define the payoffs for player 1,

$$s_1^{a_1,a_2} = \begin{cases} 1, \text{ if } \exists\ k < |B_i| \text{ s.t. } a_1 = a_1^k, a_2 = a_2^{k+1} \\ 1, \text{ if } a_1 = a^{|B_i|}, a_2 = a_2^1 \\ -1, \text{ if } \exists\ k \text{ s.t. } a_1 = a_1^k, a_2 \neq a_2^{k+1} \\ -1, \text{ if } a_1 = a_1^{|B_i|}, a_2 \neq a_2^1 \\ \epsilon, \text{ otherwise.} \end{cases} \tag{A.14}$$

For player 2,

$$s_2^{a_1,a_2} = \begin{cases} 1, \text{ if } \exists\ k > 1 \text{ s.t. } a_2 = a_2^k, a_1 = a_1^{k-1} \\ 1, \text{ if } a_2 = a_2^1, a_1 = a_1^{|B_i|} \\ -1, \text{ if } \exists\ k > 1 \text{ s.t. } a_2 = a_2^k, a_1 \neq a_1^{k-1} \\ -1, \text{ if } a_2 = a_2^1, a_1 \neq a_1^{|B_i|} \\ \epsilon, \text{ otherwise.} \end{cases} \tag{A.15}$$

For ϵ small enough, (B_1, B_2) becomes a risk dominant BRC set. ☐

Claim A.1. *For every $s \in \mathcal{R}^\sigma(B)$ and every $a_i \in B_i$ there exists a selection rule $\kappa \in \Sigma_{-i}$ so that*

$$\sum_{B_{-i},a_{-i}} P_i^\sigma(B_{-i}|s)\kappa(a_{-i}|B_{-i})(s_i^{a_i,a_{-i}} - s_i^{a_i',a_{-i}}) \geq 0, \forall\ a_i' \in A_i. \tag{A.16}$$

Proof. For every $s \in \mathcal{R}^\sigma(B)$ there exists a correlated conjecture $\kappa_{-i}\colon S \times S \to \Delta(A_{-i})$ so that

$$B_i = \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^\sigma} \mathrm{br}_i \left(\int_{S \times S} \kappa_{-i}(\cdot|y, s') \mathrm{d}\nu_\sigma(y, s'|s) \right), \tag{A.17}$$

where $\mathcal{S}_{-i}^\sigma \coloneqq \{\kappa_{-i}\colon S \times S \to \Delta(A_{-i}) :\ \forall\ y, s',\ \mathrm{supp}(\kappa_{-i}(y, s')) \subseteq \mathrm{ICR}_{-i}^\sigma(s')\}$. For every $\kappa_{-i} \in \mathcal{S}_{-i}^\sigma$, define $\hat{\kappa}_{-i}\colon \mathcal{A}_{-i} \to \Delta(A_{-i})$, which for every $B_{-i} \in \mathcal{A}_{-i}$

satisfies

$$\hat{\kappa}_{-i}(a_{-i}|B_{-i}) \coloneqq \int_S \int_{\mathcal{R}^\sigma_{-i}(B_{-i})} \kappa_{-i}(a_{-i}|y, s')\mathrm{d}\nu_\sigma(y, s'|s), \ \forall\ a_{-i} \in B_{-i}. \tag{A.18}$$

We conclude that

$$\begin{aligned}\bigcup_{\kappa_{-i}\in\mathcal{S}^\sigma_{-i}} \mathrm{br}_i\left(\int_{S\times S} \kappa_{-i}(\cdot|y, s')\mathrm{d}\nu_\sigma(y, s'|s)\right) &= \bigcup_{\kappa_{-i}\in\mathcal{S}^\sigma_{-i}} \mathrm{br}_i\left(\sum_{B_{-i}\in\mathcal{A}_{-i}} P^\sigma_i(B_{-i}|s)\hat{\kappa}_{-i}(a_{-i}|B_{-i})\right)\\ &= \bigcup_{\kappa\in\Sigma_{-i}} \mathrm{br}_i\left(\sum_{B_{-i}\in\mathcal{A}_{-i}} P^\sigma_i(B_{-i}|s)\kappa(a_{-i}|B_{-i})\right)\end{aligned} \tag{A.19}$$

□

Proof of Lemma 5.1

Proof. Let $(a_i, a'_i) \in B_i \times B'_i$. Claim A.1 above shows that for every $s \in \partial\mathcal{R}^\sigma(B, B')$, there exists a selection rule $\kappa \in \Sigma_{-i}$ so that

$$\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)\kappa(a_{-i}|B_{-i})(s_i^{a_i,a_{-i}} - s_i^{a'_i,a_{-i}}) = 0. \tag{A.20}$$

Since both a_i, a'_i are rationalizable at s, there are conjectures $\kappa, \kappa' \in \Sigma_{-i}$ so that

$$\begin{aligned}&\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)\kappa(a_{-i}|B_{-i})(s_i^{a_i,a_{-i}} - s_i^{\hat{a}_i,a_{-i}}) \geq 0, \ \forall\ \hat{a}_i \in A_i,\\ &\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)\kappa'(a_{-i}|B_{-i})(s_i^{a'_i,a_{-i}} - s_i^{\hat{a}_i,a_{-i}}) \geq 0, \ \forall\ \hat{a}_i \in A_i.\end{aligned} \tag{A.21}$$

But then we have that

$$\begin{aligned}&\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)\kappa'(a_{-i}|B_{-i})(s_i^{a'_i,a_{-i}} - s_i^{a_i,a_{-i}}) \geq 0\\ &\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)\kappa(a_{-i}|B_{-i})(s_i^{a'_i,a_{-i}} - s_i^{a_i,a_{-i}}) \leq 0.\end{aligned} \tag{A.22}$$

So there exists $\alpha \in [0, 1]$ so that

$$\sum_{B_{-i},a_{-i}} P^\sigma_i(B_{-i}|s)(\alpha\kappa'(a_{-i}|B_{-i}) + (1-\alpha)\kappa(a_{-i}|B_{-i}))(s_i^{a'_i,a_{-i}} - s_i^{a_i,a_{-i}}) = 0. \tag{A.23}$$

Moreover, for every $\epsilon > 0$ and any $\tilde{s} \in \mathcal{R}^\sigma(B')$ so that $||s - \tilde{s}_\epsilon||_2 < \epsilon$ we must have that

$$\sum_{B_{-i},a_{-i}} P_i^\sigma(B_{-i}|\tilde{s}_\epsilon)\tilde{\kappa}(a_{-i}|B_{-i})(\tilde{s}_{\epsilon i}^{a_i,a_{-i}} - \tilde{s}_{\epsilon i}^{a_i',a_{-i}}) < 0, \tag{A.24}$$

for all selection rules $\tilde{\kappa} \in \Sigma_{-i}$. Considering any convergent sequence $(\tilde{s}_{\epsilon^n})_n \to s$, continuity of beliefs implies that

$$\lim_{n\uparrow\infty} \max_{\tilde{\kappa}\in\Sigma_{-i}} \sum_{B_{-i},a_{-i}} P_i^\sigma(B_{-i}|\tilde{s}_{\epsilon^n})\tilde{\kappa}(a_{-i}|B_{-i})(\tilde{s}_{\epsilon^n i}^{a_i,a_{-i}} - \tilde{s}_{\epsilon^n i}^{a_i',a_{-i}}) = 0, \tag{A.25}$$

which concludes the proof. □

Proof of Proposition 5.2

Proof. We have shown in Claim A.1 that every $s \in \mathcal{R}^\sigma(B)$ satisfies (A.16). The converse is immediate since any selection rule $\kappa \in \Sigma_{-i}$ is a particular correlated conjecture, we thus conclude that if there exists a selection rule $\kappa \in \Sigma_{-i}$ so that

$$\sum_{B_{-i},a_{-i}} P_i^\sigma(B_{-i}|s)\kappa(a_{-i}|B_{-i})(s_i^{a_i,a_{-i}} - s_i^{a_i',a_{-i}}) \geq 0, \forall\, a_i' \in A_i. \tag{A.26}$$

then there exists $\hat{B} \in \mathcal{A}_1 \times \mathcal{A}_2$ so that $a_i \in \hat{B}_i$ and $s \in \mathcal{R}^\sigma(\hat{B})$. The result in Proposition 5.2 then follows from the continuity of $s \mapsto P_i^\sigma(B_{-i}|s)$. □

Proof of Lemma 5.2

Proof. Fix $\sigma > 0$ and a pair (η, X) that preserves best-replies. We proceed inductively on the rounds of deletion. Suppose X satisfies for all $s \in S$

$$\eta(\mathrm{ICR}^{m,\sigma}(s)) = \mathrm{ICR}^{m,\sigma}(Xs). \tag{A.27}$$

The case $m = 0$ is easy since $\mathrm{ICR}^{0,\sigma}$ is the constant map. Then for every s

$$\begin{aligned}
\bar{P}^{m,\sigma}(b|s) &:= \nu \times \nu_0 \left(\left\{(E, E', y) : B = \mathrm{ICR}^{m,\sigma}(e^{\sigma E'}y)\right\} \middle| e^{\sigma E}y = s\right) \\
&= \nu\left(\left\{(E, E') : B = \mathrm{ICR}^{m,\sigma}(e^{\sigma(E'-E)}s)\right\}\right) \\
&= \nu\left(\left\{(E, E') : \eta(B) = \mathrm{ICR}^{m,\sigma}(Xe^{\sigma(E'-E)}s)\right\}\right) \\
&= \nu\left(\left\{(E, E') : \eta(B) = \mathrm{ICR}^{m,\sigma}(Xe^{\sigma(E'-E)}X^\top Xs)\right\}\right) \\
&= \nu\left(\left\{(E, E') : \eta(B) = \mathrm{ICR}^{m,\sigma}(e^{\sigma(E'-E)}Xs)\right\}\right) \\
&= \bar{P}^{m,\sigma}(\eta(B)|Xs).
\end{aligned} \tag{A.28}$$

Suppose s, s' satisfy

$$\mathrm{ICR}^{m+1,\sigma}(s) = \mathrm{ICR}^{m+1,\sigma}(s'). \tag{A.29}$$

Then by property (5.15), we must have that

$$\mathrm{ICR}^{m+1,\sigma}(Xs) = \mathrm{ICR}^{m+1,\sigma}(Xs'). \tag{A.30}$$

□

Claim A.2. *For every $B \in \mathcal{A}_1 \times \mathcal{A}_2$ and every $\sigma > 0$, the set $\mathcal{R}^\sigma(B)$ is path-connected.*

Proof. Note that for every $\sigma > 0$, the first-order rationalizable sets $\mathcal{R}^{\sigma,1}(B)$ are connected for every $B \in \mathcal{A}_1 \times \mathcal{A}_2$. By upper-hemi continuity of best-replies, the set $\mathcal{R}^{\sigma,m}(B)$ is connected for every $B \in \mathcal{A}_1 \times \mathcal{A}_2$ and every $m \in \mathbb{N}$. □

Claim A.3. *For every $\tilde{X} \in \mathrm{P}^{2n}$ there exists permutation $\tilde{\eta}\colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathcal{A}_1 \times \mathcal{A}_2$ so that $(\tilde{\eta}, \tilde{X})$ preserves best-replies.*

Proof. For every $\tilde{X} \in \mathrm{P}^{2n}$ construct an associated coordinate permutation $\zeta\colon A_1 \times A_2 \to A_1 \times A_2$. Then we obtain an associated label permutation $\eta_\zeta\colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathcal{A}_1 \times \mathcal{A}_2$,

$$\eta_\zeta(b_1, b_2) := (\{\zeta_1(a_1, a_2) : (a_1, a_2) \in B_1 \times B_2\}, \{\zeta_2(a_1, a_2) : (a_1, a_2) \in B_1 \times B_2\}). \tag{A.31}$$

□

Proof of Lemma 5.3

Proof. Let $\mathscr{I}$ denote the collection of invariances. Let η be a permutation that acts on Z so that $(\eta, X) \in \mathscr{I}$ for some $X \in \mathrm{P}^{2n}$. Define the collection of permutation matrices, which paired with η define an invariance:

$$\mathcal{X}(\eta) := \{\tilde{X} \in \mathrm{P}^{2n} : (\eta, \tilde{X}) \in \mathscr{I}\}. \tag{A.32}$$

Let $s^* \in \mathcal{C}(C_Z)$ and suppose

$$Xs^* \neq s^*. \tag{A.33}$$

Let $\mathcal{K} \coloneqq \{(i,a) \in \{1,2\} \times (A_1 \times A_2) : (s^*)_i^a \neq (Xs^*)_i^a\}$. It can readily be shown that there exists a permutation $\omega_{\mathcal{K}} \colon \mathcal{K} \to \mathcal{K}$ and an associated matrix $X_{\omega_{\mathcal{K}}} \in \mathrm{P}^{2n}$ so that

$$\hat{X}^{(i,a),(i',a')} = \begin{cases} \pm 1, \text{ if } (i,a) \in \mathcal{K} \text{ and } \omega_{\mathcal{K}}(i,a) = (i',a'), \\ 1, \text{ if } (i,a) \notin \mathcal{K} \text{ and } (i,a) = (i',a'), \\ 0, \text{ otherwise,} \end{cases} \tag{A.34}$$

and

$$\hat{X}s^* = Xs^*. \tag{A.35}$$

Next, let $S^+, S^- \subseteq S$ be two disjoint sets so that

$$\{\hat{X}s : s \in S^+\} = S^-, \tag{A.36}$$

and $\overline{S^+} \cup \overline{S^-} = S$. Fix $\sigma > 0$ and suppose without loss of generality that $s^* \in S^+$. Consider two paths described by continuous maps $\gamma^+ \colon [0,1] \to \overline{S^+}$ and $\gamma^- \colon [0,1] \to \overline{S^-}$ satisfying the following conditions

(i) $\mathrm{ICR}^\sigma(\gamma^+(x)) = \mathrm{ICR}^\sigma(\gamma^-(x'))$, for every $x, x' \in [0,1]$,

(ii) $\gamma^+(0) = s^*$,

(iii) $\gamma^+(1) = \gamma^-(1) \in \overline{S^+} \cap \overline{S^-}$.

Since the permutation acts on Z and $s^* \in \mathcal{C}(C_Z)$, it must be that $\mathcal{R}^\sigma(\mathrm{ICR}(s^*)) \cap S^+ \neq \varnothing$ and $\mathcal{R}^\sigma(\mathrm{ICR}(s^*)) \cap S^- \neq \varnothing$. We show in Claim A.2 that each rationalizable set is path-connected and so the existence of γ^+, γ^- satisfying properties (i)-(iii) is guaranteed. Suppose now that

$$\mathrm{ICR}^\sigma(s^*) \neq \mathrm{ICR}^\sigma(\hat{X}s^*). \tag{A.37}$$

We show in Claim A.3 that for every matrix $\tilde{X} \in \mathrm{P}^{2n}$ there is a permutation of action set labels $\tilde{\eta}$ so that $(\tilde{\eta}, \tilde{X}) \in \mathscr{I}$. Since $\hat{X} \in \mathrm{P}^{2n}$ is itself part of an invariance, it must be that

$$\mathrm{ICR}^\sigma(\gamma^+(x)) \neq \mathrm{ICR}^\sigma(\hat{X}\gamma^+(x)), \ \forall\, x \in [0,1]. \tag{A.38}$$

But that contradicts property (iii) of the paths γ^+, γ^-. So it must be that

$$\mathrm{ICR}^\sigma(\gamma^+(x)) = \mathrm{ICR}^\sigma(\hat{X}\gamma^+(x)), \ \forall\, x \in [0,1]. \tag{A.39}$$

Hence $\hat{X}^\top X \in \mathcal{X}(\eta)$. But then we have that $(\eta, \hat{X}^\top X)$ is an invariance so that

$$\hat{X}^\top X s^* = s^*. \tag{A.40}$$

□

Proof of Lemma 5.4

Proof. For every $B \in \mathcal{A}_1 \times \mathcal{A}_2$, player i, and $\sigma > 0$ define

$$p^\sigma_{i,s}(B) := \nu \times \nu_0 \left(\{(E_i, E_{-i}, y) : B = \mathrm{ICR}^\sigma(e^{\sigma E_{-i}} y)\} \middle| e^{\sigma E_i} y = s \right). \tag{A.41}$$

Fix a permutation $\eta \colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathcal{A}_1 \times \mathcal{A}_2$ and some choice of $k^* \leq m_i$. Let $(s^{k^*}_\sigma)_{\sigma>0}$ be a sequence converging to $s \in \mathcal{C}(C)$ so that for every $\sigma > 0$, $s^{k^*}_\sigma \in \mathcal{G}^\sigma_i(a^{k^*}_i, a'^{k^*}_i)$. Since Z is full at s, there is $X \in \mathrm{P}^{2n}$ so that (η, X) is an invariance with invariant point s. For every $k \neq k^*$, let $s^k_\sigma := X^\sigma s^{k^*}_\sigma$ and define the limit probabilities for every $k \leq m_i$,

$$p_{i,s}(B) := \lim_{\sigma \downarrow 0} p^\sigma_{i,s^k_\sigma}(B),\ \forall\ B \in \mathcal{A}_1 \times \mathcal{A}_2. \tag{A.42}$$

Since every collision is full, by varying the permutation, we deduce from Claims 5.2 and 5.3 that

$$\sum_{k=1}^{m_i} p_{i,s^k_\sigma}(B) = 1. \tag{A.43}$$

Define the pseudo-beliefs for every $k \leq m_i$ and $a \in B_1 \times B_2$,

$$\tilde{p}_{i,k,s}(a, B) := \frac{p_{i,s}(B)}{\prod_{j=1,2} \prod_{\hat{k} \neq k} d^{a^{\hat{k}}_j, a'^{\hat{k}}_j, a_{-j}}_{s_j}}. \tag{A.44}$$

Then for every $s \in \mathcal{C}(C)$ and every $k \leq m_i$, we have that

$$\lim_{\sigma \downarrow 0} g^{a^k_i, a'^k_i, \sigma}_i(s) = \sum_{B \in \mathcal{A}_1 \times \mathcal{A}_2} \sum_{a \in B_1 \times B_2} \tilde{p}_{i,k,s}(a, B) \underbrace{\prod_{j=1,2} \prod_{\hat{k}=1}^{m_j} d^{a^{\hat{k}}_j, a'^{\hat{k}}_j, a_{-j}}_{s_j}}_{:=\pi_{a,s}} = 0. \tag{A.45}$$

The vectors $(\tilde{p}_{i,k,s})_{k \leq K}$ are all orthogonal to the same vector of products $\pi_s = (\pi_{a,s})_a$. We show that there exists a vector with entries in $\{-1, 0, 1\}$ which lies in the span of $(\tilde{p}_{i,k,s})_{k \leq K}$. Consider for instance the weighted sum of all these vectors

$$\bar{p}_{i,s}(a, B) = \sum_{k=1}^{m_i} \tilde{p}_{i,k,s}(a, B) \prod_{j=1,2} \prod_{\hat{k}=1}^{m_j} d^{a^{\hat{k}}_j, a'^{\hat{k}}_j, a_{-j}}_{s_j}. \tag{A.46}$$

By condition (A.43), $\bar{p}_s$ is a vector with entries in $\{0,1\}$. Hence the limit of every IC takes the same form:

$$\delta_C(d_s^1,\ldots,d_s^K) = \sum_{a\in A_1\times A_2} \chi_{a,C}(s) \prod_{i=1,2}\prod_{k=1}^{m_i} d_{s_i}^{a_i^k,a_i'^k,a_{-i}}, \tag{A.47}$$

for some family of binary coefficients $(\chi_{a,C}(s))_a \in \{0,1\}^{A_1\times A_2}$. By continuity, we deduce that the coefficients are in fact independent of s. ☐

Proof of Proposition 6.1

Proof. Suppose $\mathcal{C}(C)$ is a singleton collision but incentives are not aligned on C. Then there exists an un-dominated sequence of cyclic perturbations $((a_1^1,\ldots,a_1^m),(a_2^1,\ldots,a_2^m)) \in G_{C,1}\times G_{C,2}$ in C. For any $s\in\mathcal{C}(C)$, consider a perturbation in the direction that increases $d_{s_i}^{a_i^2,a_i^1,a_{-i}^1}$: Let $\xi :=(0,\ldots,1,\ldots,-1,\ldots,0)$, where the entry equal to 1 is at the coordinate $s_i^{a_i^2,a_{-i}^1}$ and the entry equal to -1 is at the coordinate $s_i^{a_i^1,a_{-i}^1}$. Since the sequence is un-dominated, there exists $\bar\epsilon>0$ so that for all $\epsilon<\bar\epsilon$, every player i and every $n\le m$,

$$a_i^n \in \mathrm{ICR}_i^\sigma(s+\epsilon\xi). \tag{A.48}$$

Indeed, the initial perturbation breaks the tie between a_{-i}^2 and a_{-i}^1 in favor of a_{-i}^2. Then we have that

$$\frac{\partial p_i^{\sigma,a_i^2,a_i^1}(a_{-i}^2|s)}{\partial d_i^{a_i^2,a_i^1,a_{-i}^1}} > 0. \tag{A.49}$$

where

$$p_i^{\sigma,a_i^2,a_i^1}(a_{-i}^2|s) := \sum_{b_{-i}\in\mathcal{A}_{-i}} P^\sigma(B_{-i}|s)\kappa(a_{-i}^2|B_{-i}), \tag{A.50}$$

where $\kappa\in\Sigma_{-i}$ solves (5.5) for $a_i=a_i^2$ and $a_i'=a_i^1$. Then $a_{-i}^2\in I^{a_i^2,a_i^1}$ implies that

$$\frac{\partial g^{a_i^2,a_i^1,\sigma}}{\partial p_i^{\sigma,a_i^2,a_i^1}(a_{-i}^2|s)} > 0. \tag{A.51}$$

Expression (A.53) then follows from the fact that the sequence is un-dominated: Every sequence of tie breaks in favor of an action outside of the cyclic sequence can be extended to break a tie in favor of some other action in the cyclic sequence. But then $\mathcal{C}(C)$ cannot be singleton collision, a contradiction.

Suppose now that incentives are aligned on C. We argue that for generic $s \in \mathcal{C}(C)$, and generic perturbation direction $d_{s_i}^{a_i,a_i',a_{-i}}$, $|\text{ICR}(s+\epsilon\xi)| = 1$, for $\epsilon > 0$ small enough and $\xi := (0, \dots, 1, \dots, -1, \dots, 0)$, where the entry equal to 1 is at the coordinate $s_i^{a_i,a_{-i}}$ and the entry equal to -1 is at the coordinate $s_i^{a_i',a_{-i}}$. Indeed, since there is no un-dominated sequence of cyclic perturbations, for any sequence of cyclic perturbations $(a_i^1, \dots, a_i^m)_{i=1,2}$, there exists $(a_1^*, a_2^*) \in A_1 \times A_2$ so that for all $n \leq m$,

$$\frac{\partial g^{a_i^*,a_i^n,\sigma}}{\partial d_{s_i}^{a_i,a_i',a_{-i}}} > 0. \tag{A.52}$$

Hence there exists $\bar{\epsilon} > 0$ so that for all $\epsilon < \bar{\epsilon}$, every player i and every $n \leq m$,

$$a_i^n \notin \text{ICR}_i^\sigma(s+\epsilon\xi). \tag{A.53}$$

And so the result follows. □

Lemma A.1 (AI)**.** *Suppose s is strictly supermodular and strictly concave. Then for all actions $(a_i, a_i') \in C_i$, $(a_{-i}, a_{-i}') \in C_{-i}$,*

$$\begin{array}{c}(a_{-i} \in I^{a_i,a_i'}, a_i \in I^{a_{-i},a_{-i}'}) \\ or \\ (a_{-i}' \in I^{a_i',a_i}, a_i' \in I^{a_{-i}',a_{-i}})\end{array} \iff \frac{d_{s_i}^{a_i,a_i',a_{-i}}}{d_{s_{-i}}^{a_{-i},a_{-i}',a_i}} > 0 \tag{A.54}$$

Proof. Recall that

$$\delta_C(d_s^1, \dots, d_s^K) = \sum_{a \in \bar{b}} \chi_{a,C} \prod_{i=1,2} \prod_{k=1}^{m_i} d_{s_i}^{a_i^k,a_i'^k,a_{-i}}. \tag{A.55}$$

Then we have that

$$\frac{\partial \delta_C}{\partial d_{s_i}^{a_i,a_i',a_{-i}}} = \prod_{k \neq k_{d_{s_i}^{a_i,a_i',a_{-i}}}} d_{s_i}^{a_i^k,a_i'^k,a_{-i}} \sum_{\tilde{a} \in A_1 \times A_2 : \tilde{a}_{-i} = a_{-i}} \chi_{\tilde{a},C} \prod_{k=1}^{m_{-i}} d_{s_{-i}}^{a_{-i}^k,a_{-i}'^k,\tilde{a}_i}, \tag{A.56}$$

where

$$k_{d_{s_i}^{a_i,a_i',a_{-i}}} := \{k \leq m_i : a_i^k = a_i, a_i'^k = a_i'\}. \tag{A.57}$$

By strict supermodularity and strict concavity we conclude that for every a_{-i} there exists unique a_i so that $\chi_{a_{-i},a_i,C} = 1$ and so that

$$\operatorname{sgn}\left(\frac{\partial \delta_C}{\partial d_{s_i}^{a_i,a'_i,a_{-i}}}\right) = \operatorname{sgn}\left(\prod_{k \neq k_{d_{s_i}^{a_i,a'_i,a_{-i}}}} d_{s_i}^{a_i^k,a'^k_i,a_{-i}} \prod_{k=1}^{m_{-i}} d_{s_{-i}}^{a^k_{-i},a'^k_{-i},a_i}\right). \tag{A.58}$$

where a_i satisfies

$$a_i \in \max_{\prec_i}\{\tilde{a}_i \in A_i : \chi_{a_{-i},\tilde{a}_i,C} = 1\}. \tag{A.59}$$

Hence

$$\frac{\frac{\partial \delta_C}{\partial d_{s_i}^{a_i,a'_i,a_{-i}}}}{\frac{\partial \delta_C}{\partial d_{s_{-i}}^{a_{-i},a'_{-i},a_i}}} = \frac{d_{s_{-i}}^{a_{-i},a'_{-i},a_i}}{d_{s_i}^{a_i,a'_i,a_{-i}}}. \tag{A.60}$$

□

Proof of Lemma 7.1

Proof. This result then follows from Lemma A.1 and Proposition 7.1. □

Proof of Theorem 7.1:

Proof. The argument in the main text establishes that if every $s \in O$ has a generalized ordinal potential on a risk dominant BRC set, the limit uniqueness holds. We thus focus on the converse. Let $O \subseteq S$ satisfy limit uniqueness and suppose $s \in S$ admits a risk dominant BRC set where the generalized ordinal potential property fails. By Lemma 2.5 in Monderer and Shapley (1996), s has a better response cycle, i.e. an ordered collection of actions for each player $\{(a_1^1, \ldots, a_1^m), (a_2^1, \ldots, a_2^m)\}$ and a player i so that for every $n \in \{1, \ldots, m-1\}$,

$$\begin{aligned} u_i(a_i^n, a_{-i}^{n+1}) &< u_i(a_i^{n+1}, a_{-i}^{n+1}), \\ u_{-i}(a_{-i}^n, a_i^n) &< u_{-i}(a_{-i}^{n+1}, a_i^n), \end{aligned} \tag{A.61}$$

and

$$\begin{aligned} u_i(a_i^m, a_{-i}^1) &< u_i(a_i^1, a_{-i}^1), \\ u_{-i}(a_{-i}^m, a_i^m) &< u_{-i}(a_{-i}^1, a_i^m). \end{aligned} \tag{A.62}$$

Let $\bar{s} \in S$ be a game where $\{(a_1^1, \dots, a_1^m), (a_2^1, \dots, a_2^m)\}$ is the maximal risk dominant best-response cycle. We consider the following path $\mathcal{P}_{s,\bar{s}} \subseteq S$,

$$\mathcal{P}_{s,\bar{s}} \coloneqq \{s_\alpha : \alpha \in [0,1]\}, \tag{A.63}$$

where $s_\alpha \coloneqq \frac{\alpha s + (1-\alpha)\bar{s}}{||\alpha s + (1-\alpha)\bar{s}||_2}$, for every $\alpha \in [0,1]$. By Lemma 4.1 we conclude that $\bar{s}$ satisfies limit multiplicity, i.e. for every player i,

$$\{a_i^1, \dots, a_i^m\} \subseteq \mathrm{ICR}_i(\bar{s}). \tag{A.64}$$

First, note that $\{(a_1^1, \dots, a_2^m), (a_2^1, \dots, a_2^m)\}$ is contained in the maximal risk-dominant BRC set of s_α, for all $\alpha \in [0,1]$. Let $\alpha^1, \dots, \alpha^L \in [0,1]$ denote the points along the path $\mathcal{P}_{s,\bar{s}}$ so that s_{α^l} lies on a collision for every $l \leq L$. We now argue that for every $l \leq L$ and every player i,

$$\{a_i^1, \dots, a_i^m\} \subseteq \mathrm{ICR}_i(s_{\alpha^l}). \tag{A.65}$$

Indeed, for $l = 1$, the collision must involve $\{a_i^1, \dots, a_i^m\}$ for each player i and so incentives are not aligned at that collision. By Theorem 6.1, s_{α^1} is not a singleton collision. We now argue that the better-response cycle breaks the indifference constraints in favor of all actions in the cycle and so $\{a_i^1, \dots, a_i^m\} \subseteq \mathrm{ICR}_i(s_{\alpha^1})$. The same argument then implies that (A.65) holds for all $l \leq L$. ☐

Proof of Proposition 9.1

Proof. The proof is by construction.

For any pair $(u, u') \in U \times U$ let $L_{u,u'} \in \mathrm{O}^{2n}$ be so that

$$\mathrm{span}(L_{u,u'}u, L_{u,u'}u') = \mathrm{span}(e_1, e_2), \tag{A.66}$$

where $e_1 \coloneqq (1, 0, \dots) \in U$ and $e_2 \coloneqq (0, 1, 0, \dots) \in U$. Then we have that

$$\varphi(u) = \underbrace{\exp\left(\theta L_{u,u'}^\top \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & & \\ \vdots & \vdots & & \ddots & \\ 0 & 0 & & & 0 \end{pmatrix} L_{u,u'}\right)}_{\coloneqq E_{u,u'}} \varphi(u') \tag{A.67}$$

where

$$\theta \coloneqq \arccos\left(\langle \varphi(u), \varphi(u') \rangle\right). \tag{A.68}$$

Denote the resulting map $E\colon (u, u') \mapsto E_{u,u'}$. Then for any $\sigma > 0$, and $y, \varepsilon \in U$,

$$\varphi(y + \sigma\varepsilon) = e^{f_{y,\varepsilon}(\sigma) E_{y+\varepsilon,y}} \varphi(y), \tag{A.69}$$

where

$$f_{y,\varepsilon}(\sigma) \coloneqq \frac{\arccos(\langle \varphi(y), \varphi(y + \sigma\varepsilon) \rangle}{\arccos(\langle \varphi(y), \varphi(\varepsilon) \rangle)}. \tag{A.70}$$

We now define a mapping that converts a pair of additive error terms $(\varepsilon_1, \varepsilon_2)$ into a pair of rotation matrices, for every $s \in S$ and player i,

$$\Gamma_{s,i}\colon (\varepsilon_1, \varepsilon_2) \mapsto (E_{s,s-\varepsilon_i}, E_{s+(\varepsilon_{-i}-\varepsilon_i),s-\varepsilon_i}) \tag{A.71}$$

Define the induced distribution over pairs of rotations

$$\tilde{\nu}_s \coloneqq \mu_1 \times \mu_2 \circ \Gamma_s^{-1}. \tag{A.72}$$

For every $\epsilon, \sigma > 0$, define

$$\hat{\mathcal{R}}_\epsilon^\sigma(B) \coloneqq \{u \in \mathcal{N}_\epsilon(S) : \overline{\text{ICR}}^\sigma(u) = B\}, \tag{A.73}$$

where $\mathcal{N}_\epsilon(S) \coloneqq \bigcup_{s \in S}\{u \in U : ||u - s||_2 < \epsilon\}$. Moreover, let

$$\hat{\mathcal{R}}_\epsilon^\sigma(B|S) \coloneqq \{\varphi(u) : u \in \hat{\mathcal{R}}_\epsilon^\sigma(B)\}. \tag{A.74}$$

Then we have that

$$\begin{aligned} &\lim_{\epsilon \downarrow 0} \lim_{\sigma \downarrow 0} \tilde{\nu}_s(\{(E_1, E_2) : \exp(f_{e^{-E_i}s, s-e^{-E_i}s}(\sigma)(E_{-i} - E_i))s \in \hat{\mathcal{R}}_\epsilon^\sigma(B|S)\}) \\ &\qquad = \lim_{\epsilon \downarrow 0} \lim_{\sigma \downarrow 0} \mu_1 \times \mu_2(\{(\varepsilon_1, \varepsilon_2) : s + \sigma(\varepsilon_{-i} - \varepsilon_i) \in \hat{\mathcal{R}}_\epsilon^\sigma(B)\}), \end{aligned} \tag{A.75}$$

and so the result follows. □