

Essays on Information Economics

by
Rafael Veiel

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN ECONOMICS

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

© 2025 Rafael Veiel. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Rafael Veiel
Department of Economics
April 18, 2025

Certified by: Stephen E. Morris
Peter A. Diamond Professor of Economics, Thesis Supervisor

Certified by: Robert M. Townsend
Elizabeth and James Killian Professor of Economics, Thesis Supervisor

Accepted by: Isaiah Andrews
Professor of Economics
Chairman, Departmental Committee on Graduate Theses

Essays on Information Economics

by

Rafael Veiel

Submitted to the Department of Economics
on April 18, 2025 in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN ECONOMICS

ABSTRACT

This thesis contains 5 chapters. Every chapter deals with the question of how information affects equilibrium behavior in strategic problems.

Chapter 1 is my job market paper “Limits of Global Games.” It considers the impact of information on equilibrium multiplicity in two-player games of strategic complementarities. 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 chapter 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.

Chapter 2 is co-authored with Olivier Gossner and is titled “Strategic Type Spaces.” We provide a strategic foundation for information: in any given game with incomplete information we define strategic quotients as information representations that are sufficient for players to compute best-responses to other players. We prove 1/ existence and essential uniqueness of a minimal strategic quotient called the Strategic Type Space (STS) in which a type is given by an interim correlated rationalizability hierarchy together with the set of beliefs over other players’ types and nature that rationalize this hierarchy 2/ that this minimal STS is a quotient of the universal type space and 3/ that the minimal STS has a recursive structure that is captured by a finite automaton.

Chapter 3 is also co-authored with Olivier Gossner and is titled “Information Design for Rationalizability.” We study (interim correlated) rationalizability in games with incomplete information. For each given game, we show that a simple and finitely parameterized class of information structures is sufficient to generate every outcome distribution induced by general common prior information structures. In this parameterized family, players observe signals of two kinds: A finite signal and a common state with additive, idiosyncratic noise. We characterize the set of rationalizable outcomes of a given game as a convex polyhedron.

Chapter 4 is co-authored with Stephen Morris and Dirk Bergemann and is titled “A Strategic Topology on Information Structures.” Two information structures are said to be close if, with high probability, there is approximate common knowledge that interim beliefs are close under the two information structures. We define an “almost common knowledge topology” reflecting this notion of closeness. We show that it is the coarsest topology generating continuity of equilibrium outcomes. We show that finite information structures are dense in the almost common knowledge topology and thus it is without loss to restrict attention to finite information structures in information design problems.

Finally, chapter 5 is a short note describing an information aggregation mechanism that can be used by players before playing a game of strategic complementarities under incomplete information. In such a game, players may have an incentive to share overly optimistic information with other players, thus inducing them to play higher actions. In this mechanism, players trade a token before playing the game. Players who want to communicate good news must purchase this worthless token and burn resources. The note shows that players only need to observe the market clearing price that arises from the token trades to aggregate their private information. Each element in a player’s private information set is encoded as a prime in the prime factorization of the market clearing price. The element that is contained in every player’s information set is identified as the prime with the highest multiplicity.

JEL Classification Codes: C72, D82

Thesis supervisor: Stephen E. Morris

Title: Peter A. Diamond Professor of Economics

Thesis supervisor: Robert M. Townsend

Title: Elizabeth and James Killian Professor of Economics

Acknowledgments

I would like to thank my advisors Stephen Morris and Rob Townsend for their continued support and guidance during my studies. Each of them helped me develop a very different set of skills.

Stephen always made sure I didn't lose sight of the big questions. He is patient, kind, and incredibly generous with his time. He encouraged me to push myself and work on hard problems. At times, when I felt like I was poking aimlessly in the dark, he saw potential in my ideas and encouraged me to keep going, and so I never lost hope. He saw my strengths and helped me use them in my research. He taught me to be patient with myself and to prioritize depth and quality above all else. He is a true inspiration for me and a fantastic mentor.

Rob has taught me that good economic research is always grounded in reality. He showed me that applicability and relevance need not come at the expense of depth and originality. In fact, he always encouraged me to be as creative as possible and follow my instinct. Rob encouraged me to study and think carefully about foundational economic questions involving the nature and behavior of money flows, commitment, the implementation of decentralized exchange, and coordination. Rob and I talked and speculated on many of these issues. My thesis focuses mainly on the issue of coordination, but with more time, I hope to work on all of these questions. When working with him, originality comes first, and that's exciting.

I would also like to thank Olivier Gossner who has been there from the beginning. His continued support and guidance helped me get into MIT and helped me become an economic theorist. He encouraged me to learn mathematics and expand my horizons. He is open-minded and sees potential in an idea that would really only take shape after many years of work. My thesis would not have been possible without him, and for that I am grateful. I also thank Ivan Werning, who provided me with guidance when I was unsure about what field to focus on in my last years of the Ph.D. I was very inspired by his way of thinking about inflation. The interactions I have had with Ivan will certainly shape my research in the coming years. I also want to thank Ian Ball for helping me improve my presentation skills and his support during the job market.

Finally, I would like to thank my friends and office mates Adam, Jaume, Isabel, Moritz, Ahmet, Marc, Bumsoo, Nagisa, Nancy, Anne and Sylvia; and my family Axel, Roxana, Alicia and Nils who have all been incredibly supportive.

Contents

<i>List of Figures</i>	11
<i>List of Tables</i>	13
1 Limits of Global Games	15
1.1 Introduction	15
1.2 Global Game Model	20
1.2.1 Model	21
1.2.2 Rationalizability and Limit Selection	23
1.3 Illustrative Example: Binary-Action	26
1.4 Main Result	30
1.5 Limit Uniqueness: Proof	32
1.5.1 Indifference Constraints	32
1.5.2 Algebraic Description of Collisions	34
1.5.3 Aligned Incentives	37
1.5.4 Limit Uniqueness for Quasi Concave, Supermodular Games	39
1.5.5 Potential Games	40
1.6 Examples	40
1.6.1 Example: Large Asymmetric Investment Games	40
1.6.2 Example: Symmetric Investment Game	41
1.6.3 Example: Is Limit Uniqueness Robust to Redistribution?	43
1.7 Discussion	44
1.7.1 Spherical State Space and Rotational Noise	44
1.7.2 Additively Separable Noise	44
1.7.3 Related Literature	45
1.7.4 Conclusion	46
2 Strategic Type Spaces	47
2.1 Introduction	47
2.1.1 Related Literature	48
2.2 Preliminaries and Notations	49
2.3 Strategic Type Spaces	49
2.3.1 Interim Correlated Best-Replies	50
2.3.2 Strategic Type Spaces	50
2.3.3 The Strategic Type Space	52
2.4 Finite Representation of STS	54

2.4.1	Automaton Representation	57
2.4.2	Example	57
2.5	Discussion	59
3	Information Design for Rationalizability	61
3.1	Introduction	61
3.2	A Technology Coordination Example	64
3.2.1	Strategic Automaton	65
3.2.2	SCAMP	67
3.2.3	Additive, Idiosyncratic Noise Representation	69
3.3	Model	70
3.4	Rationalizable Hierarchies	71
3.4.1	Revelation principle	71
3.4.2	Strategic Automaton	72
3.5	Rationalizable Distributions	72
3.5.1	SCAMP	72
3.5.2	SCAMP Sufficiency	73
3.5.3	Characterization of Rationalizable Outcomes	80
3.5.4	Sufficiency of Additive Noise Information Structures	81
3.6	Information Disclosure in Priority Auctions	82
3.6.1	Model	82
3.6.2	Complete Information Benchmark	83
3.6.3	Upper Bound: Almost full extraction	83
3.6.4	SCAMP/ANIS	84
3.7	Discussion	87
3.7.1	Structure and Complexity of SCAMP	87
3.7.2	Solution Concepts	88
3.7.3	Relation to Literature	88
4	A Strategic Topology on Information Structures	91
4.1	Introduction	91
4.1.1	Motivation and Results	91
4.1.2	Alternative Topologies and Related Literature	92
4.2	Model	93
4.2.1	Setting and Base Game	93
4.2.2	Information Structures	94
4.2.3	Solution Concept: Belief-Invariant Bayes Correlated Equilibrium	95
4.3	Main Result	96
4.3.1	The Approximate Common Knowledge Topology	96
4.3.2	Continuity of Equilibrium Outcomes	98
4.3.3	Main Result	98
4.3.4	Denseness of Finite Information Structures	106
4.3.5	Denseness of Simple Information Structures	107

4.4	General Information Structures and Bayes Nash Equilibrium	108
4.4.1	General Information Structures	108
4.4.2	Solution Concepts	108
4.4.3	Existence of Equilibria	110
4.4.4	Varying Solution Concepts and Information Structures	110
4.4.5	Correlation and Non-Redundant Information Structures	111
4.4.6	Discussion	113
4.5	Information Design	114
4.5.1	Applications	115
4.5.2	Proof of Proposition 13	115
4.6	Alternative Formulations and Related Literature	116
4.6.1	Interim Topologies	116
4.6.2	Value-Based Topology	118
4.6.3	Join Measurability	119
4.6.4	Improper Priors and Completeness	119
5	Information Aggregation Mechanisms	123
5.1	Introduction	123
5.1.1	Model	123
5.2	The Mechanism	126
5.3	Discussion	129
5.3.1	Natural Encryption	129
5.3.2	Related Literature	129
A	Proofs Chapter 1	131
A.1	Appendix: Omitted Proofs	131
A.2	Risk-Dominance and Limit Multiplicity	139
A.2.1	Example: Coordination on Surplus Splitting	139
A.2.2	Risk-Dominance	140
A.2.3	Limit Multiplicity	142
A.3	Topological Properties of Limit Selection	144
B	Proofs Chapter 2	147
B.1	Appendix	147
B.1.1	Preliminaries	147
B.1.2	Strategic Type Spaces	148
C	Proofs Chapter 3	153
C.1	Appendix	153
C.1.1	Theorem 5: Revelation Principle	153
C.1.2	Theorem 6: Sufficiency of SCAMP	154

D Proofs Chapter 4	159
D.1 Appendix	159
D.1.1 Proof of Proposition 1	159
D.1.2 Proof of Proposition 2	160
D.1.3 Proof of Proposition 9	161
D.1.4 Proof of Proposition 10	162
D.1.5 Proof of Lemma 1	162
D.1.6 Proof of Lemma 2 (Iterative Scoring Rule)	163
D.1.7 Common Belief Invariance	164
D.1.8 Proof of Proposition 14	164
D.1.9 Interim Strategic Topology	165
D.1.10 Upper Hemi Continuity of Exact BIBCE	166
D.1.11 Example Existence of BIBCE	167

List of Figures

1.1	General Signal in Carlsson and Van Damme (1993a).	15
1.2	Symmetry of the Noise Distribution	18
1.3	Global game with additively separable noise and spherical global game with rotational noise.	21
1.4	Illustration of spherical global game with two actions.	23
1.5	Rationalizable sets and their boundaries.	25
1.6	Illustration of rationalizable sets and their boundaries in binary actions case.	26
1.7	Illustration of collisions.	26
1.8	Symmetry of boundaries.	28
1.9	Path between game with best-response cycle and game with better response cycle.	39
1.10	Payoffs for different choices of α_i .	41
1.11	Budget Balanced Intervention cannot generate “opposite skewness” property.	43
2.1	For all beliefs in the region between the shaded triangle (excluded) and the sub simplex spanned by $(a, -1)$, $(b, -1)$ and $(b, 1)$, player i ’s best response is always b . For all beliefs in the region between the shaded triangle (excluded) and $(a, 1)$, player i ’s best response is always a and on the shaded triangle all beliefs induce both actions a and b as best response.	58
2.2	STS Automaton.	58
3.1	Automaton for one player in the technology example. There are 4 states and each state contains an action set. The initial state is on the left (player 1) or on the right (player 2). Double circled states are terminal ones.	66
3.2	Automaton for both players in the technology example. Each state contains an action set for player 1 (top) and for player 2 (bottom). Arrows indicate possible transitions.	66
3.3	SCAMP generating point distributions. When a single arrow leaves a state, this arrow has probability 1. Transitions may depend on the state of nature $k \in \{G, B\}$.	68
3.4	Point rationalizable payoffs (in blue), correlated equilibrium payoffs (in red) and feasible payoffs (in grey) in the coordination game.	69
3.5	Canonical Prior as Process on SCAMP automaton.	74
3.6	Averaging process to SCAMP.	74
3.7	SCAMP Automaton for $k < \bar{k}$	84

3.8	SCAMP Automaton for $k = \bar{k}$, where $\hat{k} := \bar{k} - 1$.	85
4.1	Sets $\hat{T}_\varepsilon(P, P')$ and ε -supports of P and P' .	97
4.2	Sets $\hat{T}_\varepsilon(P, P'), C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$ and ε -supports of P and P' .	99
4.3	Infection argument.	101
4.4	Domain of Grid Game.	105
5.1	Signal distributions for agents i_1, i_2 and i_3 conditional on the true state being θ .	125
D.1	Separation Argument.	165

List of Tables

1.1	Better Response Cycle.	16
1.2	Limit Uniqueness for Supermodular Payoffs.	20
1.3	Matching Pennies Game.	29
1.4	Coordination Game.	29
A.1	Coordination Game with Surplus-Splitting Component.	140
A.2	Coordination Game induced by optimistic conjectures over a, b	140

Chapter 1

Limits of Global Games

1.1 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 (1993a) 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. In its most general form, a player’s signal can be expressed as a slightly perturbed payoff matrix. Figure 1.1 shows an example of such a signal when actions are given by $\{a, b\}$ for both players: The first term in the sum is the unperturbed payoff matrix and the second term is noise, where $\sigma > 0$ allows us to scale the size of the noise.

$$S = \begin{array}{c} a \\ b \end{array} \begin{array}{cc} a & b \\ \hline 1, 1 & 1, 0 \\ \hline 0, 1 & 2, 2 \end{array} + \sigma \cdot \begin{array}{cc} \varepsilon_1^{a,a}, \varepsilon_2^{a,a} & \varepsilon_1^{a,b}, \varepsilon_2^{a,b} \\ \hline \varepsilon_1^{b,a}, \varepsilon_2^{b,a} & \varepsilon_1^{b,b}, \varepsilon_2^{b,b} \end{array}$$

Figure 1.1: General Signal in Carlsson and Van Damme (1993a).

Players have a prior over \mathbb{R}^8 , describing their ex-ante uncertainty about to the first term, and the noise is drawn from \mathbb{R}^8 , independently of the first. The induced game of incomplete information is called a “global game.” Limit uniqueness means that, for any small enough value of σ , players play a unique action at approximately all signal realizations, in every equilibrium of the corresponding Bayesian 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. With two actions, the global game selection does not depend on the distribution of the noise. Even if players’ uncertainty is restricted to a one-dimensional family of payoff matrices. For games of strategic complementarities with many actions this is not true. Frankel et al. (2003) prove limit uniqueness when players see a noisy signal of a one-dimensional parameter of the game but show that the selection may depend of the details of the noise.

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?

Main Result This paper answers that question in a two-player, many-actions generalization of the global game model in Carlsson and Van Damme (1993a) 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. Our main result can be stated as follows:

Main Result: *Limit uniqueness holds for a concave supermodular payoff matrix if and only if there are no “relevant better response cycles.”*

To make this statement precise, we need to define what we mean by “better response cycles” and what we mean by “relevant.” To illustrate this result, let’s consider the following supermodular payoff matrix¹, where both players have actions $\{a, b, c, d\}$.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	9 7	4 4	3 -1	-1 -6
<i>b</i>	8 6	7 7	10 3	7 -1
<i>c</i>	4 8	10 10	16 7	14 4
<i>d</i>	0 4	8 10	15 15	15 19

Table 1.1: Better Response Cycle.

With complete information, this game admits three pure-strategy Nash equilibria: (a, a) , (c, b) and (d, d) . The game also has what we call a better response cycle: A cyclic sequence of

¹Consider the alphabetic order on actions for each player.

action profiles, where for any two consecutive action profiles, one player is switching to an action that strictly improves her payoffs, while the action of her opponent is held fixed. As we move along the sequence, the player who improves her payoffs alternates. The following is thus a better response cycle:

$$(b, a) \rightarrow (b, b) \rightarrow (d, b) \rightarrow (d, c) \rightarrow (c, c) \rightarrow (c, a) \rightarrow (b, a)$$

Better response cycles determine the complexity of a game’s Nash equilibria: If there are no better response cycles then a Nash equilibrium can always be computed by simple, unilateral improvement dynamics.

We will say that a better response cycle is relevant if the action sets involved in the cycle are contained in a best-reply closed set of actions and players don’t have to be too certain that their opponent is playing an action involved in the cycle for them to best-reply an action involved in the cycle. Playing an action in the better response cycle is thus not very risky. This property is a generalization of Risk-Dominance introduced by [Harsanyi and Selten \(1988\)](#) as a property of equilibria.

Suppose now that players observe the payoff matrix [1.1](#) with some noise, as illustrated in [Figure 1.1](#). As we discuss later, the better response cycle described above is relevant, and so our result implies that limit uniqueness will fail. The “only if” part of our result means that the global game approach actually relies on a lot of structure - the absence of relevant cycles - to serve as a reliable selection device when we move beyond two actions and one-dimensional noise. One-dimensional noise, two actions or the absence of better-response cycles are stylized assumptions that deliver tractable models of coordination. However, in many cases it is hard to rule out a more complex model based on economic arguments alone. These are the limits of the global game approach.

What features give rise to better response cycles? In the payoff matrix [1.1](#), the cycle arises due to a simple asymmetry in payoffs between the players: For the [row player](#), overshooting - playing an action above her best-reply - is *better* than undershooting. However, for the [column player](#) overshooting is *worse* than undershooting. We show in [Section 1.6](#) that better response cycles also arise in symmetric games. [Monderer and Shapley \(1996\)](#) show that payoffs without better response cycles are precisely the payoffs that admit a generalized ordinal potential.

Global Game Model To prove our result, we 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 \(1993a\)](#) 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

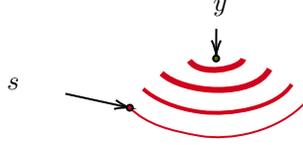


Figure 1.2: Symmetry of the Noise Distribution

any orthogonal change of basis, and both players have the same noise distribution. The former symmetry assumption is illustrated in Figure 1.2 below. The gray sphere represents the collection of all scaled payoff matrices. As described above, a player’s signal s is given by random rotation of a point on the sphere, y . Given y , the distribution of signals that a player might observe is a distribution on the sphere. Symmetry then implies that the probability of a rotation thus only varies with the magnitude of the rotation but not the direction. We illustrate this in Figure 1.2 by the constant thickness of each circle.

In our spherical model, noise is applied by a linear transformation (i.e. a rotation matrix) rather than an affine transformation as in the additively separable case. The implications of symmetry on players’ beliefs and their optimal strategies in a spherical global game can be described and analyzed in a tractable way. We therefore view our modeling choice here as serving the technical purpose of making it as tractable as possible to exploit the symmetry assumption in our model. While we use symmetry of the noise distribution heavily in our analysis, we think our results would hold more generally. In the binary action environment, it is known that as the noise shrinks, players’ beliefs become closer to a uniform distribution. In this case, symmetry holds in the limit and imposing it away from the limit would be without loss of generality. Carlsson and Van Damme (1993a) therefore don’t require symmetry to establish limit uniqueness. We do not establish a similar result in the multi-dimensional model but still think it holds.

Third, our global game model can be viewed as having no payoff uncertainty. 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. As the noise vanishes, this assumption loses its bite as the signal s and the payoff matrix y become very close.

Finally, we focus on interim correlated rationalizability (ICR), an incomplete information version of rationalizability introduced by Dekel et al. (2007a). 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 that the regions that collide are mirror images of each other. We express this property algebraically: the 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 can state this property in terms of derivatives of the above-mentioned multilinear form and show that aligned incentives are both necessary and sufficient for achieving limit uniqueness. For quasi concave and 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 \(1993a\)](#), 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 \(2003a\)](#) for a survey. 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 many-dimensional 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. [Table 1.2](#) summarizes all that is known about limit uniqueness when payoffs are supermodular. Since the literature does not offer tractable techniques to study limit uniqueness in general multi-dimensional global games, we will need to get creative.

	1-Dimensional Noise	Multi - Dimensional Noise
Two Actions	Carlsson and Van Damme (1993a) Limit Uniqueness holds	Carlsson and Van Damme (1993a) Limit Uniqueness holds
Many Actions	Frankel et al. (2003) Limit Uniqueness holds	Oury (2013) Limit Uniqueness may fail

Table 1.2: Limit Uniqueness for Supermodular Payoffs.

Organization of the Paper The rest of the paper is organized as follows: Section 1.2 introduces the spherical global game model. Section 1.3 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 \(1993a\)](#) expressed in the language of this paper. Section 1.4 contains a statement of our main result and the definitions required to state it. Section 1.5 provides the main arguments of the proof. Section 1.6 illustrates the failure of limit uniqueness in examples. Section 1.7 discusses the modeling assumption and the connection to the literature. Details of the proofs can be found in Appendix A.1.

1.2 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 1.3 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 \(1993a\)](#). 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 $s = e^{\sigma E}y$, where $e^{\sigma E}$ represents a random rotation matrix, where the rotation is of magnitude σ .

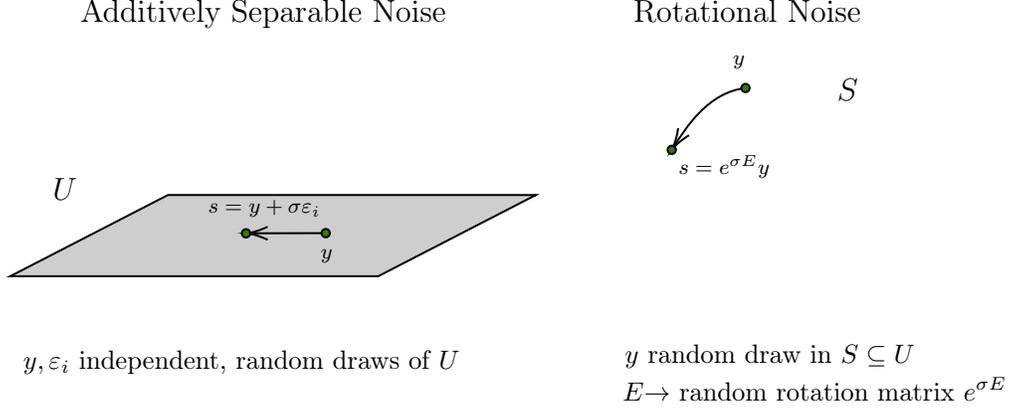


Figure 1.3: Global game with additively separable noise and spherical global game with rotational noise.

In Section 1.7.2 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 $O^n \subseteq \mathbb{R}^{n \times n}$ denote the set of orthogonal² $n \times n$ matrices and let $SO^n \subseteq O^n$ denote the set of orthogonal matrices with determinant equal to one. Note that these are rotation matrices. Every rotation matrix $R \in SO^n$ can be written as a matrix exponential³ $R = e^E$ for some skew-symmetric⁴ matrix E .

1.2.1 Model

We now introduce the global game model, which is spherical. We consider a pair of players $\{1, 2\}$, where each player⁵ 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 $U := \mathbb{R}^{A_1 \times A_2} \times \mathbb{R}^{A_2 \times A_1}$,

$$S := \{u = (u_1, u_2) \in U : \|u\|_2 = 1\}. \quad (1.1)$$

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 $n = |A_1 \times A_2|$,

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

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

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

⁵We 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. $-i \neq i$.

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.

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

$$\mathcal{E} := \{E \in [-1, 1]^{2n \times 2n} : E^\top = -E\}. \quad (1.2)$$

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} .

Bayesian Game 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

$$s = e^{\sigma E_i} y, \quad (1.3)$$

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.

Each signal, a point in S , describes payoffs for both players at every action profile. To simplify the analysis, we will study the Bayesian game, where players have common prior ν_σ , each player i receives a private signal s and i 's payoffs when (a_i, a_{-i}) is being played are given by $s_i^{a_i, a_{-i}}$. 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 $\sigma \rightarrow 0$, where the analysis simplifies.

Figure 1.4 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 1.4. 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.

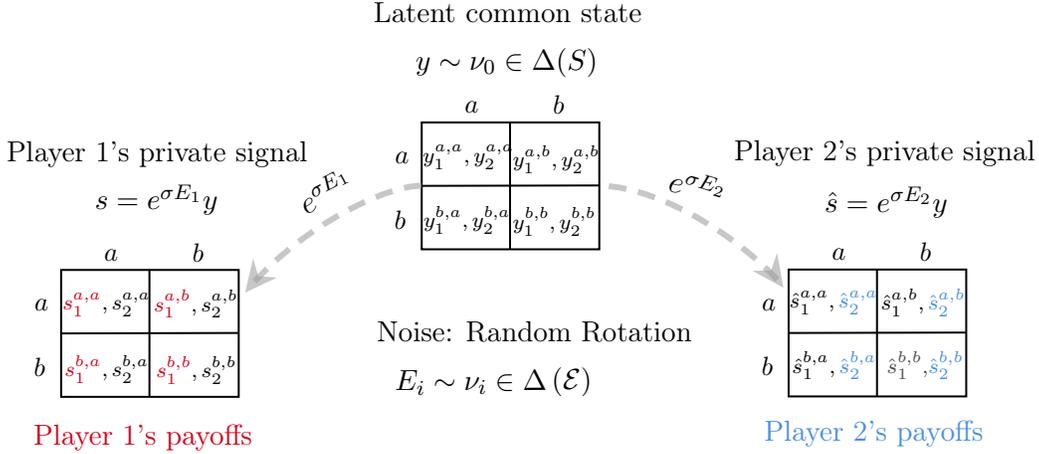


Figure 1.4: Illustration of spherical global game with two actions.

Definition 1 (Symmetry). Say that ν is symmetric if $\nu_1 = \nu_2$ and for every $E \in \mathcal{E}$, every player i and every $X \in O^n$,

$$f_{\nu_i}(X^\top E X) = f_{\nu_i}(E), \quad (1.4)$$

where f_{ν_i} is the density of ν_i .

We will assume that ν is symmetric throughout. For a more intuitive grasp on symmetry, it may be helpful to describe what symmetry of ν implies for the distribution of signals ν_σ , conditional on a latent common state y . In this case, symmetry implies that the density of signals for each player is invariant to all orthogonal transformations that preserve y . Figure 1.2 in the introduction provides an illustration of this property: the probability of a signal stays the same as you apply any rotation or reflection about the axis defined by y . Examples include the uniform distribution or the normal distribution on $\mathbb{R}^{2n \times 2n}$ restricted to \mathcal{E} .

1.2.2 Rationalizability and Limit Selection

We now state the definition of Interim Correlated Rationalizability (ICR) as defined in Dekel et al. (2007a). 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.

⁶Another 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.

Best Response Let the collection of non-empty action sets of player i be denoted by $\mathcal{A}_i := 2^{A_i} \setminus \emptyset$. 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

$$\text{br}_i(p|s) := \arg \max_{a_i} \sum_{a_{-i} \in A_{-i}} p(a_{-i}) s_i^{a_i, a_{-i}}. \quad (1.5)$$

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

$$\nu_\sigma \circ \kappa_{-i}(a_{-i}|s) := \int_{S \times S} \kappa_{-i}(a_{-i}|y, s') d\nu_\sigma(y, s'|s), \quad \forall a_{-i} \in A_{-i}. \quad (1.6)$$

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

$$\text{BR}_i^\sigma(\kappa_{-i}|s) := \text{br}_i(\nu_\sigma \circ \kappa_{-i}(\cdot|s)|s). \quad (1.7)$$

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

$$R_i^{m,\sigma}(s) := \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^{m,\sigma}} \text{BR}_i^\sigma(\kappa_{-i}|s), \quad (1.8)$$

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

Definition 2 (Rationalizability). R is given by

$$R_i^\sigma(s) := \bigcap_{m=0}^{\infty} R_i^{m,\sigma}(s). \quad (1.9)$$

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

$$\liminf_{\sigma \downarrow 0} R_i^\sigma(s) = R_i(s). \quad (1.10)$$

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

$$\lim_{\sigma \downarrow 0} \nu_\sigma(\{s \in O : \exists i \text{ s.t. } |R_i^\sigma(s)| > 1\}) = 0. \quad (1.11)$$

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

$$\mathcal{R}^\sigma(B) := \{s \in S : R^\sigma(s) = B\}. \quad (1.12)$$

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

$$\partial \mathcal{R}^\sigma(B, B') := \overline{\mathcal{R}^\sigma(B)} \cap \overline{\mathcal{R}^\sigma(B')}, \quad (1.13)$$

where $\overline{\mathcal{R}^\sigma(\cdot)}$ denotes the topological closure of $\mathcal{R}^\sigma(\cdot)$ in S . Figure 1.5 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.

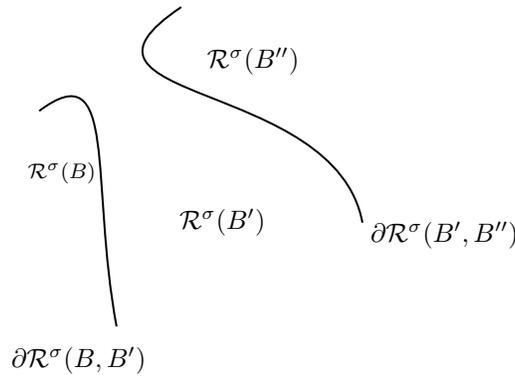


Figure 1.5: Rationalizable sets and their boundaries.

To illustrate these concepts in a more concrete setting, consider again the binary action environment. Figure 1.6 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

⁷As we show later, upper-hemi continuity properties of R 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.

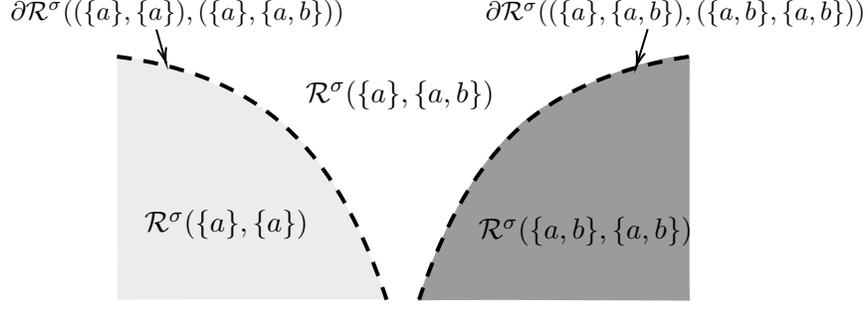


Figure 1.6: Illustration of rationalizable sets and their boundaries in binary actions case.

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 5 (Collisions). *For any collection of action set pairs $Z = \{(B_1^1, B_2^1), \dots, (B_1^m, B_2^m)\}$, define the R-collision of Z ,*

$$\mathcal{C}(Z) := \bigcap_{(B, B') \in Z} \partial \mathcal{R}(B, B'), \quad (1.14)$$

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

Figure 1.7 illustrates a collision between the boundary separating $(\{a\}, \{a\})$ from $(\{a\}, \{a, b\})$ and the boundary separating $(\{a\}, \{a, b\})$ from $(\{a, b\}, \{a, b\})$.

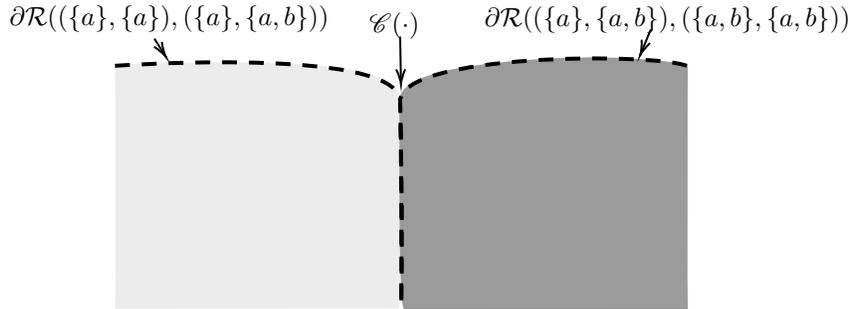


Figure 1.7: Illustration of collisions.

1.3 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 ,

$$S = \{u \in \mathbb{R}^8 : \|u\|_2 = 1\}. \quad (1.15)$$

With additively separable noise rather than rotational noise, this case⁸ was studied in [Carlsson and Van Damme \(1993a\)](#). 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 $R_i^\sigma(s)$. The partition of rationalizable sets induced by the correspondence $R^\sigma: S \rightarrow \mathcal{A}_1 \times \mathcal{A}_2$ allows each player i to compute her beliefs on the rationalizable action sets of her opponent:

$$P_i^\sigma(B_{-i}|s) = \nu_0 \times \nu \left(\{(y, E_i, E_{-i}) : B_{-i} = R_{-i}(e^{\sigma E_{-i} y})\} \mid e^{\sigma E_i y} = s \right). \quad (1.16)$$

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 1](#) that the indifference constraint

$$g_i^{1,2,\sigma}(s) = P_i^\sigma(\{a\}|s)(s_i^{1,1} - s_i^{2,1}) + P_i^\sigma(\{b\}|s)(s_i^{1,2} - s_i^{2,2}) = 0, \quad (1.17)$$

must hold at the boundary where player i switches from $\{a\}$ to $\{b\}$. Condition [\(1.17\)](#) 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 2](#) we show that at any point where boundaries collide, rationalizable sets are 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 1.8](#). 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 $\{b\}$.

⁸We discuss the connection between additively separable noise and rotational noise in [Section 1.7.2](#).

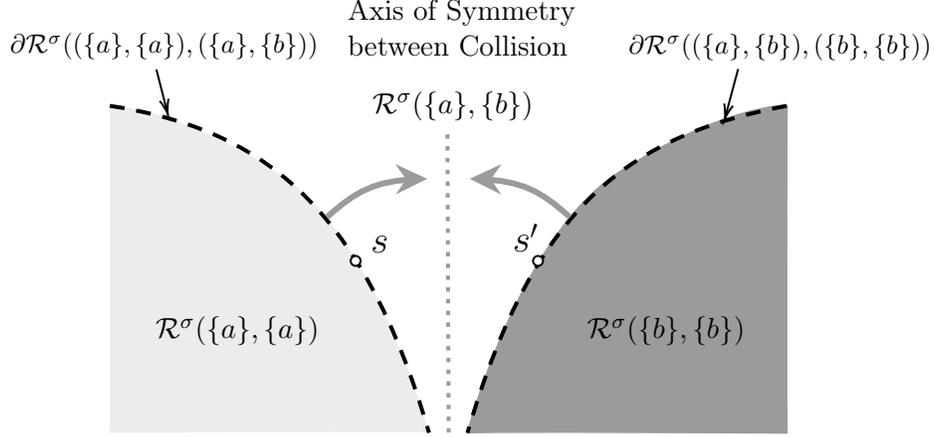


Figure 1.8: Symmetry of boundaries.

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

$$P_2^\sigma(\{a\}|s) = P_1^\sigma(\{b\}|s'), \quad (1.18)$$

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 (1.19) becomes,

$$\lim_{\sigma \downarrow 0} P_2^\sigma(\{a\}|s^*) = \lim_{\sigma \downarrow 0} P_1^\sigma(\{b\}|s^*). \quad (1.19)$$

Risk-Dominance as Multilinear Form Combining the symmetry constraint (1.19) with the necessary indifference constraints (1.17), we obtain the binary action version of Proposition 2: Points on a collision are contained in the zeros of a multilinear form given by

$$\delta(s) = \sum_{\alpha=a,b} \prod_{i=1,2} (s_i^{a,\alpha} - s_i^{b,\alpha}) = 0. \quad (1.20)$$

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 (1.20) is equivalent to both actions being risk-dominant⁹: The product of the deviation losses is exactly equal for both actions. In Proposition 2 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 (1993a) have shown that condition (1.20) is also sufficient to characterize the selection in the region of S where payoffs are supermodular.

⁹See for instance Harsanyi and Selten (1988).

Aligned Incentives To obtain a sufficient condition for collisions we study the local stability properties of the zeros of the multilinear form (1.20). We consider the derivative at a point satisfying (1.20),

$$\frac{\partial \delta}{\partial (s_i^{a,a} - s_i^{b,a})} = (s_{-i}^{a,a} - s_{-i}^{b,a}). \quad (1.21)$$

We say that s is *critical* if it satisfies (1.20) and *incentives are aligned* if

$$\frac{\partial \delta}{\partial (s_i^{a,a} - s_i^{b,a})} \frac{\partial \delta}{\partial (s_{-i}^{a,a} - s_{-i}^{b,a})} > 0. \quad (1.22)$$

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 3, we show that aligned incentives and criticality are necessary and sufficient for points to lie on a collision.

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 1.3: 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 7 we show that the presence of risk dominant¹⁰ 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 (1993a).

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

Table 1.4: Coordination Game.

In Theorem 1 we state our characterization of limit uniqueness for supermodular games with quasi concave payoffs.

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

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

$$\begin{aligned} D_i(a) &= \{s \in S : s_i^{a,\alpha} - s_i^{b,\alpha} > 0, \forall \alpha \in \{a, b\}\}, \\ D_i(b) &= \{s \in S : s_i^{a,\alpha} - s_i^{b,\alpha} < 0, \forall \alpha \in \{a, b\}\}. \end{aligned} \quad (1.23)$$

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

$$R_i^\sigma(s) = \{a_i\}, \quad R_{-i}^\sigma(s) = \arg \max_{a_{-i} \in A_{-i}} s_{-i}^{a_{-i}, a_i}. \quad (1.24)$$

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

$$Z := S \setminus \cup_i (D_i(a) \cup D_i(b)). \quad (1.25)$$

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

$$R_i(s) = \begin{cases} \{a\}, & \text{if } \delta(s) > 0, \\ \{b\}, & \text{if } \delta(s) < 0, \\ A_i, & \text{otherwise.} \end{cases} \quad (1.26)$$

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

1.4 Main Result

In this section we provide all definitions that are required to state our main result: Quasi concave, supermodular payoff matrices, Best-Reply Closed action sets, Better Response Cycles and Risk-Dominant action sets.

Quasi concave, supermodular payoff matrices Suppose the actions of each player are ordered, $A_i = \{1, \dots, N_i\}$, for every player i , where $N_i \in \mathbb{N}$. Let \preceq denote the product order on $A_1 \times A_2$. Payoffs $s \in S$ are *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 $(a'_i, a'_{-i}) \preceq (a_i, a_{-i})$,

$$s_i^{a'_i, a_{-i}} - s_i^{a_i, a_{-i}} < s_i^{a'_i, a'_{-i}} - s_i^{a_i, a'_{-i}}, \quad \forall i = 1, 2. \quad (1.27)$$

That is, switching to higher actions becomes more profitable when the other player also plays higher actions. Say that $s \in S$ is quasi concave¹¹ if for every player i , every $a_{-i} \in A_{-i}$ and every $a_i, a'_i \in A_i$ so that $s_i^{a_i, a_{-i}} \leq s_i^{a'_i, a_{-i}}$, then any $a''_i \in A_i$ so that $\min\{a_i, a'_i\} < a''_i < \max\{a_i, a'_i\}$,

$$s_i^{a''_i, a_{-i}} > s_i^{a_i, a_{-i}}. \quad (1.28)$$

Let $S^* \subseteq S$ denote the collection of quasi concave, supermodular payoffs in S . Since both conditions are described in terms of strict inequalities involving the entries of points in S , it follows immediately that S^* has a non-empty interior.

¹¹The correct name for this property is strict quasi-concavity, but to alleviate notation we will simply call it strict concavity.

Best-Reply Closed (BRC) Set A pair action (sub)sets $B_1 \subseteq A_1, B_2 \subseteq A_2$, is *Best-Reply Closed (BRC)* under payoff matrix $s \in S$ if for every player i ,

$$\bigcup_{p \in \Delta(B_{-i})} \text{br}_i(p_i | s). \quad (1.29)$$

This property states that all of player i 's best-replies to all beliefs that are supported on B_{-i} must equal B_i .¹²

Better and Best Response Cycles A pair action (sub)sets $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$, is *Better Response Cycle* under payoff matrix $s \in S$ if there exists a list of action profiles $(a_1^0, a_2^0), \dots, (a_1^m, a_2^m) \in B_1 \times B_2$, and a player i so that

(i) for every even $n \in \{0, \dots, m-1\}$,

$$s_i^{a_i^n, a_{-i}^n} < s_i^{a_i^{n+1}, a_{-i}^{n+1}} \text{ and } a_{-i}^n = a_{-i}^{n+1}, \quad (1.30)$$

(ii) for every odd $n \in \{0, \dots, m-1\}$,

$$s_{-i}^{a_{-i}^n, a_i^n} < s_{-i}^{a_{-i}^{n+1}, a_i^{n+1}} \text{ and } a_i^n = a_i^{n+1}, \quad (1.31)$$

(iii) $(a_1^0, a_2^0) = (a_1^m, a_2^m)$.

A *Best Response Cycle* under $s \in S$ is pair action (sub)sets $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$, if there exists a list of action profiles $(a_1^0, a_2^0), \dots, (a_1^m, a_2^m) \in B_1 \times B_2$, where $m > 1$ and a player i so that

(i) for every even $n \in \{0, \dots, m-1\}$,

$$a_i^{n+1} \in \text{br}_i(\mathbf{1}_{a_{-i}^n} | s), \quad (1.32)$$

(ii) for every odd $n \in \{0, \dots, m-1\}$,

$$a_{-i}^{n+1} \in \text{br}_{-i}(\mathbf{1}_{a_i^n} | s), \quad (1.33)$$

(iii) $(a_1^0, a_2^0) = (a_1^m, a_2^m) \neq (a_1^1, a_2^1)$.

Risk-Dominant Set A pair action (sub)sets $B = (B_1, B_2)$, where $B_1 \subseteq A_1$ and $B_2 \subseteq A_2$, is *Risk-Dominant* under payoff matrix $s \in S$ if

$$p_1^B(s) + p_s^B(s) \leq 1, \quad (1.34)$$

where $p_i^B(s)$ is the minimal probability player i needs to attach to B_{-i} so that she has a best-reply in B_i :

$$p_i^B(s) := \min \left\{ p_i^B \in [0, 1] : \exists p \in \Delta(A_{-i}) \text{ s.t. } p(B_{-i}) \leq p_i^B \right. \\ \left. \text{br}_i(p | s) \cap B_i \neq \emptyset \right\}. \quad (1.35)$$

¹²Note that this definition is different from that of CURB sets, which only require the best-replies to be contained in B_i .

We generalize the concept of risk-dominance for binary action games from [Harsanyi and Selten \(1988\)](#) to BRC sets. Risk-dominant action sets thus have the following interpretation: playing a best-reply from the risk-dominant set is not very risky if you are concerned that your opponent might play an action outside of the set, but you are optimistic about which action outside of the set will be chosen.

Main Result We are now ready to state the main result.

Theorem 1 (Limit Uniqueness). *Let $O \subseteq S^*$ be open. Limit uniqueness holds on O if and only if for every $s \in O$, no BRC set contains a Risk-Dominant Better Response Cycle.*

We have shown that payoff matrix 1.1 has a better response cycle given by $(\{b, c, d\}, \{a, b, c\})$. The full action sets (A_1, A_2) can readily be shown to be best-reply closed. It thus remains to show that the better response cycle is also risk dominant. For the row player action b is a best-reply to the mixture $\frac{1}{2}a + \frac{1}{2}d$, where d is outside of the cycle. For the column player action b is a best-reply to the mixture $\frac{1}{2}c + \frac{1}{2}a$, where a is outside of the cycle.

1.5 Limit Uniqueness: Proof

We prove this result in three steps: first, we describe the boundaries and collisions in terms of indifference constraints, for any fixed size of the noise. Next, we exploit symmetry to derive an algebraic expression for the limits of the indifference constraints associated to colliding boundaries. We show that collision points must be contained in the zero-set of a multilinear form. Symmetry implies that the multilinear form no longer depends on players' beliefs and is a function of payoffs only. With symmetry, we bypass what makes high-dimensional global games intractable: players' interim beliefs about the rationalizable actions of their opponent.

However, not all points in the zero-set need to be collision points of boundaries separating regions, where singletons are rationalizable. To fully characterize such collisions, and thus characterize limit uniqueness, we derive local stability properties of collision points that are sufficient to identify the singleton collisions. This is a property of payoffs, expressed in terms of the derivatives of the multilinear form derived earlier. It checks if moving away from a collision point in any direction breaks players' indifference constraints in a way that supports a unique best-reply for both players. Finally, we use supermodularity to express this condition in terms of the non-existence of better response cycles.

1.5.1 Indifference Constraints

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.

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 ,

$$P_i^\sigma(B_{-i}|s) := \nu \times \nu_0 \left(\{ (E_i, E_{-i}, y) : B_{-i} = R_{-i}^\sigma(e^{\sigma E_{-i} y}) \} \mid e^{\sigma E_i y} = s \right). \quad (1.36)$$

$P_i^\sigma(B_{-i}|s)$ 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,

$$g_i^{a_i, a'_i, \sigma}(s) := \max_{\kappa_{-i} \in \Sigma_{-i}} \sum_{B_{-i} \in \mathcal{A}_{-i}} P_i^\sigma(B_{-i}|s) \sum_{a_{-i} \in B_{-i}} \kappa_{-i}(a_{-i}|B_{-i}) (s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}}), \quad (1.37)$$

where $\Sigma_{-i} := \{ \kappa : \mathcal{A}_{-i} \rightarrow \Delta(A_{-i}) : \forall a_{-i} \in A_{-i}, \forall B_{-i} \in \mathcal{A}_{-i}, \kappa(a_{-i}|B_{-i}) > 0 \implies a_{-i} \in B_{-i} \}$ 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

$$\mathcal{G}_i^\sigma(a_i, a'_i) := \left\{ s \in S : g_i^{a_i, a'_i, \sigma}(s) = 0 \right\}. \quad (1.38)$$

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 (1.38).

Lemma 1 (Indifference Constraints). *For every $(B, B') \in \mathcal{A}$, 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*

$$s \in \bigcap_{i=1,2} \mathcal{G}_i^\sigma(a_i, a'_i). \quad (1.39)$$

The indifference constraints in (1.38) 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

$$\mathcal{G}_i(a_i, a'_i) := \lim_{\sigma \downarrow 0} \mathcal{G}_i^\sigma(a_i, a'_i). \quad (1.40)$$

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*

$$\mathcal{C}(C) := \bigcap_{i=1,2} \bigcap_{n=1}^{m_i} \mathcal{G}_i(a_i^n, a_i'^n). \quad (1.41)$$

It follows from Lemma 1 that every collision is contained in an IC-collision. Indeed, for every Z there exists

$$C_{Z,i} \subseteq \bigcup_{(B,B') \in Z} (B_i \times B'_i). \quad (1.42)$$

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

Corollary 1. *For every $Z \subseteq \mathcal{A}$ there exists $C = (C_1 \subseteq A_1 \times A_1, C_2 \subseteq A_2 \times A_2)$ so that*

$$\mathcal{C}(Z) \subseteq \mathcal{C}(C). \quad (1.43)$$

We will now study IC-collisions.

1.5.2 Algebraic Description of 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 (1.50)) implies that for any $n, n' \leq m_i$, the beliefs required to make the n th 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

$$\bar{P}^\sigma(B|s) := \nu \times \nu_0 \left(\{(E, E', y) : B = R^\sigma(e^{\sigma E'} y)\} | e^{\sigma E} y = s \right). \quad (1.44)$$

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

$$P_i^\sigma(B_{-i}|s) = \sum_{\tilde{B} \in \mathcal{A}_1 \times \mathcal{A}_2 : \tilde{B}_{-i} = B_{-i}} \bar{P}^\sigma(\tilde{B}|s). \quad (1.45)$$

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

$$\eta(R^\sigma(s)) = R^\sigma(Xs). \quad (1.46)$$

Condition (1.46) 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 $p \in \Delta(\mathcal{A}_1 \times \mathcal{A}_2)$

$$\eta(\widetilde{\text{BR}}(p|\tilde{s})) = \widetilde{\text{BR}}(p \circ \eta|X\tilde{s}), \quad (1.47)$$

where for every i ,

$$\widetilde{\text{BR}}_i(p|s) = \bigcup_{\substack{\kappa_{-i}: \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \Delta(A_{-i}) \text{ s.t.} \\ \kappa_{-i}(B_{-i}|B_1, B_2) = 1, \forall (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2}} \text{br}_i(p \circ \kappa_{-i}|s), \quad (1.48)$$

and $p \circ \kappa_{-i}(a_{-i}) := \sum_{(B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2} \kappa_{-i}(a_{-i} | B_1, B_2) p(B_1, B_2)$. Condition (1.47) is analogous to (1.46): 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 (1.49).

Lemma 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 $B \in \mathcal{A}_1 \times \mathcal{A}_2$,*

$$\bar{P}^\sigma(B | \hat{s}) = \bar{P}^\sigma(\eta(B) | X \hat{s}). \quad (1.49)$$

The result in Lemma 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 :

$$\begin{aligned} \bar{P}^\sigma(B | s) &= \nu \left(\left\{ (E, E') : B = R^\sigma(e^{\sigma(E'-E)} s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = R^\sigma(X e^{\sigma(E'-E)} s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = R^\sigma(X e^{\sigma(E'-E)} X^\top X s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = R^\sigma(e^{\sigma(E'-E)} X s) \right\} \right) \\ &= \bar{P}^\sigma(\eta(B) | X s). \end{aligned} \quad (1.50)$$

The first line follows from the definition of \bar{P}^σ in (1.45). 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: $X e^{\sigma(E'-E)} X^\top = e^{\sigma X(E'-E)X^\top}$ and symmetry of the noise distribution. When (η, X) preserves best-replies we show inductively on the rounds of elimination of R , that (η, X) is an invariance and so (1.50) holds.

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

$$X s^* = s^*. \quad (1.51)$$

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

$$(\eta(B), \eta(B')) \in Z. \quad (1.52)$$

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

The symmetry properties of beliefs established in Lemma 2 only have bite for collisions if collisions consisted of invariant points, or, if collisions are full. Lemma 3 establishes that fact.

Lemma 3 (Every Collision is full). *For every $Z \subseteq \mathcal{A}$ and $s \in \mathcal{C}(Z)$, the collision Z is full at s .*

Combining Lemmas 2 and 3 establishes that any label permutation η that acts on Z defines an invariance whose invariant points contain the collision of Z .

Proposition 1 (Symmetry of Colliding Beliefs). *For every collision $\mathcal{C}(Z)$, every $s \in \mathcal{C}(Z)$ and every permutation η that acts on Z , there exists a permutation matrix $X \in \mathbb{P}^{2^n}$ so that (η, X) is an invariance with invariant point s .*

Proof. This is an immediate consequence of Lemmas 2 and 3. \square

The symmetry of colliding beliefs established in Proposition 1 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 $\mathcal{C}(Z)$, every $s \in \mathcal{C}(Z)$, every $\sigma > 0$ and every $(B, B') \in Z$, there exist $s_{B, B'}^\sigma \in \partial\mathcal{R}^\sigma(B, B')$ so that

- (i) $\|s - s_{B, B'}^\sigma\|_2 < \sigma$,
- (ii) $\|s_{\hat{B}, \hat{B}'}^\sigma - s_{B, B'}^\sigma\|_2 < \sigma, \forall (\hat{B}, \hat{B}') \in Z$, and
- (iii) by varying the permutation η we also have that

$$\lim_{\sigma \downarrow 0} \sum_{(\hat{B}, \hat{B}') \in Z} \bar{P}^\sigma(B | s_{\hat{B}, \hat{B}'}^\sigma) = 1. \quad (1.53)$$

Property (iii) is the generalization of the symmetry property (1.19) 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

$$d_{s_i}^{a_i, a'_i, a_{-i}} := s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}}. \quad (1.54)$$

We prove Lemma 2 by exploiting the symmetry property of the noise distribution and the resulting relation (1.53). This relation across boundaries pins down the coefficients of a multilinear form, which we specify in the Lemma below.

Lemma 4 (Multilinearity of IC-Collision). *For every collection of action pairs $C_i = \{(a_i^1, a_i'^1), \dots, (a_i^{m_i}, a_i'^{m_i})\}$, and for every player i , there exists a multilinear form δ_C so that*

$$\mathcal{C}(C) = \{s \in S : \delta_C(d_s^1, \dots, d_s^K) = 0\}, \quad (1.55)$$

where for all $s \in S$,

$$\delta_C(d_s^1, \dots, d_s^K) = \sum_{a \in A_1 \times A_2} \chi_{a, C} \prod_{i=1, 2} \prod_{k=1}^{m_i} d_{s_i}^{a_i^k, a_i'^k, a_{-i}}, \quad (1.56)$$

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

$$d_s^k := \left(d_{s_i}^{a_i^k, a_i'^k, a_{-i}} \right)_{a_{-i} \in A_{-i}}. \quad (1.57)$$

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

Proposition 2 (Multilinearity). *For every $Z \subseteq \mathcal{A}$ there exist a corresponding set $C = (C_1, C_2)$, of the form $C_i = \{(a_i^1, a_i^1), \dots, (a_i^{m_i}, a_i^{m_i})\}$ for every i , and a family of binary coefficients $(\chi_{a,C})_a \in \{0, 1\}^{A_1 \times A_2}$, so that*

$$\mathcal{C}(Z) \subseteq \{s \in S : \delta_C(d_s^1, \dots, d_s^K) = 0\}, \quad (1.58)$$

where δ_C is given by (1.56).

Proof. This follows immediately from Lemma 4 and Corollary 1. \square

Proposition 2 thus provides the generalization of the bilinear form (1.20) derived in the binary action case.

1.5.3 Aligned Incentives

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

$$\{s \in S : \exists \tilde{s} \in \mathcal{C}(C) \text{ s.t. } \|s - \tilde{s}\|_2 < \epsilon\}. \quad (1.59)$$

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

$$\gamma^{a_i, a_i', a_{-i}}(d_s) := \frac{\partial \delta_C}{\partial d_{s_i}^{a_i, a_i', a_{-i}}}(d_s). \quad (1.60)$$

Define the sets of actions

$$I^{a_i, a_i'}(d_s) := \left\{ a_{-i} \in A_{-i} : \gamma^{a_i, a_i', a_{-i}}(d_s) > 0 \right\}, \quad (1.61)$$

$I^{a_i, a_i'}(d_s)$ 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 1.3,

we perturbed the zero-set of the bilinear form in equation (1.20) 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

$$G_{C,i} := \bigcup_m \{(a_i^1, \dots, a_i^m) : (a_i^k \bmod m, a_i^{(k+1) \bmod m}) \in C, \forall k < m\}. \quad (1.62)$$

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 $k < m$,

$$a_{-i}^{k+1} \in I_i^{a_i^{k+1}, a_i^k}, a_i^k \in I_{-i}^{a_{-i}^{k+1}, a_{-i}^k} \quad (1.63)$$

and

$$a_{-i}^1 \in I_i^{a_i^1, a_i^m}, a_i^m \in I_{-i}^{a_{-i}^1, a_{-i}^m}. \quad (1.64)$$

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

$$a_{-i}^{k_{-i}} \in I_i^{a_i, a_i^{k_i}}, \quad (1.65)$$

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 $l < h$, $\tilde{a}_{-i}^l \in I_i^{\tilde{a}_i^{l+1}, \tilde{a}_i^l}$, $\tilde{a}_i^{l+1} \in I_{-i}^{\tilde{a}_{-i}^{l+1}, \tilde{a}_{-i}^l}$, 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 5 (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 3 (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 5. □

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 now restrict to the class of quasi concave supermodular games.

1.5.4 Limit Uniqueness for Quasi Concave, Supermodular Games

For any collection of action pairs for each player $C = (\{(a_i^1, a_i'^1), \dots, (a_i^{m_i}, a_i'^{m_i})\})_{i=1,2}$, let $B(C) = (B_1(C), B_2(C)) \in \mathcal{A}_1 \times \mathcal{A}_2$ denote the associated pair of action sets that collect all the actions in C for each player

$$B_i(C) := \{a_i^n : n \leq m_i\} \cup \{a_i'^n : n \leq m_i\}. \quad (1.66)$$

Lemma 6 (No Cycles \implies AI). *Let s be supermodular and quasi concave. If $B(C)$ has no better response cycles under s then incentives are aligned on (C_1, C_2) .*

We can now prove our main result, Theorem 1. We proceed as follows: Let $O \subseteq S^*$ be open and suppose every $s \in O$ no BRC set contains a risk-dominant better response cycle. Then by Lemma 6 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 3 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 $s \in O$ have a risk-dominant better response cycle $B = (B_1, B_2)$ in some 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*, which we show is sufficient for \bar{s} to satisfy limit multiplicity.

We then show that every collision on the path fails aligned incentives. By Proposition 3 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 1.9 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 1.9 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.

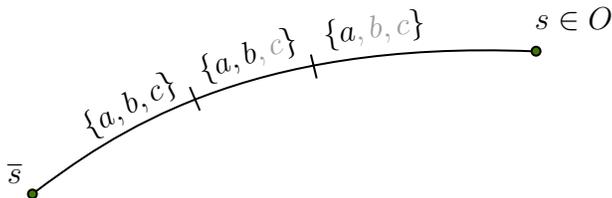


Figure 1.9: Path between game with best-response cycle and game with better response cycle.

1.5.5 Potential Games

In games with a generalized ordinal potential, the incentive of all players to switch actions can be expressed using a common function. Let $B = (B_1, B_2) \in \mathcal{A}_1 \times \mathcal{A}_2$. Say that s admits an *generalized ordinal potential on B* if there exists a function $F_s: A_1 \times A_2 \rightarrow \mathbb{R}$ so that for every player i , all actions $a_i, a'_i \in B_i$, and any $a_{-i} \in A_{-i}$,

$$s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}} > 0 \implies F_s(a_i, a_{-i}) - F_s(a'_i, a_{-i}) > 0. \quad (1.67)$$

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.

Proposition 4 (Monderer and Shapley (1996)). *A game s admits a generalized ordinal potential on B if and only if B is no better-response cycle.*

Our main result can thus also be stated in terms of the existence of a generalized ordinal potential on every risk-dominant set contained in a BRC set.

Corollary 2. *Let $O \subseteq S^*$ be open. Limit uniqueness holds on O if and only if for every $s \in O$, every Risk-Dominant set contained in a BRC set admits a generalized ordinal potential.*

Many economic situations are modeled as potential games (congestion games, aggregative action games for example). If there are good economic reasons to believe that there is a potential, then limit uniqueness holds. However, the features of the model that can break this property may sometimes be hard to rule out based on economic arguments alone.

1.6 Examples

In this section we show that the failure of limit uniqueness can result from economically interpretable conditions. We will study smooth, two-player infinite action games, which we can think of as approximations of large finite games. A sufficient condition for the existence of better response can be stated in terms of the third-order derivative of the players' payoff functions. In particular, the game need not be asymmetric. 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.

1.6.1 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 $(a_1, a_2) \in [0, 1]^2$,

$$g_i(a_1, a_2) = \alpha_i a_i - a_i^{\alpha_i} a_{-i}^{1-\alpha_i}, \quad (1.68)$$

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

$$g_i(a^*, a^*) = \frac{\alpha_i - 1}{2} 2a^*, \quad (1.69)$$

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

$$g_{-i}(a^*, a^*) = \frac{3 - \alpha_i}{2} 2a^*. \quad (1.70)$$

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 1.10 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 $\alpha_i \in (1, 2)$.

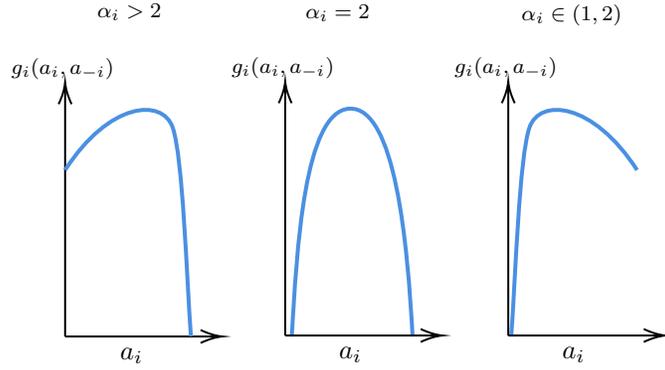


Figure 1.10: Payoffs for different choices of α_i .

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^*) :

$$\prod_{i=1,2} \frac{\partial^3 g_i}{\partial a_i^3}(a^*, a^*) < 0. \quad (1.71)$$

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.

1.6.2 Example: Symmetric Investment Game

Consider the following two-player game with symmetric, smooth payoffs:

$$g_i(a_i, a_{-i}) = a_i \left(1 + a_{-i} - \left(\frac{a_i}{a_{-i}} \right)^{f(a_{-i})} \right), \quad (1.72)$$

where actions are $a_i \in (0, \infty)$ for every player i and $f(a_{-i}) > 0$, for all a_{-i} . In this game, there are two opposing forces: On the one hand, the term $a_i + a_i a_{-i}$ means that higher actions translate to higher payoffs for player i , and even more so if player $-i$ plays higher actions. On the other hand, the term $-\left(\frac{a_i}{a_{-i}}\right)^{f(a_{-i})}$ decreases payoffs more if player i chose to invest more than player $-i$. This aspect can be interpreted as a winners curse: If i invests more than a_{-i} , then perhaps she is being overconfident. We capture this concern for overconfidence in reduced form by this term.

Consider first the case where $f(a_{-i}) = 1, \forall a_{-i}$. In that case, the game has a symmetric pure strategy Nash equilibrium given by action profile $(1, 1)$. Moreover, the game admits a generalized ordinal potential given by

$$F(a_i, a_{-i}) = \sum_{i=1,2} g_i(a_i, a_{-i}). \quad (1.73)$$

Indeed,

$$\begin{aligned} g'_i(a_i, a_{-i}) &= 1 + a_{-i} - \frac{2}{a_{-i}} a_i > 0 \\ \implies F'(a_i, a_{-i}) &= 1 + a_{-i} - \frac{2}{a_{-i}} a_i + a_i + \left(\frac{a_{-i}}{a_i}\right)^2 > 0 \end{aligned} \quad (1.74)$$

We conclude that limit uniqueness holds for any large enough, finite grid-approximation of this game.¹³

Next, suppose that the concern for overconfidence of player i gets worse when player $-i$ invests more. Consider,

$$f(a_{-i}) = a_{-i}. \quad (1.75)$$

This game has a continuum of symmetric, pure strategy Nash equilibria given $\{(a^*, a^*) : a^* > 0\}$. Moreover,

$$\frac{\partial^3 g_i}{\partial a_i^3}(a_i, a_{-i}) = -(a_{-i} + 1)a_{-i}(a_{-i} - 1) \frac{a_i^{a_{-i}-2}}{a_{-i}^{a_{-i}}}. \quad (1.76)$$

Focusing on the previous Nash equilibrium $(1, 1)$, we now have that $\frac{\partial^3 g_i}{\partial a_i^3}(1, 1) = 0$ and

$$\frac{\partial^4 g_i}{\partial a_i^3 \partial a_{-i}}(1, 1) < 0. \quad (1.77)$$

Since the third order derivative switches sign as we increase player $-i$'s actions at $(1, 1)$, we conclude that the payoffs of player i switch skewness at that point. It can be readily shown that every neighborhood around Nash equilibrium $(1, 1)$ contains a better response cycle.

¹³In particular, the potential maximizing equilibrium will be selected.

1.6.3 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 1.6.1, 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, quasi concave, 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 $g_i: A_i \times A_{-i} \rightarrow \mathbb{R}$ 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 1.6.1, 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,

$$\frac{\partial^3 g_i}{\partial a_i^3}(a_1, a_2) \geq 0, \quad \forall i \in \{1, 2\}, \quad \forall (a_1, a_2) \in A_1 \times A_2. \quad (1.78)$$

In that case, the game admits a generalized ordinal potential everywhere and so, as we argued in Example 1.6.1, 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 $f_i: A_i \times A_{-i} \rightarrow \mathbb{R}$ for every player i , so that

$$f_1(a_1, a_2) + f_2(a_1, a_2) = 0, \quad \forall (a_1, a_2) \in A_1 \times A_2. \quad (1.79)$$

Can we find a budget balanced intervention so that the induced game $\hat{g} = (g_1 + f_1, g_2 + f_2)$ satisfies (1.71) 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 (1.71), redistribution must introduce skewness in opposing directions into players' payoff functions.

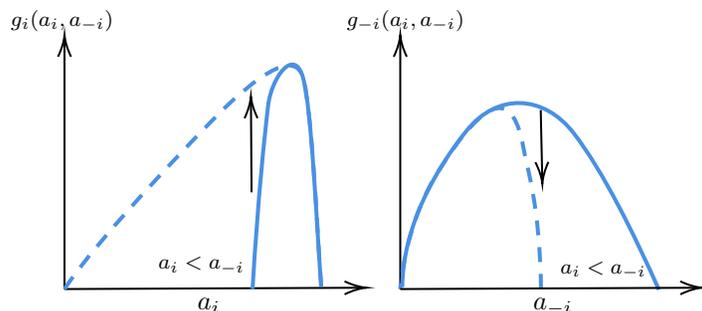


Figure 1.11: Budget Balanced Intervention cannot generate “opposite skewness” property.

Figure 1.11 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 1.11) 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.

1.7 Discussion

1.7.1 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 (1993a) 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 ???. In subsection 1.7.2 below, we show that the affine set-up in Carlsson and Van Damme (1993a) 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.

1.7.2 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

$$x = w + \sigma \varepsilon_i, \tag{1.80}$$

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 $\bar{R}_i^\sigma : U \rightarrow \mathcal{A}_i$ and let the induced limit selection be denoted $\bar{R} : U \rightarrow \mathcal{A}_1 \times \mathcal{A}_2$, so that for every $x \in U$ and player i ,

$$\liminf_{\sigma \downarrow 0} \bar{R}_i^\sigma(x) = \bar{R}_i(x). \tag{1.81}$$

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

$$\varphi(u) := \frac{u}{\|u\|_2}. \tag{1.82}$$

Proposition 5 below establishes that every classical global game can be replicated by an outcome equivalent spherical global game.

Proposition 5 (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 $u \in U$,*

$$\bar{R}(\varphi(u)) = R(\varphi(u)). \tag{1.83}$$

1.7.3 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 2. 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 $F: A_1 \times A_2 \rightarrow \mathbb{R}$ so that for every player i , for every $a_i \prec a_i^*$ there is $\bar{a}_i \succ a_i$; and for every $a_i \succ a_i^*$ there is $\underline{a}_i \prec a_i$, so that for all

$$(a_i, a'_i) \in H_i(a_i^*) := \left\{ (a_i, a'_i) : a_i \in A_i, \begin{array}{l} a'_i \in [a_i, \bar{a}_i], \text{ if } a_i \prec a_i^* \\ a'_i \in [\underline{a}_i, a_i], \text{ if } a_i \succ a_i^* \end{array} \right\}, \quad (1.84)$$

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

$$d_{s_i}^{a_i, a'_i, a_{-i}} \mu_i(a_i) \geq F(a_i, a_j) - F(a'_i, a_j). \quad (1.85)$$

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). 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.

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.

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.

1.7.4 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 3 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.

Chapter 2

Strategic Type Spaces

2.1 Introduction

For games of incomplete information, [Harsanyi \(1967\)](#) introduced type spaces as models to describe players' information on uncertain payoff-relevant parameters (i.e. states of nature), where each type is associated to a belief on states of nature and other players' types. [Mertens and Zamir \(1985\)](#) show that these type spaces can be represented in canonical models (universal type spaces) of players' hierarchies of beliefs, whose descriptions do not depend on the payoff structure of a game. Universal types should then contain payoff relevant information for all games and so become very large objects.

This paper takes back the question of how to describe players' information by taking a game as fixed. We provide a universal representation of players' payoff relevant information by switching the focus from a purely informational description of types to a strategic description. For a fixed game, we introduce strategic quotients as descriptions of players' strategic information, which allow for an economical representation of every Harsanyi type space.

In our approach, an economical representation of a Harsanyi type space is obtained by mapping its types and associated beliefs into equivalence classes. A strategic quotient is a canonical set of such equivalence classes, called strategic types, which satisfies the following two conditions for every player i :

1. For any Harsanyi type space, if two of player i 's types have same beliefs on nature and strategic types of other players, then these types belong to the same strategic type of player i .
2. If other player's behaviors depend only on their strategic types, there exists a best-response of player i that depends only on i 's strategic type.

The first is a sufficient condition for different Harsanyi types to be merged into the same strategic type. Unlike [Harsanyi \(1967\)](#), we do not require the converse of this condition. We thus allow for different beliefs over states and other players' strategic types to correspond to the same strategic type, so that strategic quotients partition Harsanyi type spaces. The second condition depends on a best-reply concept and the meaning of strategic behavior. We focus on the best-reply correspondence that underlies the solution concept of interim

correlated rationalizability (ICR), as defined in [Dekel et al. \(2007a\)](#). This condition implies that strategic quotients cannot be too coarse, as they must be closed under best-replies.

An example of a strategic quotient for any game is the universal type space of [Mertens and Zamir \(1985\)](#), and we show that every strategic quotient is a quotient of the universal type space. Since we are interested in economical information representations, we want to focus on strategic quotient that are the smallest, or coarsest. We prove existence and uniqueness of a minimal strategic quotient which we call the strategic type space (STS). We show this by first proving that all finite order ICR actions arising from any Harsanyi type can be recovered from any Strategic Quotient. We then provide a canonical construction of the set of best-reply hierarchies and show that it forms a STS when each type is associated to the set of beliefs that rationalize it as a best-response. We show that these hierarchies characterize all finite order ICR actions and deduce that our construction characterizes a unique STS up to isomorphisms.

We then analyze properties of the STS. Through careful exploration of the recursive structure underlying our construction of the STS, we show that ICR hierarchies exhibit a finite recursive structure. More precisely, we show the existence of a finite automaton associated to the underlying game, called the strategic automaton, such that ICR hierarchies are 1-1 with the set of paths on the automaton. This characterization allows us to further deduce that the STS is a compact and Hausdorff space when endowed with the product topology on ICR hierarchies. As an application of our approach, in a companion paper [?](#) we use the recursive structure to characterize the set of implementable distributions under interim correlated rationalizability.

2.1.1 Related Literature

The best-reply concept we focus on in this paper was introduced to define Interim correlated rationalizability (ICR). Rationalizability was introduced by [Bernheim \(1984\)](#); [Pearce \(1984\)](#) in games with complete information. [Dekel et al. \(2007a\)](#) generalized this concept by introducing the concept of ICR for games of incomplete information. For every type, ICR iteratively eliminates never best-replies to that type's expectation over any state contingent, correlated beliefs over other types' actions.

[Dekel et al. \(2007a\)](#) show that two Harsanyi types have the same ICR actions in all games if and only if they correspond to the same hierarchy of beliefs and hence the same point in the universal type space of [Mertens and Zamir \(1985\)](#). Therefore, ICR has been studied as a correspondence on the universal type space of [Mertens and Zamir \(1985\)](#). [?](#) characterizes ICR in terms of a common belief operator on the universal type space for global games. [?](#) first identified critical types, i.e. points of discontinuity of ICR in the universal type space of [Mertens and Zamir \(1985\)](#). They provide a topological characterization of critical types. [Dekel et al. \(2005\)](#) and [?](#) characterize the coarsest topology on the universal type space, called the strategic topology, under which ICR is continuous. [?](#) introduce the notion of *frames* as partitions of the universal type space similar to the first property of STS. They use frames as a tool to define a strategic topology of uniform convergence over games for hierarchies of beliefs. [?](#) study robustness of ICR to both higher order beliefs and payoff perturbations. The authors define *curb* collections which is closely related to the second requirement of STS, i.e. strategic closure (see Section [2.3.2](#)), defined in terms of the universal type space.

? provide an algorithm to compute hierarchies of ICR which parallels our construction of best reply hierarchies. Based on their construction, they study refinements on ICR. Finally, ? provide a characterization of critical types in terms of common belief properties in the universal type space.

Most importantly, this paper differs from the literature described above in the following way: We fix a game and introduce a canonical language to describe strategically relevant information for this game. Unlike frames and curb collections, STS are defined as universal objects which can be characterized and constructed without reference a particular Harsanyi type space.

2.2 Preliminaries and Notations

We denote the cardinality of a set Y by $|Y|$. For a family of sets $(X_i)_i$ we let $X := \prod_i X_i$ and $X_{-i} := \prod_{j \neq i} X_j$. For a family of mappings $f_i: X_i \rightarrow Y_i$, f is the map from X to Y given by $f(x) = (f_i(x_i))_i$ and f_{-i} is from X_{-i} to Y_{-i} is given by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$. Similarly, if $f: X \rightarrow Y$ and $g: Z \rightarrow W$ are mappings we denote by $(f \times g): X \times Z \rightarrow Y \times W$ the map given by $(f \times g)(x, z) = (f(x), g(z))$. On any set X we denote by id_X the identity mapping on X and omit the subscript X when there is no ambiguity. The set of Borel probability measures over a topological space X is written as Δ_X . We denote by $\text{supp } p$ the support of a probability measure p . The marginal probability on $Y \times X_m$ of a probability measure p on a product space $X = Y \times \prod_i X_i$ is denoted $\text{marg}_{Y,m}(p)$.

In commutative diagrams we describe a mapping between probability measures from Δ_X to Δ_Y which are induced by a measurable mapping from X to Y by an arrow on the subscripts as follows:

$$\begin{array}{c} \Delta_X \\ \downarrow \\ \Delta_Y \end{array}$$

Double headed arrows such as $X \rightleftarrows Y$ denote surjective mappings. The subscript i denotes a typical player from the finite set N of players. A finite set K of states of nature and, finite action sets $(A_i)_{i \in N}$ and a payoff function $u: A \times K \rightarrow \mathbb{R}^N$, are given.

2.3 Strategic Type Spaces

In this section we introduce Strategic type spaces (STS). Section 2.3.1 introduced the best-reply correspondence underlying the concept of Interim correlated rationalizability. In section 2.3.2 we introduce STS and minimal STS axiomatically. In section 2.3.3 we establish existence and uniqueness of a minimal STS, characterized as the space of best-reply hierarchies. Section 2.4 shows that the minimal STS can be represented by a finite automaton.

2.3.1 Interim Correlated Best-Replies

Dekel et al. (2007a) show that ICR can be defined as a fixed point of a best reply correspondence, which we now state.

We introduce the set $\mathcal{B}_i := 2^{A_i}$ of action subsets and define conjectures as maps $\sigma : K \times \mathcal{B}_{-i} \rightarrow \Delta(A_{-i})$. The probability $\langle \sigma, p \rangle$ over $K \times A_{-i}$ induced by a belief $p \in \Delta_{K \times \mathcal{B}_{-i}}$ and a conjecture σ is given by the relation:

$$\langle \sigma, p \rangle(k, a_{-i}) := \sum_{b_{-i} \in \mathcal{B}_{-i}} \sigma(k, b_{-i})(a_{-i}) p(k, b_{-i}).$$

Player i 's best-reply correspondence $\text{BR}_i : \Delta_{K \times \mathcal{B}_{-i}} \rightarrow \mathcal{B}_i$ is given by:

$$\text{BR}_i(p) := \bigcup_{\sigma: \text{supp } \sigma(k, b_{-i}) \subseteq b_{-i}, \forall k, b_{-i}} \left\{ \arg \max_{a_i \in A_i} \sum_{k, a_{-i}} u_i(a_i, a_{-i}, k) \langle \sigma, p \rangle(k, a_{-i}) \right\}.$$

A *Harsanyi type space* \mathcal{H} consists of a family of topological spaces $(\Theta_i)_{i \in N}$ and of continuous mappings $\pi_i : \Theta_i \rightarrow \Delta_{K \times \Theta_{-i}}$, where $\pi_i(\theta_i)$ represents type θ_i 's belief over types of other players and states of nature. As in Dekel et al. (2007a), we rely on the best-reply correspondence BR_i to define ICR on any Harsanyi type space $(\Theta_i, \pi_i)_i$ as follows: ICR of a type θ_i is given by $R_i(\theta_i) = \bigcap_m R_i^m(\theta_i)$, where $R_i^0(\theta_i) = A_i$ and $R_i^m(\theta_i)$ is i 's best response to the $\pi_i(\theta_i)$ -mixtures (i.e. an expectation $\int_{\Theta_{-i}} \sigma(k, \theta_{-i}) \pi_i(\theta_i)(k, d\theta_{-i})$) of all measurable, state and type profile contingent conjectures $\sigma(k, \theta_{-i}) \in \Delta(A_{-i})$ whose support is contained in $R_{-i}^{m-1}(\theta_{-i})$ for all θ_{-i} . We call the sequence $(R^m(\theta))_{m \geq 0}$ the *R-hierarchy* of θ .

2.3.2 Strategic Type Spaces

We define a Strategic Quotient (for ICR) as a pair $\mathcal{S} = (\mathcal{S}_i, \psi_i)_i$ consisting of an N -tuple of topological spaces \mathcal{S}_i and continuous maps $\psi_i : \Delta_{K \times \mathcal{S}_{-i}} \rightarrow \mathcal{S}_i$ which satisfy both a type space quotient requirement and a strategic requirement.

Definition 1 (Type Space Quotient). *A space $\mathcal{S} = (\mathcal{S}_i, \psi_i)_i$ is a Type Space Quotient if, for every Harsanyi type space $\mathcal{H} = (\Theta_i, \pi_i)_i$ there exist a family of maps $(\eta_i)_i$ for which the following diagram commutes:*

$$\begin{array}{ccc} \Theta_i & \xrightarrow{\pi_i} & \Delta_{K \times \Theta_{-i}} \\ \downarrow \eta_i & & \begin{array}{cc} id \downarrow & \downarrow \eta_{-i} \end{array} \\ \mathcal{S}_i & \xleftarrow{\psi_i} & \Delta_{K \times \mathcal{S}_{-i}} \end{array}$$

Definition 1 imposes a sufficient condition for two types of player i to have the same representation in \mathcal{S}_i . The two downward pointing arrows on the right of the diagram coarsen the sigma algebra of every type's beliefs. The commutativity of the diagram then requires the following: If the beliefs of two types θ_i, θ'_i coincide on $K \times \mathcal{S}_{-i}$, then η_i maps θ_i and θ'_i to the same point in \mathcal{S}_i . Note that the reverse implication is not required by the diagram. That is, two types with distinct beliefs on $K \times \mathcal{S}_{-i}$ could also be mapped to the same point in \mathcal{S}_i .

Thus, in our model, types partition beliefs¹. This contrasts with Harsanyi types spaces, where a type is associated uniquely to a belief, and with universal type spaces, where types and beliefs are homeomorphic. The universal type space together with the canonical belief maps that associate to each canonical type the corresponding belief on that state of nature as well as on other player's canonical types is thus a special case of a type space quotient.

Note also that the universal type space is the smallest object onto which information can be projected through faithful transformations in the sense of Gossner (2000a). By partitioning belief spaces we allow for type space quotients to capture coarser information than universal type spaces.

The point, $\mathcal{S}_i = \{*\}$ and constant map $\psi_i: \Delta_{K \times \mathcal{S}_i} \rightarrow \{*\}$ is another example of a type space quotient, for any game. This second example shows that type space quotients may fail to capture strategically relevant information. In our way to introduce a minimality requirement that spaces capture such information, we now define strategically closed families of behaviors.

Definition 2 (Strategic Closure). *A strategically closed family of behaviors for (\mathcal{S}, ψ) is a family \mathcal{A}_i for each player i of continuous mappings $\alpha_i: \mathcal{S}_i \rightarrow \mathcal{B}_i$ such that,*

1. \mathcal{A}_i contains the constant map equal to A_i
2. for every $\alpha_{-i} \in \mathcal{A}_{-i}$, there exists $\alpha_i \in \mathcal{A}_i$ such that the following diagram commutes:

$$\begin{array}{ccc} \Delta_{K \times \mathcal{S}_i} & \xrightarrow{\psi_i} & \mathcal{S}_i \\ id \downarrow & \downarrow \alpha_{-i} & \downarrow \alpha_i \\ \Delta_{K \times \mathcal{B}_i} & \xrightarrow{\text{BR}_i} & \mathcal{B}_i \end{array}$$

For a given pair (\mathcal{S}, ψ) , a set \mathcal{A}_i consists of correspondences α_i which map points in \mathcal{S}_i to action sets. As a minimality requirement on \mathcal{A}_i point 1 of the definition imposes that each \mathcal{A}_i contains the correspondence $s_i \mapsto A_i$ that precludes no action, for any $s_i \in \mathcal{S}_i$.

In point 2 of the definition, commutativity of the diagram imposes two requirements. First, for a family \mathcal{A} to be strategically closed, the diagram imposes a measurability requirement on \mathcal{S} : It requires beliefs that induce different best replies to a behavior in \mathcal{A}_{-i} to be associated to distinct points in \mathcal{S}_i . That is, given any profile $\alpha_{-i} \in \mathcal{A}_{-i}$, player i 's best-response correspondence to this profile, seen from $\Delta_{K \times \mathcal{S}_i}$ to A_i , is in fact \mathcal{S}_i -measurable. Second, any strategically closed family \mathcal{A} must be closed under best replies: A player's best reply to a profile in \mathcal{A}_i , viewed as a correspondence from \mathcal{S}_i to A_i is in \mathcal{A}_i .

Definition 3 (Strategic Quotient). *A Strategic Quotient is a type space quotient (\mathcal{S}, ψ) that admits a strategically closed family of behaviors.*

As an example, the universal type space together with the canonical maps is a strategic quotient, for any game: If other players' strategies are measurable wrt. their canonical types, so is a best-response. The type space consisting of a single point for each player together

¹? introduce the notion of "frames" which are partitions of type spaces that are compatible with the belief structure of the types. Frames are thus a special case of what we call type space quotients.

with the constant map is a strategic quotient when there is no uncertainty on nature, as in those games there is always a constant best-response to constant strategies of the other players. It is, however, not a strategic quotient for general incomplete information games. Take for instance the classical *electronic mail game* (Rubinstein, 1989a) or the coordination game of Section 2.4.2. In that game, the best-response of a player to the other player using either of possible strategies depends on their belief on nature, hence cannot be captured by a constant map.

The next definition formalizes the idea that one quotient is smaller than another one.

Definition 4. A space $\mathcal{S} = (\mathcal{S}_i, \psi_i)_i$ is smaller than another space $\tilde{\mathcal{S}} = (\tilde{\mathcal{S}}_i, \tilde{\psi}_i)_i$ if there exist a continuous surjection from $\tilde{\mathcal{S}}$ to \mathcal{S} so that the following diagram commutes:

$$\begin{array}{ccc} \tilde{\mathcal{S}}_i & \xleftarrow{\tilde{\psi}_i} & \Delta_{K \times \tilde{\mathcal{S}}_{-i}} \\ \downarrow & & \downarrow \text{id} \\ \mathcal{S}_i & \xleftarrow{\psi_i} & \Delta_{K \times \mathcal{S}_{-i}} \end{array}$$

In this definition a type space quotient is smaller than another if the latter admits a representation of the former. That is, all types in the former quotient can be obtained by merging types of the latter. The diagram above requires the following sufficient condition for merging types: If the beliefs of types in the latter quotient coincide on the smaller quotient then these types are merged to the same point in the smaller quotient. The definition below then identifies the minimal strategic quotient according to Definition 4².

Definition 5 (STS). A strategic quotient is called the *strategic type space (STS)* if it is smaller than every strategic quotient.

By Definition 2, all strategic quotients must distinguish types which have different best replies to some strategic behavior. Hence the STS should merge players' types whenever these types have identical best replies to all strategic behaviors from a strategically closed family. It also follows from the discussion above that a STS must be at most as large as the universal type space.

2.3.3 The Strategic Type Space

In this section we establish existence and essential uniqueness of the STS. We prove this result by characterizing the STS in terms of ICR hierarchies: First, we show the ICR hierarchies can be recovered from any Strategic Quotient (Lemma 7). We then provide a construction of \mathcal{S}_i , the set of best reply hierarchies for a game. This construction is canonical as it makes no reference to any Harsanyi type space. We show that these hierarchies coincide with all ICR hierarchies that can arise in all types in all Harsanyi type spaces (Lemma 8). We then construct a map ψ_i , which associates beliefs to best-reply hierarchies and prove that the pair

²Formally, Strategic Quotients form a category whose objects are given by the pairs \mathcal{S} satisfying strategic closure and whose morphisms are given by diagrams as in Definition 4. A minimal strategic quotient is thus a terminal object in the category.

(\mathcal{S}, ψ) is a Strategic Quotient (Lemma 9). We deduce that (\mathcal{S}, ψ) is a STS and show that it is essentially unique (Theorem 2).

Our first theorem states that every Strategic Quotient allows to recover the ICR hierarchies from any Harsanyi type.

Lemma 7 (Factorization of ICR). *For every strategic quotient $(\mathcal{S}_i, \psi_i)_i$ and every $m \in \mathbb{N}$, there exists continuous $\alpha_i^m : \mathcal{S}_i \rightarrow \mathcal{B}_i$ so that for every Harsanyi type space $(\Theta_i, \pi_i)_i$ and associated maps $(\eta_i)_i$ satisfying the diagram of Definition 1,*

$$R_i^m(\theta_i) = \alpha_i^m \circ \eta_i(\theta_i), \quad \forall \theta_i \in \Theta_i, \quad \forall i \in N$$

The proof of this result, as well as all others, is in the appendix.

We now construct the set \mathcal{S} of all hierarchies of best replies. The first level of the hierarchy is given by a player's best replies to beliefs on nature and any opponents' play. Every subsequent level of a best reply hierarchy is then obtained by computing best replies to beliefs on nature and lower levels of best reply hierarchies.

We construct inductively the sets of m -order best reply hierarchies \mathcal{S}_i^m as m -fold sequences of action set profiles. Let $\mathcal{S}_i^0 := \{A_i\}$ for every i . Given \mathcal{S}_i^{m-1} for every i , we define the subset $\mathcal{S}_i^m \subseteq \mathcal{S}_i^{m-1} \times \mathcal{B}_i$ of sequences of the form $s_i^m = (A_i, b_i^1, \dots, b_i^m)$ for which there exists a probability distribution $p_i \in \Delta_{K \times \mathcal{S}_i^{m-1}}$ satisfying

$$\text{BR}_i(\text{marg}_{K,l}(p_i)) = b_i^{l+1}, \quad \forall l < m, \quad (2.1)$$

where $\text{marg}_{K,l}(p_i)$ is the marginal probability of p_i on $K \times \prod_{j \neq i} \text{proj}_l(\mathcal{S}_j^m)$, where proj_l denotes the projection on the l -th coordinate. We define the set of player i 's best reply hierarchies as

$$\mathcal{S}_i := \{s_i \in \mathcal{B}_i^{\mathbb{N}} : s_i^m \in \mathcal{S}_i^m, \quad \forall m \in \mathbb{N}\}.$$

Lemma 8 states that the best reply hierarchies \mathcal{S} characterize all ICR hierarchies that can arise in any Harsanyi type space.

Lemma 8 (Best-Reply Hierarchies are ICR Hierarchies).

- (i) *Let $s^m \in \mathcal{B}^m$, then $s^m \in \mathcal{S}^m$ if and only if there exists a Harsanyi type space (Θ, π) and a type profile $\theta \in \Theta$ so that $s^m = (R^l(\theta))_{l \leq m}$.*
- (ii) *Let $s \in \mathcal{B}^{\mathbb{N}}$, then $s \in \mathcal{S}$ if and only if there exists a Harsanyi type space (Θ, π) and a type profile $\theta \in \Theta$ so that $s = (R^l(\theta))_{l \geq 0}$.*

For every $m \in \mathbb{N}$, we define a beliefs map $\psi_i^m : \Delta_{K \times \mathcal{S}_i^{m-1}} \rightarrow \mathcal{S}_i^m$ by

$$\psi_i^m(p_i) := (A_i, \text{BR}_i(\text{marg}_{K,1}(p_i)), \dots, \text{BR}_i(\text{marg}_{K,m-1}(p_i))).$$

Any belief p_i on $K \times \mathcal{S}_i^{m-1}$ induces, through the projection on the first m coordinates of \mathcal{S}_i , a belief p_i^m on $K \times \mathcal{S}_i^{m-1}$, thus an element $\psi_i^m(p_i^m) \in \mathcal{S}_i^m$. By definition of ψ_i^m , for every $l \leq m$, the first l elements of $\psi_i^m(p_i^m)$ coincide with $\psi_i^l(p_i^l)$. Thus, the sequence $(\psi_i^m(p_i^m))_i$ defines a unique element of \mathcal{S}_i , which we denote $\psi_i(p_i)$.

Note that once the set \mathcal{S}_{-i} of all other players' best-reply hierarchies is known, player i 's best-reply hierarchies are fully characterized by marginal beliefs and do not depend on correlations across different levels of \mathcal{S}_{-i} .

We now specify the topology on the set \mathcal{S} . Recall that strategic closure requires the Strategic Quotient to admit a closed family of continuous strategic behaviors. By construction, the coordinates of a best-reply hierarchy correspond to a closed family of strategic behaviors. We thus endow \mathcal{S} with its product topology, i.e. the coarsest topology so that all coordinate projections are continuous. Lemma 9 below states that (\mathcal{S}, ψ) is a Strategic Quotient and that \mathcal{S} is a topological quotient of the universal type space of [Mertens and Zamir \(1985\)](#).

Lemma 9 (ICR Hierarchies form the STS). *(\mathcal{S}, ψ) is a strategic quotient. Moreover, the maps η from the universal type space to \mathcal{S} are quotient maps, i.e. continuous open surjections.*

By Lemma 7 any finite order ICR hierarchy can be recovered continuously from any quotient space. By Lemma 8 the set \mathcal{S} coincides with all ICR hierarchies. Then by Lemma 9, (\mathcal{S}, ψ) is a Strategic Quotient which can be recovered from all Strategic Quotients. The product topology then ensures that (\mathcal{S}, ψ) is a quotient which is minimal. As the property of minimality is universal, every STS is homeomorphic to \mathcal{S} . Theorem 2 thus states existence and essential uniqueness of the STS:

Theorem 2 (Existence and Essential Uniqueness of STS).

(i) (\mathcal{S}, ψ) is a STS.

(ii) If (\mathcal{S}', ψ') and (\mathcal{S}'', ψ'') are STS then \mathcal{S}'' and \mathcal{S}' are homeomorphic.

2.4 Finite Representation of STS

For every truncated sequence $s^m \in \mathcal{S}^m$, define its orbit as:

$$O^m(s^m) := \{(s^m, s) \in \mathcal{B} \times \mathcal{B}^{\mathbb{N}} : (s^m, s) \in \mathcal{S}\}. \quad (2.2)$$

Denote the collection of all orbits

$$\Omega := \{O^m(s^m) : s^m \in \mathcal{S}^m, m \in \mathbb{N}\}. \quad (2.3)$$

Theorem 3. Ω is a finite set.

We break down the proof into six Claims, all proven in Appendix B.1.2. Here we provide an overview of the arguments. In order to prove Theorem 3, we will exploit a monotonicity property of BR_i according to which beliefs with smaller supports (according to stochastic dominance wrt set inclusion) admit smaller best-response sets: $\text{BR}_i(p) \subseteq \text{BR}_i(p')$ whenever p can be obtained from p' by shifting probability mass from action sets to subsets (see Lemma 11 in Appendix B.1.1). In preparation for using this property, we introduce the *best-response map on sets of sequences* for player i , $B_i : 2^{\mathcal{B}_i^{\mathbb{N}}} \rightarrow 2^{\mathcal{B}_i^{\mathbb{N}}}$ as follows

$$B_i(X) := \{\psi_i(p_i) : p_i \in \Delta(K \times X_{-i})\}. \quad (2.4)$$

Let $B(X) := \prod_i B_i(X)$. Note that by construction, $B(\mathcal{S}) = \mathcal{S}$ and so for every subset $S \subseteq \mathcal{S}$, $B(S) \subseteq \mathcal{S}$. B inherits the following monotonicity property from BR.

Monotonicity. For any $X, X' \subseteq \mathcal{B}^{\mathbb{N}}$ write $X \ll X'$ when the following two properties hold: 1) For every $x \in X$ there exists $x' \in X'$ so that for all $m \in \mathbb{N}$, $x_m \subseteq x'_m$. 2) For every $x' \in X'$ there exists $x \in X$ so that for all $m \in \mathbb{N}$, $x_m \subseteq x'_m$. Lemma 11 implies:

Claim 1. $X \ll X' \implies B(X) \ll B(X')$.

A sequence $s \in \mathcal{S}$ is *maximal after round* $m \in \mathbb{N}$ if there does not exist $\hat{s} \in \mathcal{S} \setminus \{s\}$ so that the following two properties hold: 1) $s^m = \hat{s}^m$ and 2) for all $n > m$, $s_n \subseteq \hat{s}_n$. A sequence $s \in \mathcal{S}$ is *maximal at round* $m \in \mathbb{N}$ if there does not exist $\hat{s} \in \mathcal{S} \setminus \{s\}$ so that the following two properties hold: 1) $s^{m-1} = \hat{s}^{m-1}$ and 2) $s_m \subseteq \hat{s}_m$. We use Claim 1 to establish that all maximal sequences are maximal at every round:

Claim 2. *The sequence $s \in \mathcal{S}$ is maximal after round m if and only if it is maximal at round n for all $n > m$.*

Let $\bar{\mathcal{S}}^m(s^m)$ denote the collection of sequences $\tilde{s} \in \mathcal{S}$ which are maximal after round m and satisfy $\tilde{s}^m = s^m$. Claim 2 then implies the following finiteness property of maximal sequences:

Claim 3. *There exists $L \in \mathbb{N}$ so that for every $m \in \mathbb{N}$ and every $s \in \mathcal{S}$,*

$$|\bar{\mathcal{S}}^m(s^m)| < L. \quad (2.5)$$

Say that $X \subseteq \mathcal{B}^{\mathbb{N}}$ has *converged at round* $m \in \mathbb{N}$ if for all $x \in X$ and all $n, l \geq m$, $x_n = x_l$. Claim 3 and the monotonicity of B imply the following convergence property of maximal sequences after any finite history:

Claim 4. *There exists M so that for every $m \in \mathbb{N}$ and $s \in \mathcal{S}$, the set $\bar{\mathcal{S}}^m(s^m)$ converged at round $m + M$.*

Finite Generation of the STS. We present a construction of the STS through iterated applications of the best-response operator. At each step of the construction, we allow for at most one more non-maximal transition than in the previous step. In turn, this construction will allow us to conclude on the finiteness of Ω .

We denote by $\bar{T}^{0,0}$ the set of maximal sequences $\bar{\mathcal{S}}^0(s^0)$. For every $m > 0$, let $\bar{T}^{m,0}$ denote the set of m -maximal sequences, i.e., the set of sequences that admit at most m non-maximal transitions. We also let $\bar{T}^{m,n}$ be the subset of $\bar{T}^{m,0}$ consisting of sequences that admit at most $m - 1$ non-maximal transitions before round n and that are maximal after round n . Sequences in $\bar{T}^{m,n}$ may or not have a non-maximal transition at round n . Finally, let $\bar{T}^m = \cup_n \bar{T}^{m,n}$.

Our next lemma shows that the sets $\bar{T}^{m,n}$ are constructed iteratively through the best-response operator.

Claim 5. *For every $m \geq 0, n > 0$, $\bar{T}^{m,n}$ is the set of sequences s that are maximal after round n , admit at most $m - 1$ non-maximal transitions before round n , and such that:*

$$s \in B(\bar{T}^{m,n-1}).$$

Claim 6 below establishes that \mathcal{S} can be obtained as a finite union of the sets $(\bar{T}^{m,0})_m$. The result follows from the fact that every sequence $s \in \mathcal{S}$ only makes a finite number of non-maximal transitions.

Claim 6. *There exists $N \in \mathbb{N}$ so that*

$$\bigcup_{m=1}^N \bar{T}^{m,0} = \mathcal{S}. \quad (2.6)$$

Cyclicity and Finiteness. We now show that for every m the infinite sequence of sets $(\bar{T}^{m,n})_n$ is eventually cyclic, which will establish the finiteness of Ω . First, we conclude from Claim 4 and the construction of each set $\bar{T}^{m,n}$ that the sequences in $\bar{T}^{m,n}$ converge within a bounded number of rounds after round n :

Corollary 3. *Let N satisfy (B.15). For every $m \leq N$ there exists $M_m \in \mathbb{N}$ so that for every $n \in \mathbb{N}$, $\bar{T}^{m,n}$ has converged at round $n + M_m$.*

By monotonicity of B , sequences in $\bar{T}^{m,n}$ are best-replies to sequences in $\bar{T}^{m,n-1}$ and by Corollary 3 all sequences in $\bigcup_{l \geq M_m} \bar{T}^{m,n-l}$ have converged before round n ; We conclude that for any two rounds $n > \tilde{n}$ so that

$$\{O(s^l) : s \in \bar{T}^{m,l}\} = \{O(s^{l+n-\tilde{n}}) : s \in \bar{T}^{m,l+n-\tilde{n}}\}, \quad (2.7)$$

for every $l \in \{\tilde{n} - M_m, \dots, \tilde{n}\}$, we also have that

$$\{O(s^{\tilde{n}+1}) : s \in \bar{T}^{m,\tilde{n}+1}\} = \{O(s^{n+1}) : s \in \bar{T}^{m,n+1}\}. \quad (2.8)$$

From Claim 3 we conclude that for every m , the number of sequences that are contained in $\bar{T}^{m,n} \setminus \bar{T}^{m,n-1}$ is bounded uniformly over all $n \in \mathbb{N}$.

Corollary 4. *Let N satisfy (B.15). For every $m \leq N$ there exists $L_m \in \mathbb{N}$ so that for every $n \in \mathbb{N}$,*

$$|\bar{T}^{m,n} \setminus \bar{T}^{m,n-1}| \leq L_m. \quad (2.9)$$

We conclude from both corollaries that for every $m \in \mathbb{N}$, there exists $n_m \in \mathbb{N}$ and $\tilde{n}_m < n_m$ so that (2.7) holds and so we deduce the result below:

Claim 7. *Let N satisfy (B.15). For every $m < N$ there exists $z_m \in \mathbb{N}$ and $n_m \in \mathbb{N}$ so that for all $s \in \bar{T}^{m+1,0}$*

$$O(s^n) = O(s^{n+z_m}), \quad \forall n \geq n_m. \quad (2.10)$$

Claims 7 and 6 then imply that the set of tails is finite, which is what we needed to show.

2.4.1 Automaton Representation

Define an automaton as a tuple $\mathcal{A} = (\hat{\Omega}, \hat{\beta}, \hat{\preceq}, \hat{\omega}^0)$, where $\hat{\Omega}$ is a finite set of automaton states, $\hat{\beta}: \hat{\Omega} \rightarrow \mathcal{B}$ assigns an action set profile to every state, $\hat{\preceq}$ is a binary successor relation on $\hat{\Omega}$ and $\hat{\omega}^0 \in \hat{\Omega}$ is an initial state. A path on the automaton is a sequence of states $(\omega^0, \omega^1, \omega^2, \dots)$ so that $\omega^0 = \hat{\omega}^0$ and for every $m \in \mathbb{N}$,

$$\omega^m \hat{\preceq} \omega^{m+1}. \quad (2.11)$$

Let $P_{\mathcal{A}}$ denote the set of paths. From Theorem 3 we obtain a finite automaton representation of all R-hierarchies: Let $\beta: \Omega \rightarrow \mathcal{B}$ recover the first coordinate from sequences in each orbit. Define the shift operator on sequences $\gamma: \mathcal{B}^{\mathbb{N}} \rightarrow \mathcal{B}^{\mathbb{N}}$, which for each sequence (s^1, s^2, \dots) removes the first coordinate

$$\gamma: (s^1, s^2, \dots) \mapsto (s^2, \dots). \quad (2.12)$$

Define the successor relation \preceq on Ω , where for every $\omega, \hat{\omega} \in \Omega$, $\omega \preceq \hat{\omega}$ if and only if

$$\hat{\omega} \subseteq \{\gamma(s) : s \in \omega\}. \quad (2.13)$$

Letting $\omega^0 = O^0(s^0)$, we obtain an automaton.

Theorem 4. *The tuple $\mathcal{A} = (\Omega, \beta, \preceq, \omega^0)$ is an automaton so that*

$$\{(\beta(\omega^0), \beta(\omega^1), \dots) : (\omega^0, \omega^1, \dots) \in P_{\mathcal{A}}\} = \mathcal{S}. \quad (2.14)$$

We deduce that the STS is countable, compact and Hausdorff.

2.4.2 Example

In this section we illustrate our results and construct the STS in an two players coordination Game.

Consider the following two-player game: $N = \{1, 2\}$, $K = \{-1, 1\}$ and $A_i = A = \{a, b\}$ where payoffs are given by:

	a	b
a	k, k	$-1, 0$
b	$0, -1$	$0, 0$

If player i (row player) believes that $k = 1$ with probability less than $\frac{1}{2}$ then b is a dominant action. Otherwise, neither action dominates the other. First order hierarchies of best replies in this game, \mathcal{S}_i^1 , are thus given by $\{(A, b), (A, A)\}$. The first pair corresponds to beliefs which put less than half of the probability on $k = 1$. Indeed, recall that $\mathcal{S}_{-i}^0 = \{a, b\}$ from Section 2.3.3 and consider any belief $p \in \Delta_{K \times \mathcal{S}_{-i}^0}$. Player i thus forms best replies to p -mixtures of state-contingent conjectures $\sigma: K \rightarrow \Delta_A$. In the simplex $\Delta_{K \times A}$, these p -mixtures over conjectures form geometric rectangles - the set of probabilities on $K \times A$ with constant marginal belief on K given by (p_1, p_{-1}) . The right panel of Figure 2.1 illustrates these rectangles for $p_1 < \frac{1}{2}$ and $p_1 \geq \frac{1}{2}$. The left of Figure 2.1 plots the simplex $\Delta_{K \times A}$, where the

shaded triangle with dashed contour marks the boundary of the partition induced by the best response correspondence of player i . When $p_1 < \frac{1}{2}$, the mixture of the conjectures is entirely included in the region where b_i is the unique best-response. When $p_1 \geq \frac{1}{2}$, the conjectures cross regions where a and b , or both are best-responses. Hence $\mathcal{S}_i^1 = \{(A, A), (A, b)\}$.

Figure 2.1: For all beliefs in the region between the shaded triangle (excluded) and the sub simplex spanned by $(a, -1)$, $(b, -1)$ and $(b, 1)$, player i 's best response is always b . For all beliefs in the region between the shaded triangle (excluded) and $(a, 1)$, player i 's best response is always a and on the shaded triangle all beliefs induce both actions a and b as best response.

We repeat the same procedure on \mathcal{S}_{-i}^1 . For a belief p on $K \times \mathcal{S}_{-i}^1$ of player i , let p_b denote the probability put on hierarchies ending at b , let p_1 denote the probability put on $k = 1$ and $p_{k,b}$ be the joint probability on state $k \in \{-1, 1\}$ and hierarchies ending at b . As can be seen in Figure 2.1, $2p_1 - p_b + (p_{-1,b} - p_{1,b}) < 1$ describes the portion of a rectangle associated to p_1 where b is a unique best reply for player i . Hence the set of beliefs on $K \times \mathcal{S}_{-i}^1$ so that BR_i maps to A is given by $p_1 \geq \frac{1}{2}$ and $2p_1 - p_b + (p_{-1,b} - p_{1,b}) \geq 1$. We deduce that $\mathcal{S}_i^2 = \{(A, A, A), (A, A, b), (A, b, b)\}$ corresponds to the following partition of $\Delta_{K \times \mathcal{S}_{-i}^1}$:

- (1) $2p_1 - p_b + (p_{-1,b} - p_{1,b}) \geq 1$ and $p_1 \geq 1/2$, for (A, A, A)
- (2) $2p_1 - p_b + (p_{-1,b} - p_{1,b}) < 1$ and $p_1 \geq 1/2$, for (A, A, b)
- (3) $p_1 < 1/2$, for (A, b, b)

Note that these conditions only depend i 's beliefs on K and on the last coordinate in \mathcal{S}_{-i}^1 . As the last coordinates of \mathcal{S}^1 are the same as the last coordinates of \mathcal{S}^1 we deduce that the game is indeed simple. we argue that all transitions in \mathcal{S} are described by the three rules above. The STS automaton in Figure 2.2 below illustrates the transition for coordinates in \mathcal{S}_i and \mathcal{S}_{-i} in this game:

Figure 2.2: STS Automaton.

The automaton represented in Figure 2.2 above describes the following transitions: If player i 's m -th coordinate is A , then it must be that $p_1 \geq \frac{1}{2}$ and i 's $(m + 1)$ -th coordinate in any strategic type must be one of $\beta_i(A) := \{A, b\}$. Moreover, the $(m + 1)$ -th coordinate is A if i also believes in strategic types of $-i$ whose m -th coordinate is b with low enough probability. That is, i believes that $-i$ plays b with low enough probability in round m (i.e. condition (1)). Otherwise, the $(m + 1)$ -th coordinate must be b (condition (2)). However, if i 's m -th coordinate was b , then i 's $(m + 1)$ -th coordinate must be $\beta_i(b) = b$. In this case, i 's beliefs satisfy either condition (3) or condition (2).

For this example, note first that the only possible change in the last coordinate when going from \mathcal{S}_i^1 to \mathcal{S}_i^2 is to move from A to b . A probability on $K \times \mathcal{S}_{-i}^2$ must therefore put at least as much probability on sequences ending with b than its marginal on $K \times \mathcal{S}_{-i}^1$. But under this constraint, third order types can also only move from A to b or stay unchanged. Hence the automaton above generates all the sequences in \mathcal{S} .

2.5 Discussion

Our construction is topological and the minimality property implies that the STS is endowed with the product topology on sequences. However, the automaton we have constructed in Section 2.4.2 illustrates how strategic continuity fails in the product topology. In Figure 2.2, the type space where the state in K is common knowledge can be represented by the constant sequence cycling through the full action set A forever and the sequence cycling through action b forever after the first transition. The strategic type space for the Email game in Rubinstein (1989a) can be described by attaching transition probabilities to the arrows in Figure 2.2: Transitions labeled (1), (2) and (1) have probability $1 - \epsilon$ while transitions labeled (3) and (2) have probability ϵ . The paths in this type space converge to the common knowledge path in the product topology. However, their limits do not converge. Results in ? and ? suggest that the topology of uniform convergence would ensure strategic continuity for a fixed game.

A related observation is that the sequence of approximating paths requires an ever growing type space while the common knowledge type is binary. The complexity of a type, loosely defined as the minimal size of the type space required to contain this type is therefore not continuous in the product topology. We leave the study of type-complexity for future work.

In this paper we focus on Harsanyi type spaces and not on common prior models. In a companion paper, ? represent common prior models as Markov chains on STS automata and use this representation to provide a finite characterization of rationalizable outcomes in common prior models.

Chapter 3

Information Design for Rationalizability

3.1 Introduction

In strategic interactions under incomplete information, we explore the two pivotal information design questions: How does varying information among players influence the range of distributions over outcomes, and how can we construct an information structure that yields a specific outcome distribution?

We answer these questions when outcomes are defined through the lens of Interim Correlated Rationalizability (ICR henceforth), a concept based on common certainty of rationality between players introduced by [Dekel et al. \(2007a\)](#). These answers are important from a theoretical standpoint as well as for practical applications in information design and the robustness of economic outcomes.

Information design (see e.g. [Kamenica, 2019](#); [Morris et al., 2020](#), for recent surveys) studies the impact of information on outcomes in games. In this literature, implementation is achieved through the dissemination of information to players. As in mechanism design, the question is not just what outcomes can be implemented, but also, for all possible such outcomes, to build a device that implements it.

We develop a toolbox that allows to analyze the recursive structure of ICR in any finite game. Using these tools, we then characterize rationalizable distributions and the corresponding information structures that implement them. Previous important contributions to information design under ICR, such as [Morris et al. \(2020\)](#) and economic applications (see, e.g. [Mathevet et al., 2020a](#); [Halac et al., 2021](#)), require both binary actions and supermodular payoffs. Our new methodology allow to dispense with both assumptions altogether.

The ICR solution concept builds upon the foundations laid by correlated rationalizability for complete information ([Bernheim, 1984](#); [Pearce, 1984](#)), and was expanded to incorporate incomplete information by [Dekel et al. \(2007a\)](#). It is the central concept underlying global games and outcome robustness to incomplete information.

We now discuss information design under Bayes-Nash equilibria to draw a comparison with ICR and highlight the differences in methodologies and outcomes. Information design under Bayes-Nash equilibria, as in correlated equilibria and their variations ([Aumann, 1974, 1987](#); [Forges, 1986](#); [Bergemann and Morris, 2016a](#)), assumes that the information designer not only has control over the dissemination of information among players, but can also impose

coordination on which Bayes-Nash equilibrium of the game with incomplete information they play. In contrast, information design under ICR relies only on the assumption of players' common certainty of rationality, without requiring any form of coordination on equilibrium selection. Thus, ICR provides a stronger form of information design than Bayes-Nash, which explains why fewer outcomes are implementable with the former than with the latter.

Another motivation for focusing on ICR, beyond its intrinsic power, is its fundamental role in addressing the question of robustness. The robustness literature, initiated by [Kajii and Morris \(1997\)](#), explores the persistence of equilibrium outcomes in games under slight information perturbations. This area of research conceptualizes an adversarial information designer tasked with selecting an information structure to destabilize an existing outcome. In these models, an outcome is considered destabilized whenever it is excluded with non-vanishing probability under ICR when the information perturbation is taken arbitrarily small. A comprehensive understanding of game outcomes under ICR enables precise delineation of the power of the adversarial information designer, a crucial step in understanding which outcomes are robust. The current state of the art, as represented by [Oyama and Takahashi \(2020\)](#), provides a full characterization of robust outcomes for the limited class of binary action supermodular games. We view our characterization of implementable outcomes for general games as a promising step toward overcoming these limitations.

We now turn to the description of our main result and its economic relevance, which we follow by an overview of the toolbox we develop and our proof strategy.

Our main result, [Theorem 7](#), is a complete characterization of the set of rationalizable outcome distributions for any finite game. We also provide, for each such rationalizable outcome distribution, the construction of an information structure that implements it.

The class of information structures we construct admits a simple description using finitely many parameters. We provide a representation of this class which coincides with information structures widely used in economic applications: a common state with additive idiosyncratic noise and a finite signal, where the number of signals is bounded by the data of the game. The finite signal corresponds to a recommendation of rationalizable actions for each player. The common state with additive, idiosyncratic noise generates rich higher-order beliefs and allows strategic contagion of actions. Thus, our information structures generalize those commonly studied in the Global Game literature, pioneered by [Carlsson and Van Damme \(1993a\)](#) and [Morris and Shin \(2003b\)](#). We show that the combination of these two devices is enough to generate any rationalizable outcome in any finite game.

By relying on this class of information structures, we characterize the set of rationalizable outcomes using a finite family of linear inequalities, from which it follows that it is a convex polyhedron. The structure of rationalizable outcomes is thus simple and similar to that of correlated equilibrium distributions. This linear structure is also advantageous when considering applications.

We demonstrate that our results provide a powerful tool to study a variety of economic situations where action spaces are naturally large. For example, we examine information disclosure in priority systems and provide an optimal information structure using our results. This optimal information structure can be interpreted as providing private information with a timing friction.

Our proof strategy consists of three main steps, where each step introduces new conceptual tools. In the first one, we describe the set of rationalizable outcomes by studying the entire

process of elimination of dominated strategies when the prior varies. We characterize the processes on ICR hierarchies that arise from *some* information structure through a *revelation principle*. These processes are sufficient to pin down the set of rationalizable outcomes, i.e., of outcomes that survive iterated deletion of dominated strategies. In the second step, we show that these outcomes are all induced by a finite-dimensional parametrized sub-class of processes, which we call SCAMP. Finally, in our third part, we rely on SCAMP to describe rationalizable distributions through a finite family of linear inequalities.

We now detail these three main steps: a revelation principle for ICR hierarchies, SCAMP, and the linear characterization of outcomes.

Every type in a Harsanyi type space induces a sequence of action sets obtained through the elimination of dominated strategies, henceforth referred to as the ICR hierarchy for this type. Therefore, any common prior information structure induces a probability distribution over players' ICR hierarchies and states of nature. Conversely, any probability distribution on profiles of ICR hierarchies and states of nature gives rise to a common prior information structure where the set of ICR hierarchies for a player acts as a set of types for that player.

We show that a distribution over ICR hierarchies arises from a common prior if and only if the ICR hierarchy associated with every player's type is precisely that type. A fundamental consequence is that the set of ICR hierarchies for a player forms a *canonical space* with a *revelation principle*: A distribution on ICR hierarchies is induced by *some* common prior if and only if it is induced by itself viewed as a common prior information structure. The distributions on ICR hierarchies induced by all common prior information structures thus forms a set of *canonical* information structures that is enough to induce all distributions.

Moreover, we characterize canonical information structures through a series of *obedience constraints*. There is one such constraint for every type and every round of elimination of dominated strategies, and the obedience constraints for a type at a certain level of elimination says precisely that the next level of ICR correspond to actions that survive one more round of elimination given the type's beliefs on other types and nature. It follows that these obedience constraints characterize the distributions on ICR hierarchies arising from common priors.

Now that we have described the canonical distributions on ICR hierarchies, our next step is to characterize the limit distributions arising from those, i.e., the set of rationalizable distributions. Since several distributions on ICR hierarchies may, in general, lead to the same rationalizable distribution (where, for instance, the order of elimination of dominated strategies between different rounds differs but the final result is the same), there is a certain degree of freedom in finding such a sub-class of canonical distributions.

We show that, among the class of canonical information structures, all outcome distributions are induced by a subclass which we call Simple Canonical Automaton Markov Priors (or SCAMP for short). Starting with any game, and building on a companion paper [Gossner and Veiel \(2024\)](#) we construct an automaton, given by a finite set of states Ω together with an action set in the game for each player at each state. There is an initial state at which each player is assigned their full action set. Every process on the automaton induces, through the mappings from Ω to action sets, a distribution on sequences of action sets for all players.

We say that a probability distribution on $K \times \Omega^{\mathbb{N}}$, where K denotes the set of states of nature, is a SCAMP if it: 1) satisfies the obedience constraints (in particular, it has support on profiles of ICR hierarchies), 2) is Markovian on Ω , and 3) every path on the automaton passes through at most one non-terminal cycle.

The first property ensures that the distribution of states of nature and hierarchies arises from a canonical prior; it also provides a common prior that implements it, namely the (canonical) distribution on the automaton paths itself. The second property, which builds on the recursive structure of ICR, implies that the class of information structures considered is finitely dimensional and parameterized, and simple to generate. Finally, the third property ensures that the obedience constraints at any round of elimination of dominated strategies only depend on the current ICR set and not on the history of paths leading to it. In particular, it implies a finite number of obedience constraints.

SCAMP is therefore characterized by a finite number of parameters (Markov transitions) as well as a finite number of constraints (the obedience constraints). It thus provides a finitely dimensional parametrized characterization of all rationalizable distributions.

Finally, in our third step we characterize rationalizable outcome distributions as those induced by SCAMP. We show that their set is given by a finite number of linear constraints, hence that rationalizable distributions form a polyhedron. Our methods allow to engineer, for each rationalizable distribution, a SCAMP information structure that induces it. We also show that SCAMP information structures have natural interpretations, both as Additive Noise Information Structure, which is a generalized version of global games [Carlsson and Van Damme \(1993a\)](#) information structures, and as asynchronous information structures in which the common uncertainty arises from players receiving messages at slightly different times, a generalization of [Rubinstein \(1989a\)](#).

The rest of the paper is organized as follows. In [Section 3.2](#) we illustrate our concepts and results in a game of technology adoption. We present the model in [Section 3.3](#) and introduce Strategic Automata as our main tool to represent ICR hierarchies in [Section 3.4](#). Our main results are presented in [Section 3.5](#). In [Section 3.6](#) we apply our methods and results to the design of information in common value priority systems that allows the auctioneer to extract all surplus. Finally, [3.7](#) discusses the complexity and interpretation of SCAMP and relates SCAMP to existing work.

3.2 A Technology Coordination Example

We illustrate the concepts and results of the paper in a game of technology coordination. Two players, 1 (row player) and 2 (column player), each choose between technologies a and b to participate in a joint project.

There are two states of nature. In the bad state, denoted B , the project fails and it is a dominant strategy for each player to stick to their preferred technology: Player 1 has a preference for technology b , and player 2 for a . In the good state, denoted G , the project is successful if players coordinate on the same technology, where each player benefits more from picking a different technology: Player 1 has a preference for technology a , and player 2 for b . Payoffs in that state are those of a battle of sexes.

Consider a discrete set of types T_i for each player i , and a common prior probability P over $\{G, B\} \times T_1 \times T_2$, with marginal having full support on each T_i . A triple k, t_1, t_2 is drawn

according to P , then each player i is informed of her type t_i . We denote conditional beliefs of player i by $p_i = P(\cdot|t_i)$.

Given player 1's beliefs on the state of nature, b dominates a (irrespective of player 2's choices) iff

$$p_1(B) > p_1(G).$$

Note that there are no beliefs of player 1 for which a dominates b . If player 2 plays b , b is a best-response of player 1 for every belief on the state of nature.

For player 2, a dominates b iff

$$p_2(B) > p_2(G),$$

and there are no beliefs such that b dominates a .

For $n \geq 1$, let us denote R_i^n the set of actions that survive n rounds of deletion of dominated strategies given i 's beliefs. We just have established:

$$R_i^1 = \begin{cases} b & \text{if } i = 1 \text{ and } p_i(B) > p_i(G) \\ a & \text{if } i = 2 \text{ and } p_i(B) > p_i(G) \\ ab & \text{if } p_i(B) \leq p_i(G) \end{cases}$$

where for convenience a denotes $\{a\}$, b denotes $\{b\}$ and ab denotes $\{a, b\}$. For the next levels of elimination, simple algebra shows that for player 1:

$$R_1^{n+1} = \begin{cases} a & \text{if } 3p_1(G, R_2^n=a) - p_1(G) > 2p_1(B) - p_1(B, R_2^n=b) \\ b & \text{if } p_1(B, R_2^n=a) + p_1(B) > 2p_1(G) - 3p_1(G, R_2^n=b) \\ ab & \text{otherwise} \end{cases}$$

and for player 2:

$$R_2^{n+1} = \begin{cases} a & \text{if } p_2(R_2^n=b, B) + p_2(B) > 2p_2(G) - 3p_2(R_2^n=a, G) \\ b & \text{if } 3p_2(R_2^n=b, G) - p_2(G) > 2p_2(B) - p_2(R_2^n=a, B) \\ ab & \text{otherwise} \end{cases}$$

A few remarks are in order. As already stated, at the first level, player 1 may eliminate a , but not b , while player 2 may eliminate b but not a . If player 1 doesn't eliminate b at the first level, she may eliminate a at the second level if she believes with high enough probability that player 2 eliminated a at the first level. There are no beliefs at which player 1 eliminates b at the second level while not having eliminated it at the first level. Symmetrically, player 2 may eliminate a at the second level but not at the first. More generally, if $R_1^n = ab$, for n odd we may have $R_1^{n+1} = ab$ or $R_1^{n+1} = a$ but not $R_1^{n+1} = b$ and for n even we may have $R_1^{n+1} = ab$ or $R_1^{n+1} = b$ but not $R_1^{n+1} = a$. A symmetric property holds for player 2.

3.2.1 Strategic Automaton

Call the realizations of the process $(R_i^n)_{n \in \mathbb{N}}$, player i 's ICR hierarchies. The set of possible ICR hierarchies for each player are summarized on the automaton of Figure 3.1. The state labeled with "start for P_i ", is the initial (or 0-th) level of for player i , where $R_i^0 = ab$. The

sequences of state labels starting with the initial state for player i and following the arrows, potentially ending in an absorbing state marked by a double circle, are the sequences $R_i^0 = ab$, $R_i^1, \dots, R_i^n, \dots$ that appear with positive probability in some common prior model.

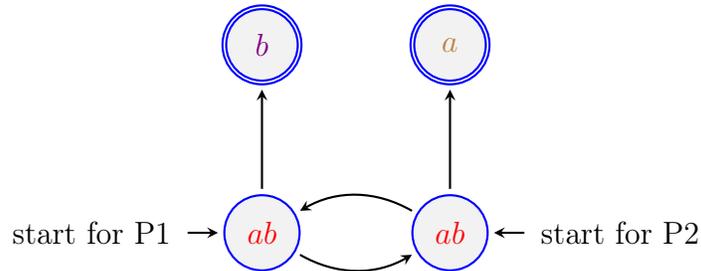


Figure 3.1: Automaton for one player in the technology example. There are 4 states and each state contains an action set. The initial state is on the left (player 1) or on the right (player 2). Double circled states are terminal ones.

Figure 3.2 allows us to visualize the possible joint ICR hierarchies for both players as the set of infinite sequences starting at the initial state.

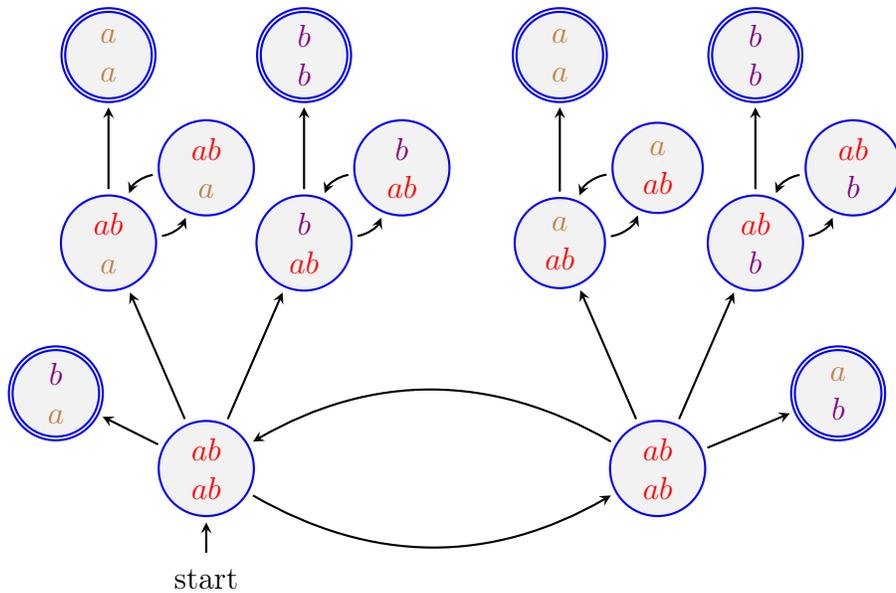


Figure 3.2: Automaton for both players in the technology example. Each state contains an action set for player 1 (top) and for player 2 (bottom). Arrows indicate possible transitions.

Let us call Ω the set of 16 states of Figure 3.2. Each state $\omega \in \Omega$ is labeled with an action set ω_i for each player i . Since every pair of types in a type space can be mapped to a path in the automaton, it follows that every prior on a type space induces a joint probability distribution on $K \times \Omega^{\mathbb{N}}$. We are now asking the question: what is the set of such possible distributions when the prior (and type space) varies?

Note that such a distribution on $K \times \Omega^{\mathbb{N}}$, call it P , is itself an information structure in which $(k, \omega^1, \dots, \omega^n, \dots)$ is drawn according to P , and player i is informed of the corresponding sequence of i -th coordinates $(\omega_i^1, \dots, \omega_i^n, \dots)$.

Our revelation principle (Theorem 5) shows that a distribution P on $K \times \Omega^{\mathbb{N}}$ arises from $(R_i^n)_n$ applied to some common prior on some type space if and only if, in the information structure with prior P , where i is informed of $(\omega_i^1, \dots, \omega_i^n, \dots)$, i 's ICR hierarchy is precisely $(\omega_i^1, \dots, \omega_i^n, \dots)$. From the derivation of ICR hierarchies from types above, this is the case when for every n , and when p_i denotes $P(k, \omega_{-i}^1, \dots, \omega_{-i}^n, \dots | \omega_i^1, \dots, \omega_i^n, \dots)$:

$$\omega_1^{n+1} = \begin{cases} a & \text{if } 3p_1(G, \omega_2^n = a) - p_1(G) > 2p_1(B) - p_1(B, \omega_2^n = b) \\ b & \text{if } p_1(B, \omega_2^n = a) + p_1(B) > 2p_1(G) - 3p_1(G, \omega_2^n = b) \\ ab & \text{otherwise} \end{cases}$$

and for player 2:

$$\omega_2^{n+1} = \begin{cases} a & \text{if } p_2(B, \omega_2^n = b) + p_2(B) > 2p_2(G) - 3p_2(G, \omega_2^n = a) \\ b & \text{if } 3p_2(G, \omega_2^n = b) - p_2(G) > 2p_2(B) - p_2(B, \omega_2^n = a) \\ ab & \text{otherwise} \end{cases}$$

These equations, which we call *Obedience Constraints*, are expressed directly on conditional beliefs p_i , and thus depend only on the probability distribution P on $K \times \Omega^{\mathbb{N}}$.

The revelation principle thus fully characterizes the possible distributions of $(k, (R_1^n)_n, (R_2^n)_n)$ that may arise in any common prior model. We represent them as distributions on $K \times \Omega^{\mathbb{N}}$. The revelation principle also tells us that any such distribution on $K \times \Omega^{\mathbb{N}}$ can be viewed as a *canonical* information structure, i.e. one in which each player i is informed of the sequence of i -th coordinates contained in each state of the automaton, and where, the n -th state contains precisely the n -th ICR sets for both players.

3.2.2 SCAMP

Now that we understand how distributions on ICR hierarchies can be obtained through the automaton, we move on to the characterization of rationalizable distributions. Recall that for a type of player i , the set of rationalizable actions is obtained as $R_i^\infty = \bigcap_n R_i^n$. We say that a distribution μ on $K \times (\{ab, a, b\})^2$ is *rationalizable* if there exists a common prior such that the induced distribution of $(k, R_1^\infty, R_2^\infty)$ is μ .

Our SCAMP revelation principle, Theorem 6 shows that rationalizable distributions are precisely those implemented by a particular type of information structure, called SCAMP, for Simple Canonical Automaton Markov Prior. A SCAMP is a process on the automaton that 1) is Markovian 2) satisfies Obedience Constraints and thus is Canonical and 3) is Simple. We now turn to an explanation of each of these properties and their consequences.

A Markov process is given by a probability on states of nature, and, for each state of the automaton and state of nature, by a transition to states on the automaton. It is thus given by a finite number of parameters only.

Consider a Markov process on the automaton of Figure 3.2. Assume that for some k , the process reaches a state where only one player has eliminated an action, such as a state in which the action sets are ab for player 1 and b for player 2. Then, either the process will cycle between the two states with the same action sets forever, or player 2 will eventually eliminate action b as well. For the sake of the example, we focus on point rationalizable

distributions, whose support is included in $K \times \{a, b\}^2$. These distributions are of interest as they are associated uniquely with an expected payoff in the game. In this case, the distribution on terminal nodes is unchanged by assuming that the first state with action sets (ab, a) transitions directly to the corresponding state with action sets (a, a) without any loops. By removing cycles at all states where only one player has eliminated an action, we obtain a Markov chain of the form in Figure 3.3. Furthermore, it is possible to show that this transformation doesn't violate the obedience constraints whenever they are satisfied by the original process.

Now, for a fixed state of nature, the process cycles between the two lower states a certain number of times, before it exits and reaches a terminal state. We call a process which only passes through a single cycle before reaching a terminal node simple. Conditional on exiting the cycle, the probability of reaching any terminal node is independent of the number of cycles. Hence the probability on terminal nodes of the Markov chain is given by the conditional probability on these nodes after 3 stages of the process.

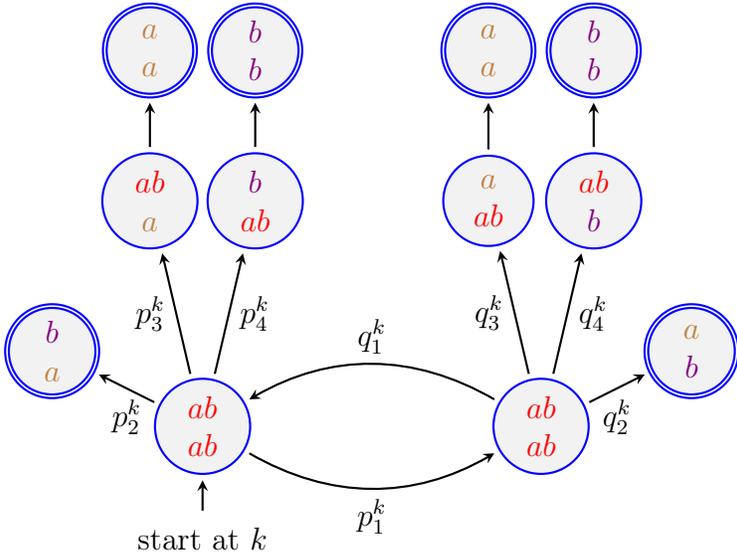


Figure 3.3: SCAMP generating point distributions. When a single arrow leaves a state, this arrow has probability 1. Transitions may depend on the state of nature $k \in \{G, B\}$.

Simplicity thus allows to compute the implemented distribution from the distribution in a finite number of iterations, 3 in this example. Furthermore, we show that whenever a distribution satisfies the OC on the first iterations of the process, there exists a SCAMP that yields the same outcome distribution on terminal nodes. Therefore, all that needs to be done is to characterize the set of possible distributions that satisfy the OC on the first iterations - in our example, the set of distributions on $(k, \omega^1, \omega^2, \omega^3)$ that satisfy OC. Since OC are linear inequalities, this yields a characterization of the set of rationalizable distributions as a (not necessarily closed) convex polyhedron. For the technology adoption game, we illustrate the corresponding payoffs generated by point distributions in Figure 3.4.

The set of point rationalizable distributions, hence their payoffs, is a subset of (agent normal form) correlated equilibria (see Forges, 1993a) and of their payoffs. Both payoff sets are subsets of the set of feasible payoffs.

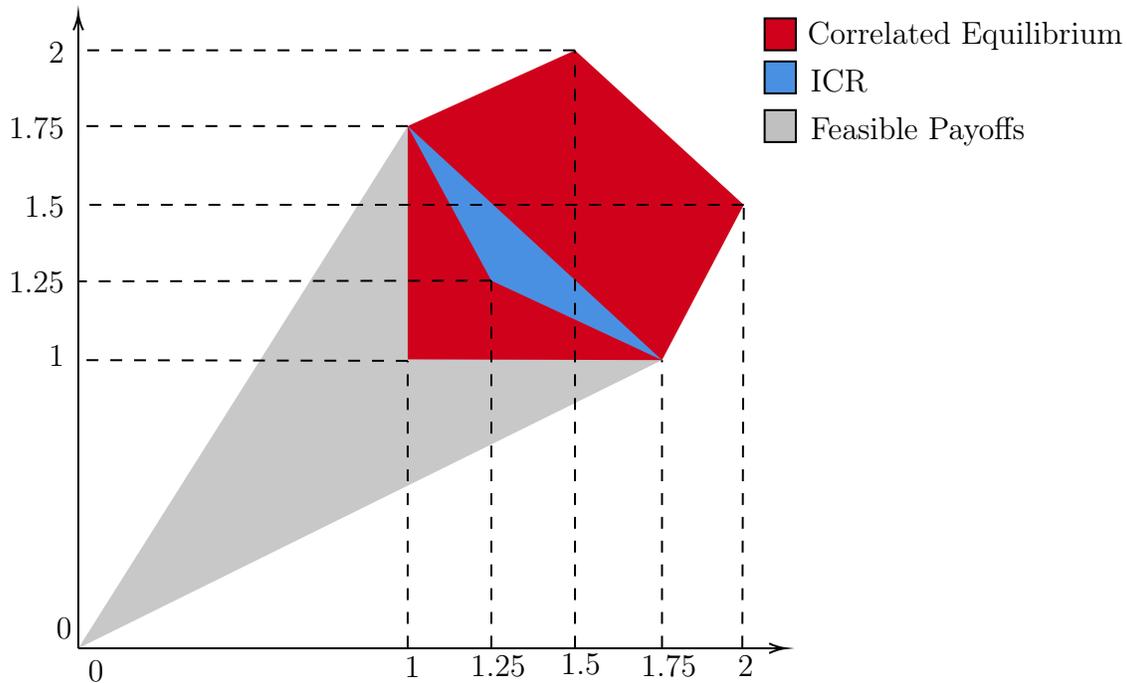


Figure 3.4: Point rationalizable payoffs (in blue), correlated equilibrium payoffs (in red) and feasible payoffs (in grey) in the coordination game.

3.2.3 Additive, Idiosyncratic Noise Representation

Paths in SCAMP are described by trees with branches containing at most one cycle. Paths embed two sources of uncertainty: 1) about branches of the tree 2) about how many iterations of the cycle the path looped through. These uncertainties have a dynamic interpretation of information dissemination in which player receive correlated private signal from a finite set corresponding to the branch and are also uncertain about the timing of these signals (the number of loops).

In our technological coordination game, each player receives a recommendation to play either a or b corresponding to the terminal state of the automaton. But a player, being informed only of the ICR sequence, doesn't know the full path of the automaton. This means that, for instance, player 1 receiving a sequence of signals $(ab, \dots, ab, a, a, \dots)$ doesn't know which of the two branches leading to action a was selected. In one of them, player 1 eliminated action b before player 2, in the second she eliminated b after player 2, and in the third she eliminated b at the same round as player 1 eliminated player 1.

The round at which a player's type transitions can be modeled as a stopping time. Note that upon receiving her action recommendation, each player assigns probability one to the other player having received her recommendation at the same time has hers \pm one round. Players' private signals are thus slightly out of sync. This is very similar to [Rubinstein \(1989a\)](#)'s email game, in which a player doesn't know whether her message to the other, or the other's confirmation was lost first. This is also similar to Global Games ([Carlsson and Van Damme, 1993a](#)), where players are uncertain about their order in eliminating dominated

strategies, but are certain about the outcome. In fact, SCAMP builds on, and generalizes, both these types of information structures.

3.3 Model

General notations For any mapping $f: X \rightarrow Y$ and any subset $E \subseteq X$ we write $f(E) := \{f(x) : x \in E\}$. For a family of sets $(X_i)_i$, we let $X = \prod_i X_i$ and $X_{-i} = \prod_{j \neq i} X_j$, for every i . Given a measurable set X , $\Delta(X)$ denotes the set of probability distributions on X . \mathbb{N} denotes the non-negative integers.

For a family of maps $(f_i: X_i \rightarrow Y_i)_i$, we let $f: X \rightarrow Y$ be given by $f(x) = (f_i(x_i))_i$ for $x \in X$ and $f_{-i}: X_{-i} \rightarrow Y_{-i}$ by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$ for $x_{-i} \in X_{-i}$.

A marginal on coordinates x_1, \dots, x_n of a distribution $P \in \Delta(\prod_\ell X_\ell)$ is denoted $\text{marg}_{x_1, \dots, x_n}(P)$.

Games with incomplete information We fix a finite set N of players and a finite set K of states of nature. We also fix a *payoff structure* u , given by a finite action set A_i and a payoff function $u_i: K \times A \rightarrow \mathbb{R}$, for each player i . A *common prior*, denoted P , is given by a family of measurable type spaces $(T_i)_{i \in N}$, a probability distribution P on $K \times T$.

A *game with incomplete information* is a pair (u, P) , where u is a payoff structure and P is a common prior.

Interim Correlated Rationalizability Given a game with incomplete information (u, P) with type spaces $(T_i)_i$, Interim Correlation Rationalizability (ICR) (Dekel et al., 2007a) is the outcome of the process of elimination of dominated strategies, or equivalently never best responses, applied to every type of every player. It is defined as follows.¹ Let \mathcal{A}_i denote the collection of non-empty subsets of A_i , and define a *conjecture for player i* as a map $\sigma_i: K \times \mathcal{A}_{-i} \rightarrow \Delta(A_{-i})$ such that the support of $\sigma_i(k, B_{-i})$ is included in B_{-i} for every $k \in K$ and $B_{-i} \in \mathcal{A}_{-i}$. A belief $p \in \Delta(K \times \mathcal{A}_{-i})$ and a conjecture σ_i induce a probability distribution $\langle \sigma_i, p \rangle \in \Delta(K \times A_{-i})$ satisfying

$$\langle \sigma_i, p \rangle(k, a_{-i}) = \sum_{B_{-i} \in \mathcal{A}_{-i}} p(k, B_{-i}) \sigma_i(k, B_{-i})(a_{-i}), \quad \forall (k, a_{-i}) \in K \times A_{-i}. \quad (3.1)$$

Player i 's best-reply map $\text{BR}_i: \Delta(K \times \mathcal{A}_{-i}) \rightarrow \mathcal{A}_i$ is defined by:

$$\text{BR}_i(p) := \bigcup_{\sigma_i} \left\{ \arg \max_{a_i \in A_i} \mathbb{E}_{\langle \sigma_i, p \rangle} u_i(\cdot, a_i, \cdot) \right\}, \quad (3.2)$$

where $\mathbb{E}_{\langle \sigma_i, p \rangle}$ is the expectation with respect to the distribution $\langle \sigma_i, p \rangle$ and the union is taken over all conjectures of player i . The ICR *hierarchy* $(R_i^m(t_i))_{m \geq 0}$ of a type $t_i \in T_i$ is defined iteratively:

1. For every $i \in N$ and $t_i \in T_i$, $R_i^0(t_i) := A_i$,

¹Our presentation slightly differs from (Dekel et al., 2007a) but the two definitions are equivalent.

2. For every $m \in \mathbb{N}$, the conditional probability $P(\cdot|t_i) \in \Delta(K \times T_{-i})$ and the (measurable) map R_{-i}^m from T_{-i} to \mathcal{A}_{-i} together induce a belief $P(\cdot|t_i) \circ (\text{id} \times R_{-i}^m)^{-1}$ on $K \times \mathcal{A}_{-i}$. Given this belief, $R_i^{m+1}(t_i)$ is defined as the set of best responses to this belief:

$$R_i^{m+1}(t_i) := \text{BR}_i(P(\cdot|t_i) \circ (\text{id} \times R_{-i}^m)^{-1}). \quad (3.3)$$

The set of rationalizable actions associated to t_i is then

$$R_i^\infty(t_i) := \bigcap_{m \in \mathbb{N}} R_i^m(t_i). \quad (3.4)$$

The *outcome distribution* μ_P on $K \times B$ induced by P through ICR is given by:

$$\mu_P := P \circ (\text{id} \times R^\infty)^{-1}. \quad (3.5)$$

Two priors P, P' are outcome equivalent if $\mu_P = \mu_{P'}$.

3.4 Rationalizable Hierarchies

In this section, we show that ICR hierarchies possess a recursive structure that is characterized by a finite automaton, and characterize the distributions of ICR hierarchies that arise as we vary the common prior.

3.4.1 Revelation principle

Every common prior P induces, through the identity on K and ICR a distribution $P_R := P \circ (\text{id} \times (R^m)_m)^{-1}$ on $K \times \mathcal{A}^\mathbb{N}$. Let \mathcal{P} denote the set of such distributions when we vary over all common priors on all possible type spaces. We define the *set of ICR hierarchies* as the set of all possible hierarchies that arise as we vary the common prior.

Definition 6 (ICR Hierarchies). *The set of ICR hierarchies is the subset $S \subseteq \mathcal{A}^\mathbb{N}$ defined by the following condition: $s \in S$ if and only if there exists $P \in \mathcal{P}$ so that $P(s) > 0$.*

Note that for every common prior P , the induced distribution $P_R \in \mathcal{P}$ can be seen itself as a common prior in which the set of types for player i is contained in $\mathcal{A}_i^\mathbb{N}$, ($k, s = (s_i^n)_{i,n}$) is drawn according to P_R , and each player i is informed of her corresponding type $s_i = (s_i^n)_n \in \mathcal{A}_i^\mathbb{N}$. We now show that the two games with incomplete information, (u, P) and (u, P_R) , are the same in the sense that

$$(P_R)_R = P_R. \quad (3.6)$$

Moreover, any distribution on $K \times \mathcal{A}^\mathbb{N}$ that satisfies this idempotency property is in \mathcal{P} . The following result thus characterizes \mathcal{P} .

Theorem 5 (Revelation Principle). *For any $P \in \Delta(K \times \mathcal{A}^\mathbb{N})$ the three conditions are equivalent:*

1. $P \in \mathcal{P}$;

2. $R(s) = s$, *P a.s.* ;

3. $P(s^0=A) = 1$ and P satisfies the family of obedience constraints:

$$s_i^m = \text{BR}_i(\text{marg}_{k, s_{-i}^{m-1}} P(\cdot, \cdot | s_i)) \text{ for a.e. } s_i = (s_i^m)_m. \quad (3.7)$$

Condition 2 establishes the pointwise idempotency property that characterizes \mathcal{P} and condition 3 provides the algebraic expression (obedience constraints (OC)) that characterizes the idempotency condition at each level of the hierarchy. An implication of this result is that every outcome distribution can be achieved by picking an information structure in \mathcal{P} , which is why we refer to it as a ‘‘Revelation Principle.’’ The proof of Theorem 5 can be found in Appendix C.1.1.

3.4.2 Strategic Automaton

An *automaton* is a triple (Ω, β, \preceq) given by a finite set of states Ω together with an action map $\beta_i: \Omega \rightarrow \mathcal{A}_i$ for every player i and a binary successor relation \preceq on states. Figure 3.2 represents an automaton, where each circle is a state, the action sets in each circle represent the labels of the action map and the arrows represent the successor relation, where each state points to its successors.

A *cycle* is an ordered collection of states $c = \{\omega^1, \dots, \omega^n\} \subseteq \Omega$ so that $\omega^h \preceq \omega^{h+1}$ for all $h < n$ and $\omega^1 = \omega^n$. A *path* is a sequence $(\omega^0, \omega^1, \dots)$ satisfying $\omega^m \preceq \omega^{m+1}$ for all $m \in \mathbb{N}$.

We define a *strategic automaton* as an automaton for which the set of paths contains the set of ICR hierarchies.

Definition 7 (Strategic Automaton). *A strategic automaton is an automaton (Ω, β, \preceq) such that:*

$$S \subseteq \{(\beta(\omega^m))_m : \forall m \geq 0, \omega^m \preceq \omega^{m+1}\}.$$

Since we only require containment, we can consider an automaton where every state is associated to an action set profile in \mathcal{A} , and $(B_i)_i \preceq (B'_i)_i$ if and only if $B_i \supseteq B'_i$ for every i . The existence of a strategic automaton is thus trivial. Figure 3.2 represents a strategic automaton in which reverse containment also holds.

3.5 Rationalizable Distributions

We now introduce SCAMP and establish the finite-dimensional structure of all rationalizable outcomes and information structures used to obtain them.

3.5.1 SCAMP

A path on a strategic automaton (Ω, β, \preceq) , gives rise for each player i to a sequence of action sets $s_i = (\beta_i(\omega^0), \beta_i(\omega^1), \dots)$. A *process* on the automaton is a probability measure $P \in \Delta(K \times \Omega^{\mathbb{N}})$ so that every sequence $(\omega^0, \omega^1, \dots)$ in its support is a path. A process on the

automaton defines a common prior, where each player i is privately informed of the sequence s_i . A process on the automaton P is *Markov* if for every $m \in \mathbb{N}$,

$$P(\omega^{m+1}|k, \omega^0 \dots \omega^m) = P(\omega^{m+1}|k, \omega^m), \quad P a.s. \quad (3.8)$$

We will thus refer to the sequence s_i as a (canonical) type. A process is *canonical* if

$$R(s) = s, \quad P a.s. \quad (3.9)$$

Canonical Markov processes are thus contained in \mathcal{P} . Moreover, for canonical Markov processes, the state ω^m at round m of a path is a sufficient statistic for the distribution over ω^{m+1} and thus also over R^{m+1} .

A path loops through a cycle if the path visits some element in the cycle more than once. A cycle is *terminal* for a path if the path visits the cycle infinitely often.

A process P is *simple* if every path in its support loops through at most one cycle that is not terminal.

Definition 8 (Simple Canonical Automaton Markov Priors, SCAMP). *A process on a Strategic Automaton is SCAMP if it is Simple, Canonical and Markov.*

3.5.2 SCAMP Sufficiency

The main result of this section is the existence of a strategic automaton on which SCAMP is sufficient to obtain all outcome distributions.

Theorem 6 (Sufficiency of SCAMP). *There exists a strategic automaton on which SCAMP induces all outcome distributions.*

To prove Theorem 6, we first restrict to a subclass of “well behaved” canonical priors, which we call “seeded” and show that this class induces all possible outcome distributions. We then construct a strategic automaton so that all seeded priors can be transformed into an outcome-equivalent SCAMP. The transformation involves averaging probabilities of transitions on the newly constructed automaton in order to make the new process Markov. With the right automaton, this procedure is very simple and is illustrated in the example below:

Example Figure 3.5 below describes the transition probabilities of a canonical prior P on a strategic automaton of an unspecified payoff structure. For every $m \in \mathbb{N}$ and every k , $p_{k,m}^1 > 0$ is the probability of making the left-most transition conditional on cycling for m rounds and conditional on state $k \in K$,

$$p_{k,m}^1 = P(s^{m+1} = (ab, a)|k, s^n = (ab, ab), \forall n \leq m). \quad (3.10)$$

The transition probabilities $p_{k,m}^2, \dots, p_{k,m}^4$ are defined similarly and are all assumed to be positive for all m and k . For every $m \in \mathbb{N}$ and state $k \in K$, let the probability of cycling m times conditional on k be given by $P(m|k) = \prod_{n \leq m} (1 - \sum_{l=1}^4 p_{k,m}^l)$. We further assume that there exists some state of nature for which this probability is positive for all m .

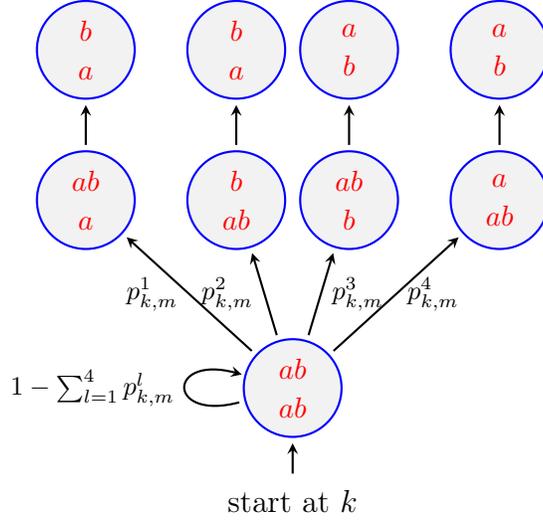


Figure 3.5: Canonical Prior as Process on SCAMP automaton.

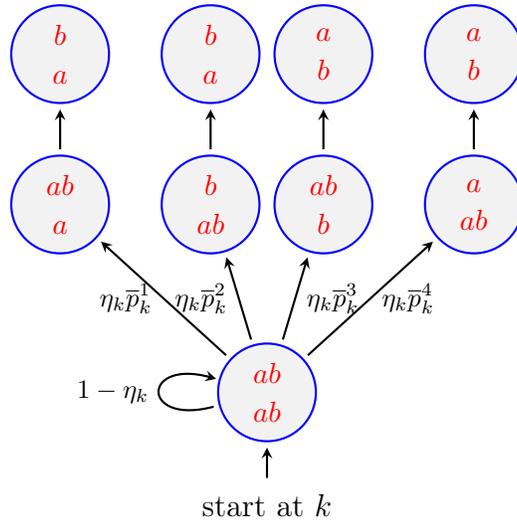


Figure 3.6: Averaging process to SCAMP.

We now construct a SCAMP, represented by the transition probabilities in Figure 3.6, which induces the same outcome distribution as the original process P . The Markov transitions are parametrized by a family $\eta_k, \bar{p}_k^1, \bar{p}_k^2, \bar{p}_k^3, \bar{p}_k^4$ and depend on k . We need $\bar{p}_k = (\bar{p}_k^1, \dots, \bar{p}_k^4)_k$ and $(\eta_k)_k$ to satisfy the following three conditions²:

²We omit symmetric obedience constraints: The type of player 1 transitioning to b at the first round, the types of player 2 making a transition away from ab at round one, the types for player 1 transitioning to b at later rounds and the types of player 2 making their transition away from ab at later rounds.

1. Outcome Equivalence:

$$\begin{aligned}\bar{p}_k^1 + \bar{p}_k^2 &= (p_{k,1}^1 + p_{k,1}^2) + \sum_{m=2}^{\infty} (p_{k,m}^1 + p_{k,m}^2) P(m-1|k), \\ \bar{p}_k^3 + \bar{p}_k^4 &= (p_{k,1}^3 + p_{k,1}^4) + \sum_{m=2}^{\infty} (p_{k,m}^3 + p_{k,m}^4) P(m-1|k).\end{aligned}\tag{3.11}$$

2. Obedience of the path where player 1 (whose action sets are in the top entry of each node in Figure 3.6) transitions to a at the first round:

$$\sum_{k \in K} P(k) \eta_k \bar{p}_k^4 (u_1(k, a, \alpha_2) - u_1(k, b, \alpha_2)) > 0, \quad \forall \alpha_2 \in \{a, b\}.\tag{3.12}$$

3. Obedience of types transitioning to a at round $m+1$ for player 1:

$$\begin{aligned}\sum_{k \in K} P(k) (\eta_k (1 - \eta_k)^{m-1} \bar{p}_k^3 (u_1(k, a, b) - u_1(k, b, b)) \\ + \eta_k (1 - \eta_k)^m \bar{p}_k^4 (u_1(k, a, \alpha_2) - u_1(k, b, \alpha_2))) > 0, \\ \forall \alpha_2 \in \{a, b\}, \quad \forall m \geq 1.\end{aligned}\tag{3.13}$$

Since P satisfies obedience, we know that

$$\sum_{k \in K} P(k) p_k^4 (u_1(k, a, \alpha_2) - u_1(k, b, \alpha_2)) > 0, \quad \forall \alpha_2 \in \{a, b\}.\tag{3.14}$$

In this example, suppose that the payoff structure has the following property³: K has five elements $\{k_{ab,a}, k_{b,ab}, k_{ab,b}, k_{a,ab}\}$, where for every $\alpha \in \{a, b\}$, at state $k_{ab,\alpha}$, action α is a dominant action for player 2 while player 1 has no dominant action; and at state $k_{\alpha,ab}$, action α is a dominant action for player 1 while player 2 has no dominant action.

Consider setting $\eta_k = 1$ at every state $k \in K$ where some player has a dominant action - and setting $\eta_k = \zeta \in (0, 1)$ otherwise. The resulting Markov chain will have cycles at states of nature where there is no dominant action for either player. By convexity of the OCs, condition (iii) is satisfied when choosing $\bar{p}_k^l = p_k^l + \sum_{m=2}^{\infty} p_{k,m}^l P(m-1|k)$ at each $l = 1, \dots, 4$ and any such choice of $(\eta_k)_k$. Note that this choice of \bar{p} also satisfies condition (i), which is independent of $(\eta_k)_k$. Given our choice of \bar{p} , there must be ζ small enough so that

$$\zeta \bar{p}_k^1 \leq p_k^1, \quad \forall k \text{ s.t. } \min_{\alpha_2 \in \{a, b\}} (u_1(k, a, \alpha_2) - u_1(k, b, \alpha_2)) < 0.\tag{3.15}$$

Hence $\eta_k \in \{\zeta, 1\}$ for all k and some small enough constant $\zeta \in (0, 1)$ defines a SCAMP that is outcome equivalent to P .

An important challenge in the general construction is that not all priors are as well behaved as the one in the example above. The paths of a prior P may loop through multiple

³Through the obedience constraints, the positivity assumptions on the transition probabilities of Figure 3.5 require that K contain at least those five states.

cycles, they may pass through several states before and after each cycle, and some transitions may only occur at even rounds while others only occur at odd rounds, for example. As we have seen in the technology coordination example of Section 3.2, the automaton also needs to be rich enough for the procedure to work, containing possibly different states with the same action profile label. To approach this problem, we restrict attention to a well-behaved subclass of canonical priors, which we call “seeded priors”. We then construct an automaton that is rich enough so that the averaging procedure described in the example above yields an outcome-equivalent SCAMP for every prior in the class.

Semi-Lattice Priors

We first restrict attention to a well-behaved subclass \mathcal{P}^* of priors in \mathcal{P} , which we call *Semi-Lattice priors*. Priors in this class are supported by a collection of hierarchies that have a semi-lattice property whenever sequences are “comparable.” To make this property precise we start by defining two notions of comparability for hierarchy-pairs. We say that two hierarchies $s, \hat{s} \in S$ are *jointly comparable* if they pass through the same action-set profiles, i.e.

$$\{s^m : m \in \mathbb{N}\} = \{\hat{s}^m : m \in \mathbb{N}\}. \quad (3.16)$$

They are *individually comparable* if, for every player i

$$\{s_i^m : m \in \mathbb{N}\} = \{\hat{s}_i^m : m \in \mathbb{N}\}. \quad (3.17)$$

We now introduce the semi-lattice property on priors by defining the following partial order on comparable hierarchies

$$s \subseteq_S \hat{s} \iff s_i^m \subseteq \hat{s}_i^m, \forall m \in \mathbb{N}, \forall i. \quad (3.18)$$

For every $k \in K$, let $S_P(k) := \{s \in S : P(k, s) > 0\}$. A prior $P \in \mathcal{P}$ is *jointly (individually) closed from below* if for any $k \in K$ and any jointly (individually) comparable sequences $s, \hat{s} \in S_P(k)$, there exists $\underline{s} \in S_P(k)$ that is jointly (individually) comparable to both s and \hat{s} , and satisfies

$$\underline{s} \subseteq_S s \text{ and } \underline{s} \subseteq_S \hat{s}. \quad (3.19)$$

Definition 9 (Semi-Lattice Prior). *A prior $P \in \mathcal{P}$ is a semi-lattice prior if it is both jointly and individually closed from below.*

we let $\mathcal{P}^* \subseteq \mathcal{P}$ denote the collection of semi-lattice priors.

Lemma 10 (Semi-Lattice Prior). *For every $P \in \mathcal{P}$ there exists an outcome equivalent, Semi-Lattice Prior $P^* \in \mathcal{P}^*$.*

We prove Lemma 10 by applying Knaster-Tarski’s fixed point theorem and exploiting a key monotonicity property of BR, which we state in Claim 11 in the appendix.

Fix any $P \in \mathcal{P}^*$. For $k \in K$ and every $s \in S_P(k)$, let $S_P(k, s) \subseteq S_P(k)$ denote the collection of hierarchies in $S_P(k)$ that are jointly comparable to s . Since P is a semi-lattice prior, there exists a unique map $\chi_k : S_P(k) \rightarrow S_P(k)$ satisfying

$$\chi_k(s) \in \inf_{\subseteq_S} S_P(k, s). \quad (3.20)$$

Define the collection of jointly minimal hierarchies

$$\underline{S}_P(k) := \{\chi_k(s) : s \in S_P(k)\}, \text{ and } \underline{S}_P := \bigcup_{k \in K} \underline{S}_P(k). \quad (3.21)$$

Lemma 11 (Finiteness). *There exists finite $\underline{S} \subseteq S$ and $m^* \in \mathbb{N}$ so that*

(i) *for every $P \in \mathcal{P}^*$, $\underline{S}_P \subseteq \underline{S}$,*

(ii) *for all $s \in \underline{S}$ and all $n \in \mathbb{N}$, $s^{m^*+n} = s^{m^*}$.*

For any $l \in \mathbb{N}^{m^*}$ and any $s \in \underline{S}_P$, define the sequence $s \otimes l \in S$, obtained from s by making l^n extra copies of s^{n-1} , for every $n \in \{1, \dots, m^*\}$:

$$s \otimes l := (s^0, \underbrace{s^0, \dots, s^0}_{l^1 \text{ times}}, \dots, s^{m^*-1}, \underbrace{s^{m^*-1}, \dots, s^{m^*-1}}_{l^{m^*} \text{ times}}, s^{m^*}, \dots). \quad (3.22)$$

Bootstraps

Let \mathcal{L} denote the collection of mappings $\ell : K \times \underline{S} \rightarrow \{0, 1\}^{m^*}$ so that for every $(k, s) \in K \times \underline{S}$ and $m < m^* - 1$,

$$\ell^m(k, s) = \ell^{m+2}(k, s) = 1 \implies \ell^{m+1}(k, s) = 1. \quad (3.23)$$

Moreover, let \mathcal{T} denote the collection of all maps $\tau : \mathcal{L} \times K \times \underline{S} \rightarrow \{0, \dots, m^*\}$. For any family of maps $L \subseteq \mathcal{L}$, $(k, s) \in K \times \underline{S}$, and $\tau \in \mathcal{T}$, define

$$\langle \tau \rangle_L(k, s) := \sum_{\ell \in L} \ell(k, s) \tau(\ell, k, s). \quad (3.24)$$

Given $L \subseteq \mathcal{L}$, profile $(\tau_i)_i \in \mathcal{T}^I$ and $(k, s) \in K \times \underline{S}$ define the induced hierarchies

$$\mathcal{I}_{L, (\tau_i)_i}(k, s) := \{(\hat{k}, \hat{s}) \in K \times \underline{S} : \hat{s} \in \underline{S}_P(\hat{k}), \forall i [\hat{s} \otimes \langle \tau_i \rangle(\hat{k}, \hat{s})]_i = [s \otimes \langle \tau_i \rangle(k, s)]_i\}. \quad (3.25)$$

For any $k \in K$ and $s \in S_P(k)$, define the minimal information set

$$\underline{S}_P(s) := \{(\hat{k}, \chi_{\hat{k}}(\hat{s})) \in K \times S : \hat{s} \in S_P(\hat{k}), \exists i \text{ s.t. } P(\hat{k}, \hat{s}|s_i) > 0\}. \quad (3.26)$$

Definition 10 (Bootstrap). *A collection of maps $L \subseteq \mathcal{L}$ is a bootstrap for P if for every $(k, s) \in K \times \underline{S}$ satisfying $s \in S_P(k)$, there exists a profile $(\tau_i)_i \in \mathcal{T}^I$, and $s^* \in S_P(k)$ satisfying $\chi_k(s^*) = s$ so that,*

$$\mathcal{I}_{L, (\tau_i)_i}(k, s) = \underline{S}_P(s^*). \quad (3.27)$$

Let \mathcal{L}_P denote the collection of bootstraps for P . The example below illustrates why there is typically more than one bootstrap for a prior P .

For any $(k, s), (\hat{k}, \hat{s}) \in K \times \underline{S}$ and player i , say that two bootstraps $L, L' \in \mathcal{L}_P$ have *overlap* at $(k, s, \hat{k}, \hat{s}, i)$ if there exists $\tau \in \mathcal{T}$ so that

$$[\hat{s} \otimes \langle \tau \rangle_L(\hat{k}, \hat{s})]_i = [s \otimes \langle \tau \rangle_{L'}(k, s)]_i \quad (3.28)$$

Lemma 12 (Overlap). $L, L' \in \mathcal{L}_P$ have overlap at $(k, s, \hat{k}, \hat{s}, i)$ if and only if they have overlap at (k, s, k, s, i) and $(\hat{k}, \hat{s}, \hat{k}, \hat{s}, i)$.

By appealing to Lemma 12, we will say that $L, L' \in \mathcal{L}_P$ are *overlapping* at $(k, s, i) \in K \times \underline{S} \times I$ if L and L' have overlap at (k, s, k, s, i) .

Let \sqsubseteq be the partial order on \mathcal{L}_P , where $L \sqsubseteq L'$ if and only if

(i) for every $\ell \in L$ there exists $\ell' \in L'$ so that for every $(k, s) \in K \times \underline{S}$,

$$\ell(k, s) \leq \ell'(k, s), \quad (3.29)$$

(ii) for every $\ell' \in L'$ there exists $\ell \in L$ so that for every $(k, s) \in K \times \underline{S}$,

$$\ell(k, s) \leq \ell'(k, s). \quad (3.30)$$

Lemma 13 (Semi-Lattice Structure of Overlapping Bootstraps). *For every L, L' that are overlapping at (k, s, i) , there exists $L'' \in \mathcal{L}_P$ so that L and L'' are overlapping at (k, s, i) and*

$$L'' \sqsubseteq L \text{ and } L'' \sqsubseteq L'. \quad (3.31)$$

SCAMP

We thus obtain an automaton with states

$$\Omega_P^* := \{(L, m, s) : L \in \mathcal{L}_P, s \in \underline{S}, m \leq m^*\}. \quad (3.32)$$

For any $L \in \mathcal{L}_P$ and $\ell \in L$, the cyclic transition is then given by the following two coordinates

$$\begin{aligned} \overline{m}_\ell(k, s) &:= \begin{cases} \max\{m \in \{0, \dots, m^*\} : \ell^m(k, s) = 1\}, & \text{if } \max_{h \leq m^*} \ell^h(k, s) > 0, \\ -1, & \text{otherwise,} \end{cases} \\ \underline{m}_\ell(k, s) &:= \begin{cases} \min\{m \in \{0, \dots, m^*\} : \ell^m(k, s) = 1\}, & \text{if } \max_{h \leq m^*} \ell^h(k, s) > 0, \\ -1, & \text{otherwise.} \end{cases} \end{aligned} \quad (3.33)$$

For every $L \in \mathcal{L}_P$, every $k \in K$ and $s \in \underline{S}$, define the set of sequences

$$\mathcal{S}_P(L, k, s) := \{\hat{s} \in S_P(k) : \exists (\tau_i)_i \in \mathcal{T}^I \text{ s.t. } \mathcal{I}_{L, (\tau_i)_i}(k, s) = \underline{S}_P(s^*)\}. \quad (3.34)$$

Fix any $\eta > 0$ and define a Markov chain $\overline{P}_k : \Omega_P^* \rightarrow \Delta(\Omega_P^*)$ as follows

$$\overline{P}_k(L', m', s' | L, m, s) := \begin{cases} \mu_L(s|k) \rho_L(m'|k, m, s), & \text{if } m = 0, (L, s) = (L', s') \\ \rho_L(m'|k, m, s), & \text{if } 0 < m < m^*, (L, s) = (L', s') \\ 1, & \text{if } m = m' = m^*, (L, s) = (L', s'), \end{cases} \quad (3.35)$$

where

$$\mu_L(\underline{s}|k) := P(\mathcal{S}_P(L, k, s) | k)$$

$$\rho_L(m'|k, m, s) := \begin{cases} \eta, & \text{if } \exists \ell \in L \text{ s.t. } m = \overline{m}_\ell(k, s), m' = m + 1 \\ 1 - \eta, & \text{if } \exists \ell \in L \text{ s.t. } m = \overline{m}_\ell(k, s), m' = \underline{m}_\ell(k, s) \\ 1, & \text{if } m \neq \overline{m}_\ell(k, s), \forall \ell \in L, m' = m + 1 \\ 0, & \text{otherwise.} \end{cases} \quad (3.36)$$

We thus obtain a successor relation on Ω_P^* given by \preceq_P^* , which for every $\omega, \omega' \in \Omega_P^*$ satisfies

$$\omega \preceq_P^* \omega' \iff \exists k \in K \text{ s.t. } \bar{P}_k(\omega'|\omega) > 0. \quad (3.37)$$

Moreover, we obtain an action-set label in the natural way: For every $(L, m, s) \in \Omega_P^*$,

$$\beta_P^*(L, m, s) = s^m. \quad (3.38)$$

Then for every $L \in \mathcal{L}_P, \underline{s} \in \underline{S}$ and $\alpha: L \times K \times \underline{S} \rightarrow \mathbb{N}_{>0}$, we obtain the following SCAMP:

$$\bar{P}_\eta \left(k, L, \underline{s} \otimes \sum_{\ell \in L} \alpha(\ell, k, \underline{s}) \ell(k, \underline{s}) \right) = (1 - \eta)^{\sum_{\ell \in L} \alpha(\ell, k, \underline{s})} \eta^{d_P(L, k, \underline{s})} \mu_L(\underline{s}|k) \text{marg}_k P(k), \quad (3.39)$$

where

$$d_P(L, k, \underline{s}) := \left| \left\{ \ell \in L : \sum_{m \leq m^*} \ell^m(k, \underline{s}) > 0 \right\} \right|. \quad (3.40)$$

Obedience The set of conjectures on automaton states Σ is the set of random selections $\sigma: K \times \mathcal{A}_{-i} \rightarrow \Delta(A_{-i})$, where for every k, B , $\sigma(B_{-i}|k, B_{-i}) = 1$. For every $(k, s) \in K \times S$, player i and round m , define the minimal and maximal conjectured payoff increments between $a_i, a'_i \in A_i$ respectively as,

$$\begin{aligned} \underline{u}_{i, a_i, a'_i}^m(k, s) &:= \min_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s_{-i}^m} \sigma(a_{-i}|k, s^m) (u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})), \\ \bar{u}_{i, a_i, a'_i}^m(k, s) &:= \max_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s_{-i}^m} \sigma(a_{-i}|k, s^m) (u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})). \end{aligned} \quad (3.41)$$

The expected payoff increments under any prior $P \in \mathcal{P}$ are denoted

$$\underline{U}_{i, P}^m(s_i, a_i, a'_i) := \sum_{k \in K, \hat{s} \in \underline{S}_{i, P}(k, s_i)} \underline{u}_{i, a_i, a'_i}^m(k, \hat{s}) P(k, \hat{s}), \quad (3.42)$$

Say that P satisfies *sub-obedience* at round m and s_i if for every $a'_i \in s_i^m \setminus s_i^{m+1}$, there is $a_i \in s_i^{m+1}$ so that

$$0 < \underline{U}_i^m(s_i, a_i, a'_i). \quad (3.43)$$

Say that P satisfies *extreme obedience* if P satisfies obedience at $m = 1$ and in the limit: For P -almost every s and any player i ,

$$\lim_{m \uparrow \infty} s_i^m = \lim_{m \uparrow \infty} \text{BR}_i(\text{marg}_{k, s_{-i}^{m-1}} P(\cdot, \cdot | s_i)). \quad (3.44)$$

The following result follows immediately from the monotonicity of BR_i :

Lemma 14. *If $P \in \Delta(K \times S)$ satisfies sub-obedience and extreme obedience then $P \in \mathcal{P}$.*

Finally, by the semi-lattice structure of bootstraps we conclude that if P satisfies extreme obedience, so does \bar{P}_η , for $\eta > 0$ small enough:

Lemma 15 (Extreme Obedience). *There exists $\eta > 0$ small enough so that \bar{P}_η satisfies extreme obedience.*

3.5.3 Characterization of Rationalizable Outcomes

We now show that the set of outcomes is characterized by a convex polyhedron. Let m^* be the length of branches used in the construction of the Markov process above. A distribution $p \in \Delta(K \times \bar{\Omega}^{m^*})$ satisfies OCs on $\bar{\Omega}^{m^*}$ if for every $0 < m \leq m^*$ and all $s_i = (\bar{\omega}_i^0, \dots, \bar{\omega}_i^{m^*})$,

$$s_i^m = \text{BR}_i(\text{marg}_{k, s_{-i}^{m-1}}(p(\cdot, \cdot | s_i))). \quad (3.45)$$

Let $\mathcal{O}^{m^*} \subseteq \Delta(K \times \bar{\Omega}^{m^*})$ denote the set of distributions on $K \times \bar{\Omega}^{m^*}$ that satisfy OCs. Let $X^* \subseteq K \times \bar{\Omega}^{m^*}$ be the set of $(k, (\bar{\omega}^0, \dots, \bar{\omega}^{m^*}))$ so that $\bar{\omega}^{m^*}$ is terminal. For $p \in \Delta(K \times \bar{\Omega}^{m^*})$, define the conditional terminal probability $\bar{p} \in \Delta(K \times \mathcal{A})$,

$$\bar{p}(k, B) = \sum_{(\bar{\omega}^0, \dots, \bar{\omega}^{m^*}) \in X^* : \bar{\beta}(\bar{\omega}^{m^*}) = B} p((k, (\bar{\omega}^0, \dots, \bar{\omega}^{m^*})) | X^*). \quad (3.46)$$

The conditional terminal probability \bar{p} satisfies limit-obedience if for every b in its support and every player i ,

$$B_i = \text{BR}_i(\bar{p}(\cdot, \cdot | B_i)). \quad (3.47)$$

Let $\mathcal{O}^\infty \subseteq \Delta(K \times \mathcal{A})$ denote the set of probabilities \bar{p} satisfying $B_i = \text{BR}_i(\bar{p}(\cdot, \cdot | B_i))$ for all i and B in the support of \bar{p} , i.e. satisfying limit-obedience.

Let $\mathcal{O} \subseteq \mathcal{O}^\infty$ denote the set of conditional terminal probabilities satisfying limit-obedience which are obtained from distributions in $p \in \mathcal{O}^{m^*}$, i.e.

$$\mathcal{O} = \{\bar{p} : p \in \mathcal{O}^{m^*}\} \cap \mathcal{O}^\infty. \quad (3.48)$$

The relative closure of this set is a convex polyhedron:

Lemma 16 (Linearity of \mathcal{O}). *The relative closure of the set \mathcal{O} is a convex polyhedron.*

Let $\mathcal{O}^* \subseteq \Delta(K \times \mathcal{A})$ denote the set of *Rationalizable Distribution*, i.e. the set of all outcome distributions that can arise under canonical priors

$$\mathcal{O}^* = \{\nu_P : P \in \mathcal{P}\}. \quad (3.49)$$

Theorem 7 (Rationalizable Distributions). *\mathcal{O} coincides with the set of all Rationalizable distributions, $\mathcal{O} = \mathcal{O}^*$. Its relative closure is a convex polyhedron.*

Our main result regarding rationalizable distributions follows from our sufficiency of SCAMP and the stationarity property of “simple” Markov processes. Every distribution in \mathcal{O}^{m^*} induces an average distribution on the branches of the automaton just like a prior in \mathcal{P} . So the averaging procedure used to establish the sufficiency of SCAMP in Theorem 6 can be used to map every distribution in \mathcal{O}^{m^*} into a SCAMP. Conversely, every SCAMP induces a distribution in \mathcal{O}^{m^*} through its marginal probability on the first m^* rounds.

3.5.4 Sufficiency of Additive Noise Information Structures

The automaton structure is derived from the recursive properties of ICR and proved useful in establishing finite dimensional characterizations of outcome distributions and information structures that obtain them. SCAMP bear strong similarities to the email game by [Rubinstein \(1989a\)](#). They can also be represented as information structures that are widely used in economic applications: Asynchronous information⁴ provision (as in [Abreu and Brunnermeier, 2003](#)), and a common state with additive noise (as in [Carlsson and Van Damme, 1993a](#)). Our main result thus establishes that those classes of information structures are “the same” and also sufficient to induce all outcomes in any finite game.

For a profile of finite sets of signals $(X_i)_i$, an *Additive Noise Information Structure* (ANIS) consists of 1) $\mu \in \Delta(K \times \prod_i X_i)$, where each X_i is a finite set of signals, and 2) a random time $\theta \in \mathbb{N}$ with individual delays $\tau_i \in \{0, \dots, n_i\}$, where $n_i \leq |\Omega|$. Given a signal profile $x \in \prod_i X_i$ and state of nature $k \in K$, a random time θ with a geometric distribution and a profile of bounded times $(\tau_i)_i$ are drawn independently of each other. Every player is privately informed of the vector $x_i \in X_i$ and her vector of times $\theta + \tau_i$. Every player is then privately informed of a pair $z_i = (x_i, y_i)$, where

$$y_i = \theta + \tau_i. \tag{3.50}$$

Theorem 8 (Sufficiency of ANIS). *For every finite game, there exists a profile of finite sets $(X_i)_i$ so that the set of Additive Noise Information Structures on $(X_i)_i$ induces all outcome distributions.*

Theorem 8 is an immediate consequence of Theorem 6: Every SCAMP can be decomposed into a random draw of an automaton branch, consisting of the set of all states visited by a path, and a random exit time for every cyclic state. Moreover, the support of every type’s beliefs when exiting a cycle is bounded. At every state of nature and every branch, there is a common exit time $\theta \in \mathbb{N}$. Each player is informed of some player-specific exit time $\theta + \tau_i$, where SCAMP implies that τ_i is bounded. With her exit time, each player i also receives a signal x_i about the state of nature and the automaton branch through her action labels $\bar{\beta}_i$.

ANIS are the combination of a one-dimensional additive noise information structure and a finite set of signals (as large as there are branches on the automaton). There are a few superficial differences to the original Global Game introduced by [Carlsson and Van Damme \(1993a\)](#). Firstly, in a global game the random draw θ is directly payoff-relevant, while in our set-up it is payoff-irrelevant but correlated with the payoff relevant state k . Since the set of payoff-relevant states is finite, we need to disentangle types from payoff-relevant states. Secondly, the noise term can be correlated with the payoff-relevant state in our setting while it is independent in a Global Game. [Oyama and Takahashi \(2011\)](#) provide an example where the dependence of the noise on payoff relevant states is necessary to induce all outcomes. Lemma 8 thus establishes, that up to an additional finite signal, a global game with state-dependent noise is sufficient for ICR.

SCAMP also has a straightforward dynamic interpretation: At every cycle along a branch, a profile of signals is drawn at some random time with geometric distribution. Players receive their private signal asynchronously within some bounded time window. With asynchronous

⁴See [Morris \(2014\)](#) for a more complete overview.

arrivals of private signals, players may believe other players received their signal earlier, which allows non-cyclic types (i.e. seeds), who receive their signal without delay, to infect other types. The asynchronicity can be by design or interpreted as a timing friction.

3.6 Information Disclosure in Priority Auctions

In this section, we use SCAMP to study information design in priority systems.

In many situations where customers are served sequentially, such as boarding and check-in at airlines, delivery services for packages, food, or groceries like Uber, Amazon, and Doordash, companies can charge their clients one or several priority tiers that determine the order in which they are served.

How useful it is for customers to pay for higher tiers depends both on the quality of the service, total demand, and how many other customers access priority tiers. Companies can then disseminate information and price tiers strategically in order to increase their revenue.

Customers can pay for higher priority, but if too many of them buy the higher priority class, then the utility of the service is diminished. In the specification below, the externality can vary with the quality. As we will show below, having a higher externality for higher quality allows the seller to extract approximately all the surplus via information design. In particular, we construct a SCAMP/ANIS to implement an information structure that extracts approximately all the surplus for the seller.

3.6.1 Model

Two buyers want to use a service of unknown quality⁵ $k \in K = \{1, \dots, \bar{k}\}$, where $\bar{k} > 0$. The quality is distributed according to a full support distribution $P_K \in \Delta(K)$. Each buyer $i \in \{1, 2\}$ chooses a priority $a_i \in A = \{0, \dots, n\}$ at a cost normalized to a_i , where $n > 2$. The service can be consumed by both buyers simultaneously, but each buyer prefers to be serviced first. Buyer i 's utility for the service with quality k when bidding a priority a_i and j bid priority a_j is given by:

$$u_i(k, a_i, a_j) = \begin{cases} k - a_i, & \text{if } a_j < a_i, \\ \alpha_k k - a_i, & \text{if } 0 < a_i = a_j, \\ -a_i, & \text{otherwise,} \end{cases} \quad (3.51)$$

where $0 < \alpha_k \leq 1$ represents an externality of consuming the service last (β_k) and simultaneously (α_k) when the quality is k and the utility of consuming last is set to zero.

Since rationalizability is set-valued, we must specify a selection rule to evaluate rationalizable outcome distributions. We take an adversarial approach by considering the worst possible selection from the perspective of the seller. The *revenue guarantee* for the seller at a canonical prior $P \in \mathcal{P}$ with outcome distribution ν_P is defined as

$$V(P) = \sum_{k, B} \min_{a_i \in B_i, a_j \in B_j} (a_i + a_j) \nu_P(k, b). \quad (3.52)$$

⁵We let K be an interval of \mathbb{N} to simplify the analysis.

Assume that for every k there is $a \in A$ such that $a \leq \alpha_k k$, that is, when the quality is commonly known to be k , there is a choice of priority, so that consuming simultaneously is better than consuming last. We then set \bar{a}_k the least priority with this property: $\bar{a}_k = \max\{a \in A : a \leq \alpha_k k\}$.

3.6.2 Complete Information Benchmark

As a benchmark, we compute the outcomes when the quality is publicly observed. Let $R_i(k) \subseteq A$ denote the set of player i 's ICR actions when quality k is publicly announced. The result below characterizes rationalizable bids when the externality of simultaneous consumption is either low (i.e. α_k is large), or high.

Lemma 17. $R_i(k) = \begin{cases} 0, \dots, \bar{a}_k & \text{if } \alpha_k > \frac{k-1}{k} \\ \bar{a}_k & \text{otherwise.} \end{cases}$

Proof. Assume first $\alpha_k > \frac{k-1}{k}$ and let $a_j \leq \bar{a}_k$. In this case, playing $a_i = a_j$ is strictly preferred to playing $a_i > a_j$ in state k . Playing $a_i = a_j$ is weakly preferred to $a_i < a_j$ if $\alpha_k k \geq a_j$. We deduce that every $a_i \leq \bar{a}_k$ is a best-reply to itself in state k , hence is rationalizable. Moreover, every $a > \bar{a}_k$ is dominated in state k by \bar{a}_k .

Now assume $\alpha_k \leq \frac{k-1}{k}$. In this case, bid a_1 dominates 0, and once 0 is eliminated a_2 dominates a_1 . We recursively eliminate dominated strategies $a_1, \dots, \bar{a}_{k-1}$ by successive outbidding. \bar{a}_k is a best-reply as long as $\alpha_k k \geq a_j$. \square

3.6.3 Upper Bound: Almost full extraction

We provide an upper bound on the payoff guarantee. The bound below corresponds to approximate full surplus extraction: it requires that both players play the highest action \bar{a}_k that would be rationalizable under common knowledge at each state k . Note that, by linearity of the payoffs and the required consistency with a common prior, signals that persuade the buyers of higher quality are always offset by signals that persuade the buyer of a lower quality. Hence the seller cannot, in expectation, induce bids that are higher than the average of $(\bar{a}_k)_k$. By definition, this extracts as much surplus from the buyer as possible. As we make the grid of possible bids finer, the surplus retained by the buyers, i.e. $\bar{a}_k - \alpha_k k$, shrinks to zero.

Lemma 18. For every $P \in \mathcal{P}$,

$$V(P) \leq \sum_{k \in K} 2\bar{a}_k P_K(k).$$

Proof. Let $P \in \mathcal{P}$. For every player i and type s_i , no bid above $\mathbb{E}_P(\alpha_k k | s_i)$ is rationalizable. To see this, note that for any choice $a_j > \mathbb{E}_P(\alpha_k k | s_i)$ of player $j \neq i$, no action greater than $a_j - 1$ is a best response for s_i . Since

$$\begin{aligned} V(P) &\leq \sum_{s \in S} (\mathbb{E}_P(\alpha_k k | s_i) + \mathbb{E}_P(\alpha_k k | s_j)) P(s) \\ &\leq 2\mathbb{E}_P(\alpha_k k) = \sum_{k \in K} 2\bar{a}_k P_K(k). \end{aligned} \tag{3.53}$$

□

We show that for certain parameters, there is a SCAMP/ANIS that actually attains the upper bound derived above: if $\alpha_{\bar{k}} \leq \frac{\bar{k}-1}{\bar{k}}$ and $\alpha_k > \frac{k-1}{k}$ for all $k < \bar{k}$, then the upper bound is attained by an ANIS information structure where every player i is privately informed of the following signal:

$$z_i = (a_k, \theta_k + \varepsilon_i), \quad (3.54)$$

where $a_k \in K$, $\theta_k \in \mathbb{N}$ and $\varepsilon_i \in \{0, 1\}$ is player and quality specific noise. The first component a_k corresponds to a private action recommendation similar to the canonical signals of a correlated equilibrium and the second component $\theta_k + \varepsilon_i$ is the additive noise component, similar to the canonical signals in global games.

3.6.4 SCAMP/ANIS

In Figures 3.7 and 3.8 we show an automaton that attains the upper bound provided above: all terminal nodes have singleton action labels given by \bar{a}_k , for every k and every player. Indeed, each automaton state in Figures 3.7 and 3.8 contains an action set for each player of the form $\{a, \dots, \bar{a}_k\}$, where $a \leq \bar{a}_k$. These action sets are identified in the figure by their minimal action. For example the node containing “ $\begin{smallmatrix} a_k \\ a_k + 1 \end{smallmatrix}$ ” represents an automaton state labeled with action sets $\{a_k, \dots, \bar{a}_k\}$ for player 1 and $\{a_k + 1, \dots, \bar{a}_k\}$ for player 2. In order to show that the payoff bound can be attained by some SCAMP/ANIS, it will be enough to find transition probabilities on the automaton so that OC holds. We will provide conditions for parameters $(\alpha_k)_k$ so that the automaton depicted below admits a SCAMP.

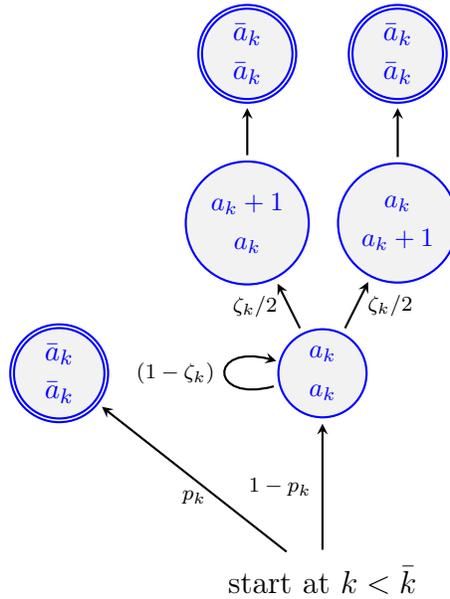


Figure 3.7: SCAMP Automaton for $k < \bar{k}$

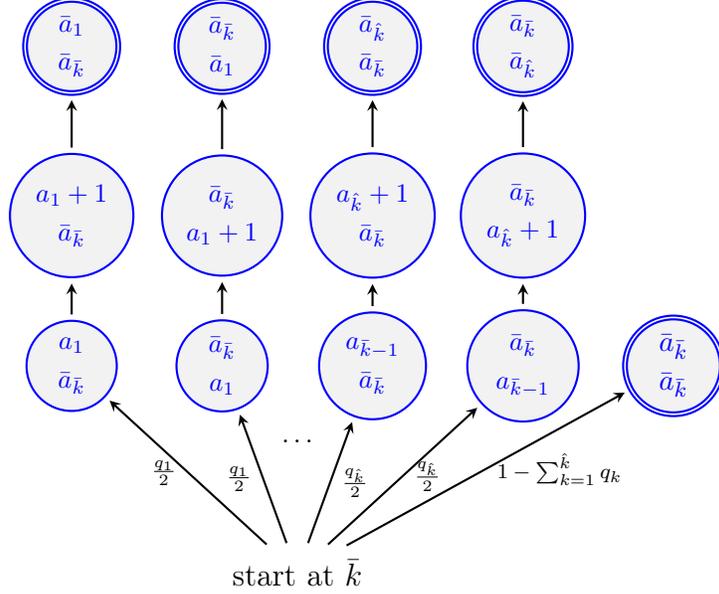


Figure 3.8: SCAMP Automaton for $k = \bar{k}$, where $\hat{k} := \bar{k} - 1$.

The SCAMP automaton in Figures 3.7 and 3.8 has three positive parameters for every quality $k < \bar{k}$: ζ_k, p_k and q_k , where $\sum_{k < \bar{k}} q_k \leq 1$ and $p_k \in [0, 1]$ for every $k < \bar{k}$. The $(q_k)_{k < \bar{k}}$ parametrizes the transition probabilities in Figure 3.8 when the state is \bar{k} . Positive q_k means that any player observing a sequence of the form $(a_k, a_k + 1, \dots, \bar{a}_k, \bar{a}_k, \dots)$ is uncertain if the quality is k or \bar{k} . By picking the right values of p_k, q_k , we can make OC hold for that sequence and allow ζ_k to be arbitrarily small - a feature required for the OC of cyclic types. The parameter ζ_k appears on the left part of the automaton and parametrizes the probability of exiting the cycle. After observing a sequence $(a_k, a_k, a_k + 1, \dots, \bar{a}_k, \bar{a}_k, \dots)$, a player knows that the state is k , but may not know if their opponent knows this.

Proposition 6 below states that, under certain conditions on the parameters $(\alpha_k)_k$, the automaton in Figures 3.7 and 3.8 can be populated with transition probabilities $(\zeta_k, q_k, p_k)_{k < \bar{k}}$ so that OC holds for all sequences. This establishes our (approximate) surplus extraction result:

Proposition 6 (Optimal Information Disclosure: ANIS). *Suppose $\alpha_{\bar{k}} \leq \frac{\bar{k}-1}{\bar{k}}$ and $\alpha_k > \frac{k-1}{k}$ for all $k < \bar{k}$. Then for every P_K the ANIS described in (3.61) attains the upper bound payoff guarantee in Lemma 18.*

The relevant conditions for a SCAMP/ANIS to extract (approximately) all of the surplus are twofold and can be explained using Lemma 17: first, the externality of simultaneous consumption at the highest quality \bar{k} , i.e. $1 - \alpha_{\bar{k}}$, must be large enough. This means that at the highest level of quality, players always prefer to outbid their opponent so as to avoid a tie, whenever outbidding can be profitable given a player's first order beliefs. Our SCAMP ensures that, conditional on the quality actually being highest, some type of a player will be informed of that fact. Second, when the quality is not \bar{k} , the externality, i.e. $1 - \alpha_k$ is low. Under complete information, Lemma 17 implies that all bids below \bar{a}_k are rationalizable because simultaneous consumption is less taxing. However, if players believe with enough probability

that their opponent only plays the highest rationalizable bid, they will match that bid. This information structure is thus a straightforward generalization of the Electronic Email game in [Rubinstein \(1989a\)](#): A few types (those who believe the quality to be highest) infect all other types and make them play the maximal bid. Unlike the single agent information design problem, as discussed in [Kamenica and Gentzkow \(2011\)](#), our information structure does not achieve the payoff bound only by introducing uncertainty about payoff-relevant parameters, but instead uses higher-order uncertainty to select the highest rationalizable action that would be played under common knowledge. Note that this example requires that there are more than two bids, i.e. $n > 2$. It can readily be shown that in the binary bid game, bidding zero is in fact a risk dominant equilibrium and so constructing an email game (as in [Rubinstein \(1989a\)](#)) would not allow the seller to extract all the surplus in that case.

ANIS admits a natural dynamic interpretation: Players receive a private signal about the quality a_k asynchronously within some time window around θ_k . The signal a_k contains a minimal bid recommendation. Signals indicating the highest quality arrive early.

Obedience Constraints We now state the OC and check for which set of parameters the OC have a solution. For every $k \in K$ and $m \in \mathbb{N}$, define the type

$$s_{k,m,i} = (0, \dots, \underbrace{a_k, \dots, a_k}_{m \text{ times}}, a_k + 1, \dots, \bar{a}_k, \bar{a}_k, \dots). \quad (3.55)$$

We start by deriving conditions that ensure OC of types that never cycle hold:

First-order Beliefs given $s_{k,1,i}$ for $k \neq \bar{k}$ are proportional to

$$P(k', s_{k,1,i}) = \begin{cases} \frac{\zeta_{k'}}{2} P(k'), & \text{if } k' = k \\ \frac{q_k}{2} P(k'), & \text{if } k' = \bar{k} \\ 0, & \text{otherwise.} \end{cases} \quad (3.56)$$

and for $k = \bar{k}$,

$$P(k', s_{k,1,i}) = \begin{cases} P(\bar{k}), & \text{if } k' = \bar{k} \\ 0, & \text{otherwise.} \end{cases} \quad (3.57)$$

To ease notation, let $Q_k = P(\bar{k} | s_{k,1,i})$. To check OC, we require outbidding to be a best-reply if and only if $a_j < \alpha_k k$. The expected payoff when bidding $\alpha_k k$ and tying (left hand side) should be equal to the expected payoff when bidding $\alpha_k k + 1$ and winning (right hand side):

$$\begin{aligned} Q_k \alpha_{\bar{k}} \bar{k} + (1 - Q_k) \alpha_k k - \alpha_k k &= Q_k \bar{k} + (1 - Q_k) k - \alpha_k k - 1 \\ Q_k (1 - \alpha_{\bar{k}}) \bar{k} + (1 - Q_k) (1 - \alpha_k) k &= 1 \\ Q_k &= \frac{1 - (1 - \alpha_k) k}{(1 - \alpha_{\bar{k}}) \bar{k} - (1 - \alpha_k) k} \end{aligned} \quad (3.58)$$

Since $1 > k(1 - \alpha_k)$ for all $k < \bar{k}$ and $1 < \bar{k}(1 - \alpha_{\bar{k}})$, we set q_k and ζ so that

$$Q_k = \frac{q_k / 2P(\bar{k})}{\zeta_k / 2P(k) + q_k / 2P(\bar{k})}. \quad (3.59)$$

That is, $q_k = 2^{\frac{1-(1-\alpha_k)k}{P(k)}}$ and $\zeta_k = 2^{\frac{(1-\alpha_{\bar{k}})\bar{k}-1}{P(k)}}$. We assume that ζ_k is arbitrarily small by letting $\alpha_{\bar{k}}$ be arbitrarily close to one.

We proceed with the OCs at $s_{k,m,i}$ for $m > 1$:

$$P(k, s_{k,m,i}, s_{k,n,-i}) = \begin{cases} P(k)(1 - \zeta_k)^m \frac{\zeta_k}{2}, & \text{if } n = m \\ P(k)(1 - \zeta_k)^{m-1} \frac{\zeta_k}{2}, & \text{if } n = m - 1 \\ 0, & \text{otherwise.} \end{cases} \quad (3.60)$$

Bidding $a_i = a_k$ when player j also bids a_k at round $n \leq m$ requires

$$\begin{aligned} P(k)(1 - \zeta_k)^m \frac{\zeta_k}{2} (\alpha_k k - a_i) &\geq ((1 - \zeta_k)^{m-1} (k - a_i - 1) + (1 - \zeta_k)^m \alpha_k k - a_i - 1) P(k) \frac{\zeta_k}{2} \\ \alpha_k k - a_i &\geq \frac{(1 - \zeta_k)^{m-1}}{(1 - \zeta_k)^m} (k - a_i - 1) + \alpha_k k - a_i - 1 \\ 1 &\geq \frac{1}{1 - \zeta_k} (k - a_i - 1) \end{aligned}$$

which for ζ_k small enough holds whenever $2 \geq k - a_k$. So bidding $a_i \geq a_k + 1$ at round $n \geq m + 1$ requires

$$\begin{aligned} (1 - \zeta_k)^m \frac{\zeta_k}{2} (k - a_i) + (1 - \zeta_k)^{m-1} \frac{\zeta_k}{2} (\alpha_k k - a_i) &> (1 - \zeta_k)^m \frac{\zeta_k}{2} (\alpha_k k - a_i - 1) \\ (1 - \zeta_k)(k - a_i) + (\alpha_k k - a_i) &> (1 - \zeta_k)(\alpha_k k - a_i) - (1 - \zeta_k) \\ (1 - \zeta_k)k + \zeta_k \alpha_k k &> a_i - 1 + \zeta_k \end{aligned}$$

Which for ζ_k small enough holds when $k - a_k > -1$.

Consider the ANIS on $K \times Z_1 \times Z_2$, where Z_i is the set of possible signals that can arise from (3.54), so that for each $k \in K$,

$$\begin{aligned} \Pr(\theta_k = n|k) &= \begin{cases} \zeta_k^* (1 - \zeta_k^*)^n, & \text{if } k \neq \bar{k} \\ \mathbf{1}_{n=0}, & \text{otherwise.} \end{cases} \\ \Pr(a_{\bar{k}}|k) &= \begin{cases} (1 - q_k^*), & \text{if } \tilde{k} = k \\ q_k^*, & \text{if } \tilde{k} = \bar{k} \\ 0, & \text{otherwise.} \end{cases} \\ \Pr(\varepsilon_{i,k} = 0|k) &= \frac{1 - \zeta_k^*}{2 - \zeta_k^*} \\ q_k^* &= \frac{1 - (1 - \alpha_k)k}{P(\bar{k})}, \quad \zeta_k^* = \frac{(1 - \alpha_{\bar{k}})\bar{k} - 1}{P(k)} \end{aligned} \quad (3.61)$$

3.7 Discussion

3.7.1 Structure and Complexity of SCAMP

Consider the strategic automaton used to establish Theorem 6. In the course of the construction, we show that its size is bounded. OCs of SCAMP may nevertheless be

complex due to cycles allowing paths to be “out of sync”. However, the property of “simple” in SCAMP implies that OCs only contain a bounded number of terms. A process P has *bounded obedience constraints* if there is $L \in \mathbb{N}$ so that for every player i every type s_i in the support of P and every round m ,

$$|\{(k, v) \in K \times \overline{\Omega}^{\mathbb{N}} : P(k, v|s_i) > 0\}| \leq L. \quad (3.62)$$

The following result is then an immediate consequence of “simple”:

Corollary 5 (Bounded OC of SCAMP). *Every SCAMP has bounded obedience constraints.*

In the example of Section 3.2 we moved from an automaton with multiple cycles (Figure 3.2) to an automaton with only one cycle (Figure 3.3), where we only kept the lowest cycle. Any SCAMP on this automaton has the property that the beliefs of every type are finitely supported: A player who transitions from ab to a at round m knows that the cycle was left within two rounds after or before round m . However, the cycling probability may depend on the branch and the state of nature.

3.7.2 Solution Concepts

The subsequent discussion motivates the requirement for ICR in information design as opposed to that of correlated equilibria.

In the framework of correlated equilibria, an information structure implements a distribution of outcomes if there is at least one Nash equilibrium in the corresponding Bayesian game that results in the desired distribution. It is acknowledged, however, that multiple equilibria might exist within the Bayesian game, potentially leading to alternative outcome distributions. Implicit in this approach is the assumption that the information designer is not only responsible for the dissemination of information but also possesses the capability to direct players toward coordination on a specific Nash equilibrium.

Rationalizability does not assume coordination on an equilibrium but only common certainty of Bayesian rationality, hence iterative deletion of (strictly) dominated strategies. It is a set solution concept, and in general, more than one outcome can survive the iterative deletion of dominated strategies. Rationalizable distributions include the set of distributions that are implementable in dominant strategies. This concept is thus both weaker than implementation in dominant strategies, and stronger than Nash implementation.

Our results are built on the recursive structure of ICR. For supermodular games, ICR and Bayes-Nash equilibrium (BNE) make the same predictions regarding extremal outcomes. More generally, Liu (2015a) shows that a subjective version of Belief invariant Bayes correlated equilibrium is equivalent to ICR. We thus expect analogous results and structures to arise under Belief Invariant Bayes Correlated Equilibrium and leave that extension and the extension to BNE in general games for future work.

3.7.3 Relation to Literature

SCAMP can be viewed as a generalization of the e-mail game in Rubinstein (1989a), where a type corresponds to the number of emails sent before a message was lost. Similar models of

information were used in [Oyama and Takahashi \(2020\)](#), [Morris et al. \(2020\)](#) and [Halac et al. \(2021\)](#) for instance. These models have the advantage of being economical in parameters and are examples of SCAMP for binary action games with different interpretations. In the model of global games ([Carlsson and Van Damme, 1993a](#); [Morris and Shin, 2003b](#)), a type to a player corresponds to a (usually unidimensional) noisy signal about the underlying state of nature. These models offer good interpretability and can easily yield comparative statics such as the relative precision of players' signals. However different in nature, one being a discrete model, and the other one continuous, the information structure in the email game and in global games appear to be deeply connected. In fact, both models allow to encompass rich systems of higher-order beliefs and capture contagion phenomena through which dominant actions unravel.

At a more fundamental level, both models exhibit the same relationship between types. In the email game, when a player stops receiving message confirmations, either the message sent to the other player was lost, or the subsequent confirmation from that player was lost, and this relationship is essentially independent of the number of messages sent. In global games, when conditioning on the observation of a signal, typical signals received by other players are either slightly higher, or slightly lower, and this relationship is here too independent of the value of the signal observed. Both models satisfy a form of translation invariance of beliefs, which is captured by our Markov property. With SCAMP we isolate this aspect of information design and show that it is all you need to generate all rationalizable outcomes in any finite game.

Chapter 4

A Strategic Topology on Information Structures

4.1 Introduction

4.1.1 Motivation and Results

Players make choices in a game where payoffs depend on some unknown state of the world. Optimal strategic behavior will then depend on players' beliefs about the states, their beliefs about others' beliefs, and so on. A (common prior) information structure of the players can be represented as a probability distribution over the players' beliefs and higher-order beliefs about the state. We are interested in how the information structure impacts equilibrium outcomes in all finite games. We define a notion of closeness between information structures, by saying that two information structures are close if their equilibrium outcomes are close in all finite games. To make such a definition precise, we must define a topology on a fixed set of information structures. We will take as our fixed set of information structures the set of common priors on the universal type space. In this set, players' private information (i.e. their type) is identified by a hierarchy of beliefs. To describe our topology on information structures, we will take a topology on hierarchies of beliefs as given. Hierarchies of beliefs are endowed with a natural topology¹, the product topology. In this topology, two hierarchies are close if their beliefs are close up to some finite order. However, it is known that equilibrium outcomes are sensitive to the infinite tails of higher-order beliefs ([Rubinstein \(1989b\)](#) and [Carlsson and Van Damme \(1993b\)](#)). Our main result is a characterization of the coarsest topology on information structures generating continuity of equilibrium outcomes.

Following [Monderer and Samet \(1989\)](#), an event is said to be *common p -belief* if everyone believes it with probability at least p , everybody believes with probability at least p that everyone believes it with probability at least p , and so on. An event is said to be *approximate common knowledge* if it is common p -belief with p close to 1. Now two information structures are close in the almost common knowledge topology if, with high ex ante probability, there is approximate common knowledge that their higher-order beliefs are close (in the product

¹We say that the product topology on hierarchies is natural because it is under this topology that the set of hierarchies satisfy their universal property (see [Mertens et al. \(2015\)](#) for details).

topology). Our main result (Theorem 1 in Section 1) establishes that the almost common knowledge topology is the coarsest topology generating continuity of equilibrium outcomes.

Our definition of information structures excludes any (payoff-irrelevant) correlating devices that players might have access to. In the language of [Mertens and Zamir \(1985\)](#), we restrict attention to *non-redundant* information structures. However, our main result allows the correlating device to appear within the equilibrium solution concept: we study *belief-invariant Bayes correlated equilibria* where players' action choices can be correlated but only when the correlation does not alter players' beliefs and higher-order beliefs about the state. In this set-up, we show that finite information structures are dense in the strategic topology. While we think our notions of information structures and equilibrium are the most natural for our main exercise, we also spell out (in Section 4.4) how our results change if we allow more general information structures with redundancies and Bayes Nash Equilibrium (which rules out correlation in the solution concept). We show that the set of Bayes Nash equilibria under any general information structure (with redundancies) can be approximated by the set of belief-invariant Bayes correlated equilibria of a finite, non-redundant information structure. One implication is that it is without loss to focus on finite, non-redundant information structures in solving information design problems.

4.1.2 Alternative Topologies and Related Literature

[Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#), building on [Monderer and Samet \(1989\)](#), identified topologies on information structures that generated continuity of equilibrium outcomes. And the topologies identified both have the same flavor as the almost common knowledge topology. The key difference is that both papers fix (different) sets of information structures with the restriction that each profile of types gives rise to a distinct state, and is not associated with higher-order beliefs. This makes the topology hard to interpret. It also makes both directions of the proof easier than in our problem. We describe, following our main result (Theorem 1 in Section 1), which steps of our proof relate to this early work and which are novel to this paper. We discuss the joint measurability condition implicit in [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#) in Section 4.6.3. Other differences are that the earlier work focuses on Bayes Nash equilibrium and restricts attention to countable information structures to ensure existence; we allow uncountable information structures (sets of the universal type space) and we ensure equilibrium existence by incorporating redundancies in our solution concept. But we also show how our results apply to Bayes Nash equilibrium and general information structures in Section 4.4. The earlier work measures closeness of equilibrium outcomes by players' expected payoffs in equilibrium; we note in Section 4.6.2 that our results would be unchanged if we used this approach.

[Dekel et al. \(2006\)](#) defined an interim strategic topology on hierarchies of beliefs under the solution concept of interim correlated rationalizability (ICR). Two belief hierarchies were said to be close in the interim strategic topology if, in any game, an action that was ICR at one hierarchy was approximately ICR at the other hierarchy. It is well-known (e.g., from the work of [Rubinstein \(1989c\)](#), [Carlsson and Van Damme \(1993a\)](#) and [Weinstein and Yildiz \(2007\)](#)) that closeness in the product topology is not sufficient for close strategic behavior. [Chen et al. \(2017\)](#) characterized the interim strategic topology of [Dekel et al. \(2006\)](#) and showed that it requires closeness of beliefs about some tail events (“frames”). However, the connection to

approximate common knowledge has been unclear. We show that two information structures are close in our almost common knowledge topology if and only if there is a high ex ante probability that hierarchies are close in the interim strategic topology.

4.2 Model

This section will introduce the model used in our main analysis, with four main ingredients.

1. We will hold fixed a finite set of players, a finite set of (payoff-relevant) *states* and a probability distribution over the states.
2. We define a *base game* to consist of players' possible actions and their payoff functions; i.e., how each player's payoff depends on the action profile chosen and the state.
3. An *information structure* will consist of a probability distribution over the state and players' beliefs and higher-order beliefs about the state, with the appropriate marginal over states.
4. The equilibrium solution concept will be the *belief-invariant Bayes correlated equilibrium* (BIBCE), see Definition 8 in [Bergemann and Morris \(2016b\)](#). This is a joint distribution over the information structure and the players' actions such that players' actions are best responses and measurable with respect to players' beliefs and higher-order beliefs about the state, but allowing payoff-irrelevant correlation of actions.

Our model of an information structure is restrictive in that it rules out players observing multiple signals giving rise to the same beliefs and higher-order beliefs. Equivalently, it rules out what [Mertens and Zamir \(1985\)](#) labelled *redundant* types, i.e., players observing payoff-irrelevant signals through which they can correlate their behavior. On the other hand, our equilibrium solution concept allows players to observe payoff-irrelevant correlating devices in choosing actions. Thus we have made a modelling choice to put correlation devices in the solution concept rather than the information structure.

In order to relate our work to the literature and to applications, we will later discuss (in Section 4.4) how our results easily translate to a setting with general information structures (allowing *redundancies*) and the more relaxed solution concept of Bayes Nash equilibrium, i.e., if we move correlation possibilities from the solution concept to the information structure.

4.2.1 Setting and Base Game

There is a finite set of players I and a finite set of (payoff) states Θ . Throughout the paper, we will fix a prior $\mu \in \Delta(\Theta)$. Without loss, we will maintain the assumption that μ has full support.

A *base game* then describes players' actions and payoffs: thus a base game is a tuple $\mathcal{G} = ((A_i)_{i \in I}, (u_i)_{i \in I})$, where A_i is a finite set of actions for player i and $u_i : A_i \times A_{-i} \times \Theta \rightarrow [-M, M]$ is a payoff function for player i , where $A_{-i} := \prod_{j \neq i} A_j$ and $M > 0$ is an exogenous payoff bound.

4.2.2 Information Structures

We follow [Harsanyi \(1967-68\)](#) and [Mertens and Zamir \(1985\)](#) in identifying signals or types with players' beliefs and higher-order beliefs, or hierarchies of beliefs. We first formally define the set of hierarchies of beliefs.

Definition 1. (*Hierarchies of Beliefs*) The profile of players' hierarchies of beliefs, $(\mathcal{T}_i)_{i \in I}$ is defined recursively as follows: For every i , let $\mathcal{T}_i^0 := \{*\}$ be a singleton and let $\mathcal{T}_i^1 := \Delta(\Theta)$. Given profiles $(\mathcal{T}_i^{m-1})_{i \in I}$ for $m > 1$, define for every i ,

$$\mathcal{T}_i^m := \left\{ ((\tau_i^1, \dots, \tau_i^{m-1}), \tau_i^m) \in \mathcal{T}_i^{m-1} \times \Delta(\mathcal{T}_{-i}^{m-1} \times \Theta) : \text{marg}_{\mathcal{T}_{-i}^{m-2} \times \Theta}(\tau_i^m) = \tau_i^{m-1} \right\}.$$

Let \mathcal{T}_i denote the set of sequences $\tau_i = (\tau_i^m)_m$ so that for every $\bar{m} \in \mathbb{N}$, the truncated sequence $(\tau_i^m)_{m \leq \bar{m}}$ belongs to $\mathcal{T}_i^{\bar{m}}$.

We will follow [Mertens and Zamir \(1985\)](#) in imposing the product topology on hierarchies \mathcal{T}_i . Let d_{Π} be a metric on $\Omega := \mathcal{T} \times \Theta$ inducing the product topology on $\mathcal{T} := \prod_{i \in I} \mathcal{T}_i$ and the discrete topology on Θ .² However, we later discuss (Section 4.6.1) what properties of the product topology on hierarchies is important for our arguments. [Mertens and Zamir \(1985\)](#) show that for every $(\tau = (\tau_i)_{i \in I}, \theta) \in \Omega$ and every $i \in I$ there is a unique belief $\tau_i^* \in \Delta(\mathcal{T}_{-i} \times \Theta)$ so that for all $m \in \mathbb{N}$, $\tau_i^m = \text{marg}_{\mathcal{T}_{-i}^{m-1} \times \Theta}(\tau_i^*)$, where $\tau \mapsto \tau^* = (\tau_i^*)_{i \in I}$ is a homeomorphism. Let \mathcal{B} denote the Borel sigma-algebra on Ω . We will refer to Ω as the “universal state space”, with typical element $\omega = (\tau, \theta)$, where $\tau \in \mathcal{T}$. Now an information structure is just a probability distribution on the universal state space that respects the prior on states and the labelling of hierarchies.

Definition 2. (*Information Structure*). An information structure P is a Borel probability measure on Ω that satisfies two conditions: (i) [consistency] the marginal of P on Θ is μ ; (ii) [labelling] for every player i , there is a version of the conditional probability $P_i: \mathcal{T}_i \rightarrow \Delta(\mathcal{T}_{-i} \times \Theta)$ of P so that

$$\tau_i^* = P_i(\tau_i), \quad P\text{-a.s.}$$

We will write $\mathcal{P} \subseteq \Delta(\Omega)$ for the set of (common prior) information structures. We will sometimes describe these as *non-redundant* information structures because we do not allow multiple types or signals of a player giving rise to the same hierarchies of beliefs. This is to contrast them with the more general *redundant* information structures discussed in Section 4.4.

An information structure is *finite* if it has finite support. A hierarchy $\tau_i \in \mathcal{T}_i$ is *finite* if it is in the support of a finite information structure. The set of finite types in \mathcal{T}_i is denoted by \mathcal{T}_i^* and the set of finite states $\Omega^* \subseteq \Omega$ is given by

$$\Omega^* := \{(\tau, \theta) \in \Omega : \forall i, \tau_i \in \mathcal{T}_i^*\}.$$

²The product topology on \mathcal{T} is the coarsest topology so that projections $\text{proj}_{\mathcal{T}^m}: \mathcal{T} \rightarrow \mathcal{T}^m$ are continuous, where for every m , \mathcal{T}^m is endowed with the weak topology. [Dekel et al. \(2006\)](#) provide a metric that induces the product topology on \mathcal{T} , which for any discount factor $\eta \in (0, 1)$ can be described by $d_{\Pi}((\tau, \theta), (\hat{\tau}, \hat{\theta})) = \mathbf{1}_{\theta \neq \hat{\theta}} + \sum_{n=1}^{\infty} \eta^n \max_i d_i^n(\tau_i^n, \hat{\tau}_i^n)$, where d_i^n is a product metric inducing the weak topology on \mathcal{T}_i^n .

4.2.3 Solution Concept: Belief-Invariant Bayes Correlated Equilibrium

Together, a base game \mathcal{G} and an information structure P define a game of incomplete information (\mathcal{G}, P) . Now we define our main equilibrium solution concept. We will be allowing players' action choices to be correlated. So players' action choices will be described by a *decision rule*, mapping states and hierarchies to distributions over action profiles; thus for any incomplete information game (\mathcal{G}, P) , a decision rule is a measurable map $\sigma : \mathcal{T} \times \Theta \rightarrow \Delta(A)$, where $A := \prod_{i \in I} A_i$ and $\Delta(A)$ is endowed with the Euclidean topology.

An information structure P and decision rule σ jointly induce a measure $\sigma \circ P \in \Delta(A \times \mathcal{T} \times \Theta)$ in the natural way. We will be interested in *outcomes* specifying a joint distribution over actions and states $\nu \in \Delta(A \times \Theta)$. Decision rule σ *induces* outcome ν_σ if ν_σ is the marginal of $\sigma \circ P$ on $A \times \Theta$. For every player i , a decision rule σ and hierarchy $\tau_i \in \mathcal{T}_i$ induce a belief $\sigma \circ \tau_i \in \Delta(A \times \mathcal{T}_{-i} \times \Theta)$, which for every measurable set $E \subseteq \mathcal{T}_{-i} \times \Theta$ and action profile $a \in A$ satisfies $\sigma \circ \tau_i(\{a\} \times E) = \int_E \sigma(a | \tau_{-i}, \tau_i, \theta) d\sigma \circ \tau_i^*$.

Our notion of equilibrium will be (one version of) incomplete information correlated equilibrium. Two properties will be key.

Definition 3. A decision rule σ is ε -obedient if for every player i , action a_i and deviation a'_i ,

$$\int \sum_{a_{-i} \in A_{-i}} (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) d\sigma \circ \tau_i(a_i, a_{-i}, \tau_{-i}, \theta) > -\varepsilon, \text{ a.s.}$$

Thus if we interpret the decision rule $\sigma : \mathcal{T} \times \Theta \rightarrow \Delta(A)$ as a recommendation rule for a mediator in the game, ε -obedience requires that players will not lose more than ε by following the recommendation.

Importantly, this is an interim version of ε -obedience: we are requiring that the decision rule is almost always approximately optimal.³

Definition 4. A decision rule σ is belief-invariant if, for every $a_i \in A_i$, the marginal probability $\sigma(a_i \times A_{-i} | (\tau_i, \tau_{-i}), \theta) = \sigma_i(a_i | \tau_i)$ does not depend on (τ_{-i}, θ) .⁴

Definition 5. A decision rule σ is an ε -belief-invariant Bayes correlated equilibrium (ε -BIBCE) of (\mathcal{G}, P) if it satisfies ε -obedience and belief invariance.

We will say that a decision rule is a BIBCE if it is a 0-BIBCE.

³An ex-ante notion of ε -obedience would require interim ε -obedience only with probability $1 - \varepsilon$, and thus allows players to choose actions that are not ε best responses with probability $1 - \varepsilon$. The ex-ante notion of ε -obedience is much more permissive: Engl (1995) has established the lower hemi-continuity of ex ante ε -Bayes Nash equilibrium in games of incomplete information with respect to weak convergence of priors on a fixed type space.

⁴Lehrer et al. (2010) and Forges (2006) introduced the notion of belief-invariance to the study of incomplete information correlated equilibrium: Forges (1993b) (implicitly) and Lehrer et al. (2010) and Forges (2006) (explicitly) imposed the restriction in their definitions of the belief-invariant Bayesian solution. This latter solution imposes (implicitly or explicitly) joint feasibility, so that the decision rule depends on the profile of players' signals. Liu (2015b) showed that a subjective version of BIBCE is equivalent to interim correlated rationalizability. A belief-invariant decision rule can be interpreted as a payoff-irrelevant correlating device (where player i observes private signal a_i). Now BIBCE is equivalent to Bayes Nash equilibria if players are allowed to observe a correlating device.

4.3 Main Result

In this section, we first define an “approximate common knowledge” topology on information structures and then show it is the coarsest topology to generate continuity of BIBCE outcomes (our main result).

4.3.1 The Approximate Common Knowledge Topology

In words, we will say that two information structures are close if each assigns high ex ante probability to the event that there is approximate common knowledge that their interim beliefs are close. An event is “approximate common knowledge”, if for some p close to 1, everyone believes the event with probability at least p , everyone believes with probability at least p that everyone believes the event with probability at least p , and so on.... (Monderer and Samet (1989)). The subtlety in formalizing this definition is how we define the event that interim beliefs are close.

We first define ε -neighborhoods, the set of universal states ε -close to a given universal state $\omega \in \Omega$.

Definition 6. For every $\varepsilon > 0$ and $\omega \in \Omega$, the ε -neighborhood of ω is given by

$$\mathcal{N}_\varepsilon(\omega) := \{\omega' \in \Omega : d_\Pi(\omega, \omega') < \varepsilon\}.$$

For any $\varepsilon > 0$, we define the ε -support of $P \in \mathcal{P}$, to be the ε -neighborhoods of universal states whose ε -neighborhoods are assigned positive probability by P :

$$\text{supp}_\varepsilon(P) := \bigcup_{\omega \in \Omega: P(\mathcal{N}_\varepsilon(\omega)) > 0} \mathcal{N}_\varepsilon(\omega).$$

Thus $\text{supp}_\varepsilon(P)$ is an ε -expansion of the support of P . For any two priors P, P' , we define the intersection of their ε -supports:

$$\hat{T}_\varepsilon(P, P') := \text{supp}_\varepsilon(P) \cap \text{supp}_\varepsilon(P').$$

This is the set of universal states where interim beliefs are close to interim beliefs in the support of both priors. Following Monderer and Samet (1989), for every probability $p \in [0, 1]$ and event $E \in \mathcal{B}$, we define $B^p(E)$ to be the set of universal states where all players assign probability at least p to the event E . Thus

$$B^p(E) := \{(\tau, \theta) \in E : \forall i, \tau_i^*(E_{-i}) \geq p\}$$

where $E_{-i} := \text{proj}_{\mathcal{T}_{-i} \times \Theta}(E)$. For every $m \in \mathbb{N}$, let $[B^p]^m(E) := B^p \circ \dots \circ B^p(E)$ denote the m -fold application of B^p ; $[B^p]^m(E) := B^p \circ \dots \circ B^p(E)$ is the set of universal states where all players assign probability at least p to all players assigning at least probability p (m times) to event E being true. Now an event is said to be *common p -belief* if this is true for all m .

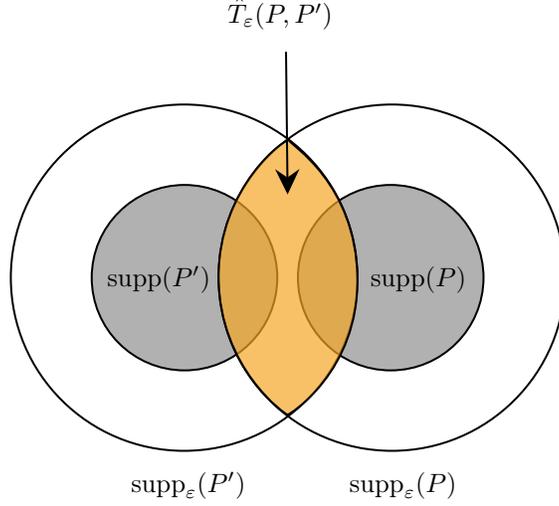


Figure 4.1: Sets $\hat{T}_\varepsilon(P, P')$ and ε -supports of P and P' .

Definition 7. (*Common p -Belief*) For any event $E \in \mathcal{B}$, the event that E is common p -belief, $C^p(E)$, is defined as

$$C^p(E) := \bigcap_{m \in \mathbb{N}} [B^p]^m(E).$$

An event is said to be approximate common knowledge if it is common p -belief for p close to 1. We can now define the approximate common knowledge (ACK) distance⁵ between any two priors, $P, P' \in \mathcal{P}$.

Definition 8. (*Approximate Common Knowledge Distance*) For every $P, P' \in \mathcal{P}$, let

$$d^{ACK}(P, P') := \inf \left\{ \varepsilon \geq 0 : \begin{array}{l} P \left(C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) \geq 1 - \varepsilon \\ P' \left(C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) \geq 1 - \varepsilon \end{array} \right\}.$$

Thus two priors are close if, under both priors, there is high probability that there is approximate common knowledge that their interim beliefs are close.

Definition 9. (*Approximate Common Knowledge Topology*) The approximate common knowledge (ACK) topology is the topology generated by open sets $\{P' \in \mathcal{P} : d^{ACK}(P, P') < \varepsilon\}_{P \in \mathcal{P}}$.

While d^{ACK} fails the triangle inequality, the topology is metrizable.

Proposition 1. *The ACK topology is metrizable.*

The proof (in the Appendix) constructs a metric that generates the same topology. Because the topology is metrizable, it is characterized by the convergent sequences.

Definition 10. A sequence $(P^k)_k$ in \mathcal{P} ACK-converges to P if and only if $d^{ACK}(P^k, P) \rightarrow 0$.

⁵A distance is a map $d : \mathcal{P} \times \mathcal{P} \rightarrow [0, \infty)$ which is zero on the diagonal.

It is useful to consider what the topology looks like in some special cases. If there is only one payoff relevant state (i.e., $|\Theta| = 1$), then there is a unique universal state, where there is common knowledge of the state. If there is only one player, then an information structure is given by a distribution over first order beliefs about Θ , and the topology reduces to the weak topology on $\Delta(\Delta(\Theta))$. Similarly, if there are many players but types are perfectly correlated. If types are independent, then an information structure is given by a distribution in $\Delta(\Delta(\Theta)^I)$, and the topology again reduces to the weak topology. If we restrict to information structures with a finite support of a fixed size, then an information structure is a probability distribution on a finite set and the topology reduces to the weak topology. So in order for the ACK topology to be interesting, there must be at least two states, at least two players, an unbounded number of types and neither independence or perfect correlation.

4.3.2 Continuity of Equilibrium Outcomes

We now define continuity of equilibrium outcomes (for BIBCE). For every $\varepsilon > 0$ and (\mathcal{G}, P) , let $\mathcal{B}^\varepsilon(\mathcal{G}, P)$ denote the collection of all ε -BIBCE under (\mathcal{G}, P) . We write $\mathcal{O}_\varepsilon(\mathcal{G}, P)$ for the set of outcomes that are within ε of outcomes induced by an ε -BIBCE of (\mathcal{G}, P) , so

$$\mathcal{O}_\varepsilon(\mathcal{G}, P) := \{v \in \Delta(A \times \Theta) : \exists \sigma \in \mathcal{B}^\varepsilon(\mathcal{G}, P) \text{ s.t. } \|v - v_\sigma\|_2 < \varepsilon\},$$

where $\|\cdot\|_2$ is the Euclidean norm on outcomes. Now we will say that two information structures are strategically close for base game \mathcal{G} if the sets of BIBCE outcomes are close.

Definition 11. (*Strategic Distance*) For every $P, P' \in \mathcal{P}$ and base game \mathcal{G} , let

$$d^*(P, P'|\mathcal{G}) := \inf \left\{ \varepsilon \geq 0 : \begin{array}{l} \{\nu_\sigma : \sigma \in \mathcal{B}^0(\mathcal{G}, P)\} \subseteq \mathcal{O}_\varepsilon(\mathcal{G}, P') \\ \{\nu_\sigma : \sigma \in \mathcal{B}^0(\mathcal{G}, P')\} \subseteq \mathcal{O}_\varepsilon(\mathcal{G}, P) \end{array} \right\}.$$

A topology on \mathcal{P} generates continuity of equilibrium outcomes if for every $\varepsilon > 0$, every base game \mathcal{G} and every $P \in \mathcal{P}$, the set $\{P' \in \mathcal{P} : d^*(P, P'|\mathcal{G}) < \varepsilon\}$ is open.

4.3.3 Main Result

The following is our main result.

Theorem 1. *The approximate common knowledge topology is the coarsest topology on \mathcal{P} that generates continuity of equilibrium outcomes.*

We will first sketch the main ideas in the proof, distinguishing between steps that are inherited from earlier work, and how, and which steps are and must be novel. To establish that approximate common knowledge convergence implies strategic convergence (i.e. that the ACK topology generates continuity of strategic outcomes), one must show that if two information structures are close enough in the approximate common knowledge topology, then in any game, the set of ε -BIBCE outcomes will be close. This builds on the arguments in [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#) (which in turn build on an argument in [Monderer and Samet \(1989\)](#)). These two papers fix an equilibrium strategy profile under one information structure and show that there is an approximate equilibrium

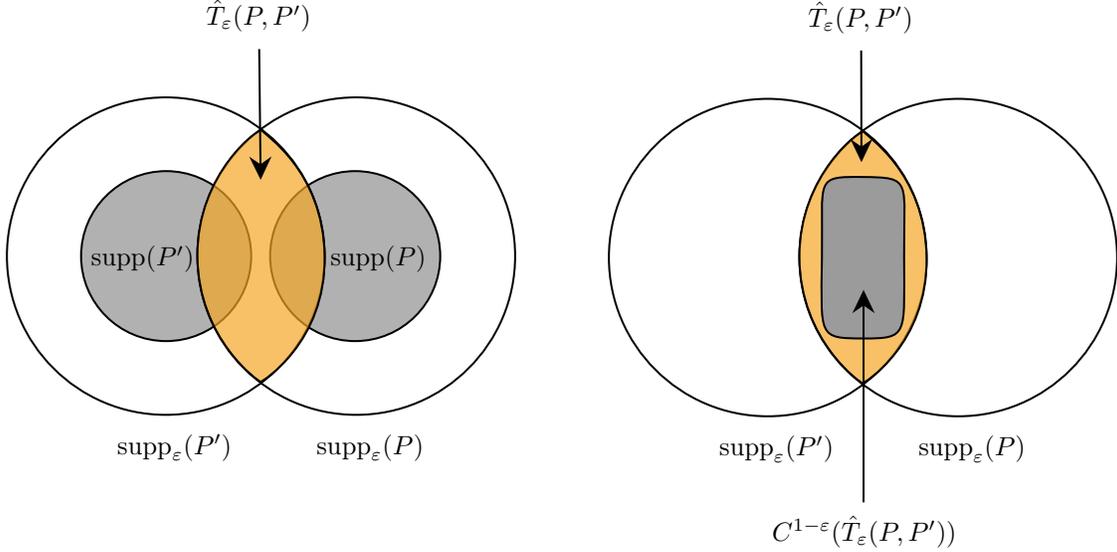


Figure 4.2: Sets $\hat{T}_\varepsilon(P, P')$, $C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$ and ε -supports of P and P' .

in any nearby information structure where the strategy profile is unchanged on the event where there is approximate common knowledge that interim beliefs are close, but can vary arbitrarily elsewhere. This proof strategy relies on a well-defined notion of holding the strategy profile fixed on the approximate common knowledge event. In our context, there is no well-defined notion of holding the strategy profile fixed on the approximate common knowledge event. In particular, if the two information structures are minimal (they do not have common knowledge subsets), the supports of distinct information structures will be disjoint, as illustrated in Figure 4.2. So instead we will continuously extend the equilibrium decision rule σ under one information structure P to its ε -support $\text{supp}_\varepsilon(P)$ and thus to the event:

$$\hat{T}_\varepsilon(P, P') := \text{supp}_\varepsilon(P) \cap \text{supp}_\varepsilon(P'),$$

and thus the approximate common knowledge event:

$$C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P')) \subseteq \hat{T}_\varepsilon(P, P').$$

This is one place where we are exploiting properties of the product topology on types to show that the extended decision rule satisfies approximate obedience on $C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$ and hence is an approximate equilibrium under P' . This argument goes through with any refinement of the product topology as we will discuss later.

To establish that strategic convergence implies approximate common knowledge convergence, it is enough to show that, for any two information structures that are not close in the approximate common knowledge topology, one can construct a game where an equilibrium outcome under one information structure is not close to any approximate equilibrium under the other information structure. [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#) do this by showing that if two information structures P and P' are not close, there is an “infecting” event D^0 under one of the information structures, say P , such that there is no

approximate common knowledge event on the complement of $C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$. One can then construct a binary coordination game in the spirit of the email game of [Rubinstein \(1989b\)](#) where there is a unique equilibrium under P because there is a dominant strategy on the “infecting” event, but there are multiple equilibria under P' . This strategy is not available to us because we cannot assume that there is a dominant strategy on the “infecting” event because payoffs must be measurable with respect to the payoff states Θ .

Instead, we show that if two information structures are not close, there exists an “infecting” event D^m that is measurable with respect to m -order beliefs. For m large enough, this event has the property that some types in the support of P that are not in $\hat{T}_\varepsilon(P, P')$ are closer to any element in D^m than to any element in $\hat{T}_\varepsilon(P, P')$. Hence, for m large enough, some types can be excluded from $C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$ based on their beliefs on a m -order-measurable event. From event $D_\varepsilon^0 = D^m$ we obtain a cover $(D_\varepsilon^n)_{n \in \mathbb{N}}$ of the complement of $C^{1-\varepsilon}(\hat{T}_\varepsilon(P, P'))$ recursively, where

$$D_\varepsilon^n = \text{supp}(P) \setminus B^{1-\varepsilon}(\text{supp}(P) \setminus D_\varepsilon^{n-1}),$$

for all $n \geq 1$. Now we can construct a first game where players are incentivized to announce their approximate m -th order beliefs on a finite grid, (for m large enough chosen as a function of ε) and build an email game style binary action coordination game on top of that. For the first component of the game, we can construct a game with an iterated scoring rule with the property that it is ε -rationalizable for players to truthfully report finite-order beliefs, from a finite grid which are closest in d_Π . This scoring rule is also used in [Dekel et al. \(2006\)](#) and [Gossner and Mertens \(2020\)](#). This game alone cannot be used to induce different outcomes in P and P' since on the set $\hat{T}_\varepsilon(P, P')$, reporting the same approximate finite-order beliefs is an ε -BIBCE for the types in both priors. As we cannot a priori rule out the possibility of $\hat{T}_\varepsilon(P, P')$ containing the support of P' the scoring rule is not suitable for separating outcomes. Players therefore also choose an additional action: Either action zero or action one. No matter what additional action is chosen by the opponent, action one is the unique, strict best-reply for players who reported themselves to be in D_ε^0 . All other types will match action one if they believe with probability at least ε that their opponent also chose action one. Hence iterative deletion of dominated strategies implies that action one is played on D_ε^n . [Figure 4.3](#) illustrates the region D_ε^0 and the role of the scoring rule. Type profiles in the red region will play action one and an action represented by a circle in the right panel of the Figure. No type in the support of P' will play an action corresponding to a type in D_ε^0 . Action one will infect all type profiles in the orange shaded region who are also in the support of P but will not infect types in the support of P' .

Hence, in every ε -BIBCE the types in orange and red shaded regions who are in the support of P will play action 1. There is a ε -BIBCE where all types play action 0 under P' . This establishes that the sets of outcomes of the priors are at least ε apart. In this second part of the proof, we exploit properties of the product topology to ensure that we can cover the red region D_ε^0 with a finite grid and there is a finite game where players find it ε -optimal to report the closest element in the grid.

We now report two Lemmas before proceeding with the formal proof. Let $\mathcal{N}_\delta(E) := \bigcup_{\hat{\omega} \in E} \{\omega \in \mathcal{T} \times \Theta : d_\Pi(\hat{\omega}, \omega) < \delta\}$ denote the δ -ball around an event $E \in \mathcal{B}$.

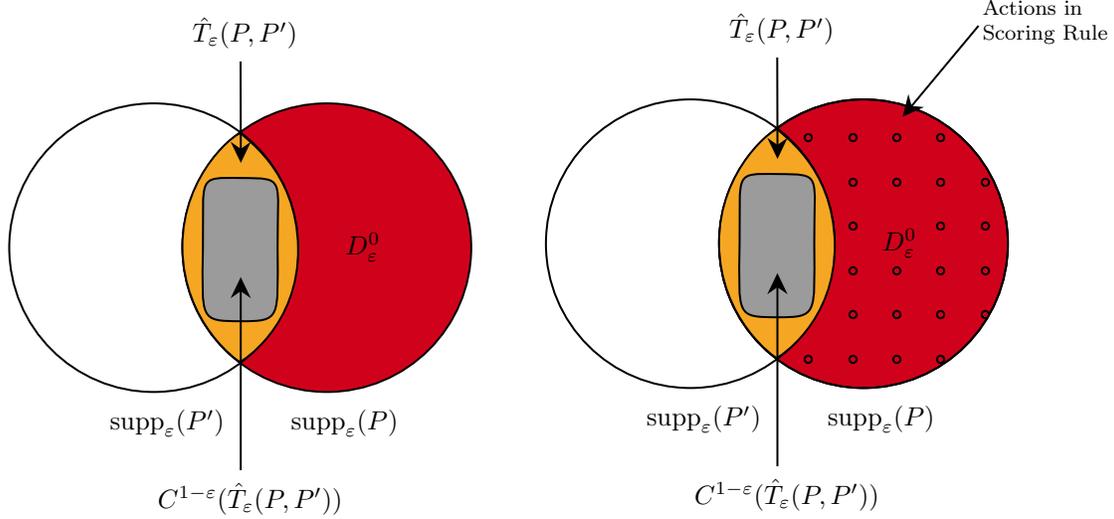


Figure 4.3: Infection argument.

Lemma 1. For every $\varepsilon > 0$ and any event $E \in \mathcal{B}$, there is a finite event $G_\varepsilon \subseteq \Omega \setminus \mathcal{N}_\varepsilon(E)$ so that for every $\omega \in \Omega$,

$$\min_{\omega_g \in G_\varepsilon} d_\Pi(\omega, \omega_g) < \varepsilon \implies \omega \in \Omega \setminus E$$

and

$$\omega \in \Omega \setminus \mathcal{N}_\varepsilon(E) \implies \min_{\omega_g \in G_\varepsilon} d_\Pi(\omega, \omega_g) < \varepsilon.$$

Lemma 1 implies that no type in the support of P' will be ε -close (in the product topology) to a grid element in the red shaded region D_ε^0 and every type in the support of P that is also in the red shaded region is ε -close (in the product topology) to a grid element in the red shaded region D_ε^0 . Lemma 2 below then establishes that there exists a finite game, where it is uniquely ε -rationalizable for every type to report the closest grid element (which is also less than ε -close).

Lemma 2. For every $\varepsilon > 0$, there is a finite set of action profiles $A \subseteq \mathcal{T}$ and payoffs $u_i^\varepsilon: A \times \Theta \rightarrow \mathbb{R}$ for every player i , so that for every information structure P , every ε -obedient decision rule σ satisfies

$$\sigma(a|\tau, \theta) > 0 \implies a \in \arg \min_{a \in A} d_\Pi((\tau, \theta), (a, \theta)) \text{ and } d_\Pi((\tau, \theta), (a, \theta)) < \varepsilon.$$

We now state two known results which we also use in our proof. The existence of BIBCE was established by [Stinchcombe \(2011\)](#), who also established it for general information structures, with redundancies as we later report in [Proposition 8](#).

Proposition 2. (*Existence of BIBCE*) There exists a BIBCE for every (\mathcal{G}, P) .

Theorem A in [Stinchcombe \(2011\)](#) has established the existence of a more demanding notion of incomplete information correlated equilibrium, when one looks at correlated equilibrium of the agent normal form of the game of incomplete information. [Forges \(1993b\)](#)

called these “agent normal form correlated equilibria.” Since every agent normal form correlated equilibrium induces an outcome equivalent BIBCE, existence of equilibria is guaranteed. We first establish that ACK-convergence is sufficient for continuity of equilibrium outcomes.

Proposition 3. *The set $\Omega^* \subseteq \Omega$ is dense in the product topology on Ω .*

Proof. Lipman (2003) shows that finite common prior types are dense in the product topology on Ω . \square

Proposition 4. *For every base game \mathcal{G} and every $\varepsilon > 0$, there is $\delta > 0$ so that if $d^{ACK}(P, P') < \delta$ then $d^*(P, P'|\mathcal{G}) < \varepsilon$.*

Proof. Let $\Omega_P := \cap_{\epsilon > 0} \text{supp}_\epsilon(P)$, $\Omega_{P'} := \cap_{\epsilon > 0} \text{supp}_\epsilon(P')$ and $\hat{\Omega}_\delta := C^{1-\delta}(\hat{T}_\delta(P, P'))$. It is without loss of generality to assume that $\Omega_P \cap \Omega_{P'} = \emptyset$. Fix a base game \mathcal{G} with payoffs given by $u_i: A \times \Theta \rightarrow [-M, M]$ for each player i . Suppose $d^{ACK}(P, P') < \delta = \frac{\varepsilon}{6M(u)}$, where $M(u) := \max\{M, |A \times \Theta|\} < \infty$ and recall that $2M \geq \max_i \sup_{a_i, a'_i, a_{-i}, \theta} |u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)|$. By Lemma 1, there is a finite set G_δ so that $\hat{\Omega}_\delta \subseteq \mathcal{N}_\varepsilon(G_\delta)$ and so for every $\omega \in \hat{\Omega}_\delta$, $\min_{g \in G_\delta} d_\Pi(\omega, g) \leq \delta$. Let $\zeta_\delta: \hat{\Omega}_\delta \rightarrow G_\delta$ be any measurable map satisfying $d_\Pi(\omega, \zeta_\delta(\omega)) \leq \delta$, $\forall \omega \in \hat{\Omega}_\delta$. We show that every decision rule which is obedient under P , admits a decision rule arbitrarily close to it which is $6M\delta$ -obedient under P' . Consider a decision rule $\sigma: \Omega_P \rightarrow \Delta(A)$ satisfying for every τ_i and for every $a'_i \in A_i$,

$$\int_{\Omega_P} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \sigma(a|\omega) P(d\omega|\tau_i) \geq 0,$$

where $\Delta u_i(a, a'_i, \theta) = u_i(a, \theta) - u_i(a'_i, a_{-i}, \theta)$ and $\omega_\theta = \text{proj}_\Theta(\omega)$. We have existence of such a decision rule from Proposition 2. We now extend σ to $\hat{\Omega}_\delta$ as follows: For every $\omega \in \hat{\Omega}_\delta \cup \Omega_P$, define the δ -extension

$$\sigma_\delta(\omega) := \begin{cases} \int \sigma(\omega') P(d\omega'|\zeta_\delta(\omega)) & \text{if } \omega \in \hat{\Omega}_\delta, \\ \sigma(\omega) & \text{if } \omega \notin \hat{\Omega}_\delta. \end{cases}$$

For all i and $\tau \in \text{proj}_{\mathcal{T}}(\Omega_P)$,

$$\begin{aligned}
& \left| \int_{\hat{\Omega}_\delta \cup \Omega_P} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) (\sigma_\delta(a|\omega) - \sigma(a|\omega)) \tau_i(d\omega) \right| \\
&= \left| \sum_{g \in G_\delta, a \in A} \left(\int_{\zeta_\delta^{-1}(g)} \Delta u_i(a, a'_i, \omega_\theta) (\sigma_\delta(a|\omega) - \sigma(a|\omega)) \tau_i(d\omega|\zeta_\delta^{-1}(g)) \right) \tau_i(\zeta_\delta^{-1}(g)) \right| \\
&\leq 2M \left| \sum_{g \in G_\delta, a \in A} \left(\int_{\zeta_\delta^{-1}(g)} (\sigma_\delta(a|\omega) - \sigma(a|\omega)) \tau_i(d\omega|\zeta_\delta^{-1}(g)) \right) \tau_i(\zeta_\delta^{-1}(g)) \right| \\
&\leq \sum_{g \in G_\delta, a \in A} 2M \left| \left(\int \left(\int \sigma(a|\omega') P(d\omega'|\zeta_\delta^{-1}(g), \tau'_i) \right) P(d\tau'_i|\zeta_\delta^{-1}(g)) \right) - \int \sigma(a|\omega') \tau_i(d\omega|\zeta_\delta^{-1}(g)) \right| \tau_i(\zeta_\delta^{-1}(g)) \\
&= \sum_{g \in G_\delta, a \in A} 2M \left| \int \left(\int \sigma(a|\omega') (\tau'_i(d\omega'|\zeta_\delta^{-1}(g)) - \tau_i(d\omega'|\zeta_\delta^{-1}(g))) \right) P(d\tau'_i|\zeta_\delta^{-1}(g)) \right| P(\zeta_\delta^{-1}(g)|\tau_i) \\
&\leq 2M\delta.
\end{aligned}$$

Since σ is obedient, we conclude that σ_δ is $2M\delta$ -obedient. By construction, for every $(\tau', \theta') \in \hat{\Omega}_\delta \cap \Omega_{P'}$ there exists $(\tau, \theta) \in \Omega_P$ so that $d_\Pi((\tau', \theta'), (\tau, \theta)) < \delta$, which directly implies that

$$\int_{\hat{\Omega}_\delta \cup \Omega_P} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \sigma_\delta(a|\omega) (dP(\omega|\tau_i) - dP'(\omega|\tau'_i)) < 2M\delta.$$

By the $2M\delta$ -obedience of σ_δ established above, we have that

$$\int_{\hat{\Omega}_\delta} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \sigma_\delta(a|\omega) dP'(\omega|\tau'_i) \geq -4M\delta.$$

Moreover, $(\tau', \theta') \in \hat{\Omega}_\delta$ implies that for any measurable $\hat{\sigma}: \Omega_{P'} \setminus \hat{\Omega}_\delta \rightarrow \Delta(A)$, we have $\int_{\Omega_{P'} \setminus \hat{\Omega}_\delta} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \hat{\sigma}(a|\omega) P'(d\omega|\tau'_i) \geq -2M\delta$ and so

$$\int_{\Omega_{P'}} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \sigma_\delta(a|\omega) P'(d\omega|\tau'_i) \geq -6M\delta. \tag{4.1}$$

So σ_δ satisfies $6M\delta$ -obedience under P' if restricted to types τ'_i that belong to $\hat{\Omega}_\delta$. We now argue that there exists a $6M\delta$ -obedient decision rule under P' that agrees with the extension σ_δ on $\hat{\Omega}_\delta$. For every player i , let $D_i := \left\{ \tau \in \mathcal{T} : \tau_i(\hat{\Omega}_\delta) \leq 1 - \delta \right\}$ and note that if $P'(\hat{\Omega}_\delta) < 1$ then $P'(\bigcup_{i \in I} D_i) > 0$.

Consider the auxiliary payoffs $\tilde{u}: \Omega \times A \rightarrow \mathbb{R}^I$, defined for every $(\tau, \theta) \in \Omega$ and $a \in A$ as follows

$$\tilde{u}_i((\tau, \theta), a) := \begin{cases} u_i(a_{-i}, a_i, \theta), & \text{if } \tau_i(\hat{\Omega}_\delta) \leq 1 - \delta \\ \mathbf{1}_{\sigma_\delta(a|\tau, \theta) > 0}, & \text{otherwise.} \end{cases}$$

By Proposition 2 we deduce that the incomplete information game (\tilde{u}, P') admits an obedient decision rule $\bar{\sigma}$ that coincides with σ_δ on $\hat{\Omega}_\delta$. Now, consider the combined decision rule, which for every $\omega \in \Omega_P$ satisfies

$$\sigma'(\omega) := \begin{cases} \sigma_\delta(\omega) & , \text{ if } \omega \in \hat{\Omega}_\delta \\ \bar{\sigma}(\omega) & , \text{ if } \omega \in \Omega_{P'} \setminus \hat{\Omega}_\delta. \end{cases}$$

We have shown that for every player i and every type $(\tau, \theta) \in \Omega_{P'}$ satisfying $\tau_i(\hat{\Omega}_\delta) > 1 - \delta$, the combined decision rule σ' satisfies $6M\delta$ -obedience. Finally, by obedience of $\bar{\sigma}$ for every player i and $(\tau', \theta) \in \Omega_{P'}$ so that $\tau'_i(\hat{\Omega}_\delta) \leq 1 - \delta$,

$$\int_{\Omega_{P'}} \sum_{a \in A} \Delta \tilde{u}_i(a, a'_i, \omega) \bar{\sigma}(a|\omega) P'(d\omega|\tau'_i) = \int_{\Omega_{P'}} \sum_{a \in A} \Delta u_i(a, a'_i, \omega_\theta) \bar{\sigma}(a|\omega) P'(d\omega|\tau'_i) \geq 0. \quad (4.2)$$

Combining (4.1) and (4.2) we deduce that σ' is $6M\delta$ -obedient under P' . We will now show that $\nu_{P', \sigma'}$ is close to $\nu_{P, \sigma}$. For any $(a, \theta) \in A \times \Theta$,

$$\begin{aligned} & |\nu_{P, \sigma}(a, \theta) - \nu_{P', \sigma'}(a, \theta)| \\ & \leq \left| \sum_{g \in \mathcal{G}_\delta} \left(\int \sigma(a|\tau', \theta) P(d\tau', \theta|g) \right) (P(d\tau, \theta) - P'(d\tau, \theta)) \right| \\ & \quad + |P'(\hat{\Omega}_\delta|\theta) - P(\hat{\Omega}_\delta|\theta)| \leq 2\delta. \end{aligned}$$

Hence $\sum_{a, \theta} |\nu_{P, \sigma}(a, \theta) - \nu_{P', \sigma'}(a, \theta)|^2 < |A \times \Theta| 4\delta^2$ and so $d^*(P, P'|\mathcal{G}) < 6M(u)\delta$. \square

We now establish that failure of ACK-convergence implies a failure of convergence of equilibrium outcomes.

Proposition 5. *For every $\varepsilon > 0$, if $d^{ACK}(P, P') \geq \varepsilon$ then there is a game \mathcal{G} so that $d^*(P, P'|\mathcal{G}) \geq \varepsilon$.*

Proof. We now establish that for all $\varepsilon \in (0, 1/2)$ so that for all P, P' satisfy $d^{ACK}(P, P') \geq \varepsilon$, we also have $d^*(P, P'|\mathcal{G}) \geq \varepsilon$ for some game \mathcal{G} : If convergence fails in our metric, then there must be some game on which ex-ante strategic convergence fails. Then we must find such a game \mathcal{G} . The condition $d^{ACK}(P, P') \geq \varepsilon$ means that $P(\hat{\Omega}_\varepsilon) \leq 1 - \varepsilon$, $P'(\hat{\Omega}_\varepsilon) \leq 1 - \varepsilon$ or both. Suppose that $P(\hat{\Omega}_\varepsilon) \leq 1 - \varepsilon$ and $P'(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$. First, note that $P(\hat{T}_\varepsilon(P, P')) < 1$. Indeed, if $P(\hat{T}_\varepsilon(P, P')) = 1$, we also have that $P(\hat{\Omega}_\varepsilon) = 1$, which is a contradiction. Let $D_{\varepsilon, P} := \text{supp}_\varepsilon(P) \setminus \hat{T}_\varepsilon(P, P')$ and $D_{\varepsilon, P}^c := \Omega \setminus \mathcal{N}_\varepsilon(D_{\varepsilon, P})$. From Lemmas 1 and 2 we conclude that there is m and z , an associated game $\hat{\mathcal{G}} = (A^{m, z}, (u_i^{m, z})_i)$, where the finite collection of action profiles takes the form $A^{m, z} = \hat{D}_{\varepsilon, P} \cup \hat{D}_{\varepsilon, P}^c$, with $\hat{D}_{\varepsilon, P} \subseteq D_{\varepsilon, P}$ a finite ε -grid, $\hat{D}_{\varepsilon, P}^c \subseteq D_{\varepsilon, P}^c$ a finite ε -grid, so that for every ε -BIBCE, $\hat{\sigma}' \in \mathcal{B}^\varepsilon(\hat{\mathcal{G}}, P')$ and every $\omega \in \Omega_{P'}$, $\hat{\sigma}'(\hat{D}_{\varepsilon, P}|\omega) = 0$. Moreover, for every $\omega \in D_{\varepsilon, P}$ and every ε -BIBCE, $\hat{\sigma} \in \mathcal{B}^\varepsilon(\hat{\mathcal{G}}, P)$, $\hat{\sigma}(\hat{D}_{\varepsilon, P}|\omega) = 1$.

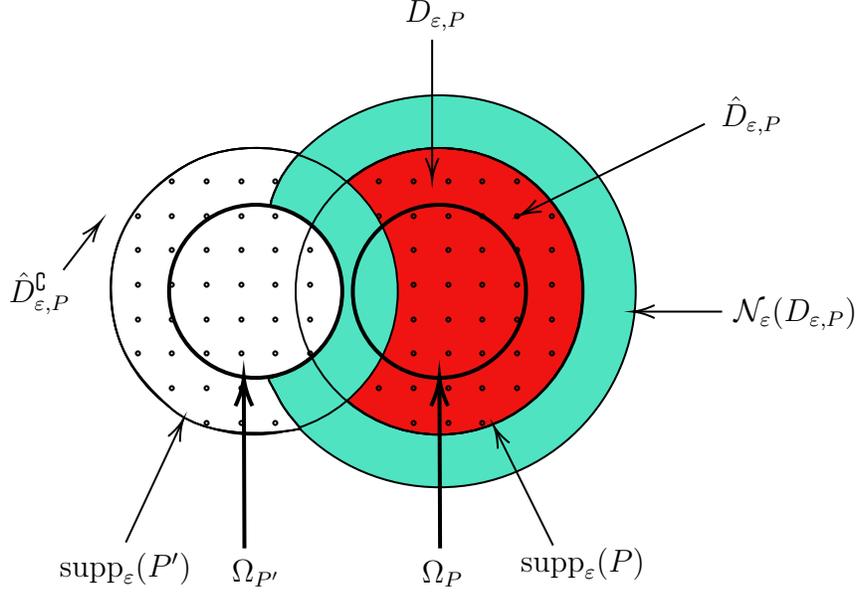


Figure 4.4: Domain of Grid Game.

Figure 4.4 above illustrates the sets we just defined. The area enclosed by left the two bold circle represents the support of P' , $\Omega_{P'}$ and the area enclosed by the right bold circle represents the support of P , Ω_P . The red shaded area represents the set $D_{\varepsilon,P}$ where the dots on top of the red shaded area represents $\hat{D}_{\varepsilon,P}$, the actions chosen by type profiles in $D_{\varepsilon,P}$. The green shaded area represents $\mathcal{N}_{\varepsilon}(D_{\varepsilon,P}) \setminus D_{\varepsilon,P}$, which contains no action in $A^{m,z}$. Finally, the dots with white background represent the remaining actions, i.e. the set $\hat{D}_{\varepsilon,P}^c$. Every type in $\Omega_{P'}$ will pick an action from that set. Define $D_P^m := \Omega_P \cap D_{\varepsilon,P}$ and for every player i , let $D_{P,i}^m := \{\tau_i : \exists (\hat{\tau}, \theta) \in \Omega_P, ((\tau_i, \hat{\tau}_{-i}), \theta) \in D_P^m\}$. Define the sets $D_{P'}^m$ and $D_{P',i}^m$ for prior P' analogously. Define the sequence $(D_P^{m+n})_{n \in \mathbb{N}}$, recursively for every $n \geq 1$,

$$D_{P,i}^{m+n} := \Omega_P \setminus B_i^{1-\varepsilon}(\Omega_P \setminus D_P^{m+n-1}), \quad D_P^{m+n} := \prod_{i \in I} D_{P,i}^{m+n}.$$

Let $\hat{D}_{P,i} := \{\tau_i : \exists (\tau_{-i}, \theta) \in \mathcal{T}_{-i} \times \Theta \text{ s.t. } (\tau, \theta) \in \hat{D}_{\varepsilon,P}\}$. Consider the action set $A_i^* := \{a^*, a^{**}\}^I \times A^{m,z}$ and let payoffs of \mathcal{G} be given by

$$\hat{u}_i(a^0, d, \theta) = \begin{cases} \varepsilon \mathbf{1}_{a_i^0 = a^*} + u_i^{m,z}(d, \theta) & \text{if } d_i \in \hat{D}_{P,i} \\ \mathbf{1}_{a_i^0 = a^*} + u_i^{m,z}(d, \theta) & \text{if } d_i \notin \hat{D}_{P,i} \text{ and } \exists j \neq i \text{ s.t. } a_j^0 = a^* \end{cases}$$

In this game, playing $d_i \in \hat{D}_{P,i}$ and a^* is a dominant action for every type in $D_{P,i}^m$. Under payoffs $u^{m,z}$ and prior P' , reporting a grid element in D_P^m has cost at least $\varepsilon > 0$. Hence every type under prior P' rationalizes both a^{**} and a^* while no type under prior P in the set $D_{P,i}^m$ rationalizes a^{**} . Now consider $\tau_i \in D_{P,i}^{m+1}$. Again, playing $a_i^0 = a^*$ is uniquely ε -rationalizable since all types in $D_{P,i}^m$ play a^* : Indeed, for every $\tau_i \in D_{P,i}^{m+1}$, there is a player i so that $\tau_i(D_{P,-i}^m) \geq \varepsilon$. For every $n \in \mathbb{N}$, there is a player i_n so that $D_{P,i_n}^{m+n} \neq \emptyset$. Proceeding

inductively we obtain that $a_i^0 = a^*$ remains the unique ε -best-reply for some player at type profiles in $\Omega_P \setminus \hat{\Omega}_\varepsilon$. Deduce that there is a ε -BIBCE where all types of all players play a^{**} under P' while for all ε -BIBCE under P , P assigns at least probability ε to type profiles where some player's type plays a^* . Hence $d^*(P, P'|\mathcal{G}) \geq \varepsilon$. A symmetric argument shows that whenever 1) $P(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$ and $P'(\hat{\Omega}_\varepsilon) \leq 1 - \varepsilon$ or 2) $P(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$ and $P'(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$ implies that outcomes are ε apart under a similarly constructed game. Note that under $P(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$ and $P'(\hat{\Omega}_\varepsilon) > 1 - \varepsilon$, the game with payoffs $u^{m,z}$ is enough to separate the outcomes of P and P' .

The grid game alone would not work for our proof: Indeed, we could construct a sequence of priors P^k for any prior P so that $d^{ACK}(P^k, P) > \varepsilon$ for all k but where $P^k(T_\varepsilon(P^k, P)) \uparrow 1$, i.e. the infecting event has diminishing ex-ante probability. If we used the grid game only, we would have that $d^*(P, P'|\hat{\mathcal{G}}) \downarrow 0$. Based on this game alone, the ACK topology would be too strong. \square

Now Theorem 1 follows immediately from Propositions 4 and 5.

We will conclude this section by discussing a key property of ACK topology (denseness of finite information structures), which will be important for our applications.

4.3.4 Denseness of Finite Information Structures

It is often convenient to work with finite information structures. We have:

Proposition 6. *Finite information structures are dense in \mathcal{P} under the ACK topology.*

Proof. Fix any information structure $P \in \mathcal{P}$. For any $\delta > 0$ we construct a finite information structure P_δ so that $d^{ACK}(P, P_\delta) < \delta$. To do so, we construct a finite grid G_δ of state and type pairs whose δ -neighborhood covers Ω_P , i.e. $\Omega_P \subseteq \mathcal{N}_\delta(G_\delta)$ and for all $\omega \in \Omega_P$, $\min_{g \in G_\delta} d_\Pi(\omega, g) \leq \delta$. Consider any partition on Ω_P given by the pre-image of any map $\zeta: \Omega_P \rightarrow G_\delta$, where for all $\omega \in \Omega_P$, we have that $d_\Pi(\omega, \zeta(\omega)) \leq \delta$. Consider the information structure P_δ given by $P \circ \zeta^{-1}$. This information structure is finite and has the property that for any $\tau \in \text{supp}(P_\delta)$ there exists $\tau' \in \text{supp}(P)$ so that $d_\Pi(\tau, \tau') < \delta$ and vice versa; for all $\tau' \in \text{supp}(P)$ there is $\tau \in \text{supp}(P_\delta)$ so that $d_\Pi(\tau, \tau') < \delta$. Indeed, for every player i and $\omega \in \Omega_P$ and $(\tau, \theta) = \zeta(\omega)$, beliefs of any measurable event $E \in \mathcal{B}$ at τ_i are given by,

$$P \circ \zeta^{-1}(E|\tau_i) = \int P(E|\hat{\tau}_i)P(d\hat{\tau}_i|\zeta_i^{-1}(\tau_i)),$$

where $\zeta_i^{-1}(\tau_i) := \bigcup_{(\hat{\tau}, \hat{\theta}) \in G_\delta: \hat{\tau}_i = \tau_i} \left\{ \tilde{\tau}_i : \exists (\tilde{\tau}_{-i}, \tilde{\theta}) \text{ s.t. } (\tilde{\tau}, \tilde{\theta}) \in \zeta^{-1}(\hat{\tau}, \hat{\theta}) \right\}$. First order beliefs of $P \circ \zeta^{-1}$ take the form $P \circ \zeta^{-1}(\theta|\tau_i) = \int P(\theta|\hat{\tau}_i)P(d\hat{\tau}_i|\zeta_i^{-1}(\tau_i))$ and so for every $\tilde{\tau}_i \in \zeta_i^{-1}(\tau_i)$, first-order beliefs of $P \circ \zeta^{-1}(\cdot|\tau_i)$ and $P(\cdot|\tilde{\tau}_i)$ are no more than δ apart in the Euclidean topology on $\Delta(\Theta)$. Deduce that $P \circ \zeta^{-1}(\cdot|\tau_i)$ and $P(\cdot|\tilde{\tau}_i)$ have δ -close hierarchies of beliefs in the product topology on \mathcal{T}_i . Hence $\hat{T}_\delta(P, P_\delta) = \text{supp}_\delta(P)$ and so $d^{ACK}(P, P_\delta) < \delta$. For δ small enough, there is $\hat{\zeta}: \mathcal{T} \rightarrow \mathcal{T}$, so that our choice of ζ above satisfies $\zeta(\tau, \theta) = (\hat{\zeta}(\tau), \theta)$ and so the marginal of $P \circ \zeta^{-1}$ on Θ coincides with that of P . Hence finite information structures are dense in \mathcal{P} under the strategic topology. \square

4.3.5 Denseness of Simple Information Structures

A finite information structure P is a *first-order belief* information structure if each type has a distinct first-order belief. An information structure P is *simple* if it is both a finite and a first-order belief information structure. We denote the collection of simple information structures

$$\mathcal{P}^* := \left\{ P \in \mathcal{P} : |\text{supp}(P)| < \infty, \forall (\tau, \theta), (\hat{\tau}, \hat{\theta}) \in \text{supp}(P), \forall i, \tau_i \neq \hat{\tau}_i \implies \tau_i^1 \neq \hat{\tau}_i^1 \right\}$$

It is often convenient to work with simple information structures. We have:

Proposition 7. *Simple information structures are dense in \mathcal{P} under the ACK topology.*

Proof. We have shown in Proposition 6 that finite information structures are dense in \mathcal{P} under the strategic topology. We have shown this by constructing a finite prior $P_\delta \in \mathcal{P}$, for any $P \in \mathcal{P}$ and any $\delta > 0$ so that $d^{ACK}(P, P_\delta) < \delta$. It remains to show that P_δ is close to a simple information structure. Let $\Omega_{i, P_\delta}(\tau_i) := \{\hat{\tau}_i : P_\delta(\hat{\tau}_i) > 0, \hat{\tau}_i^1 = \tau_i^1\}$. For any $\varepsilon > 0$ let $\rho_{i, \varepsilon} : \mathcal{T}_i \times \Theta \rightarrow \Delta(\{0, 1\})$ have the property that $\rho_{i, \varepsilon}(\hat{\tau}_i, \theta) = \rho_{i, \varepsilon}(\hat{\tau}_i, \theta')$ for all $\theta, \theta' \in \Theta$ and all $\hat{\tau}_i$ in the support of P_δ satisfying $|\Omega_{i, P_\delta}(\hat{\tau}_i)| = 1$. For any $\hat{\tau}_i$ in the support of P_δ satisfying $|\Omega_{i, P_\delta}(\hat{\tau}_i)| > 1$ let $\|\rho_{i, \varepsilon}(\hat{\tau}_i, \theta) - \rho_{i, \varepsilon}(\hat{\tau}_i, \theta')\|_2 < \varepsilon$ so that for all distinct $\tilde{\tau}_i, \bar{\tau}_i \in \Omega_{i, P_\delta}(\tau_i)$ and $s_i, s'_i \in \{0, 1\}$,

$$\frac{\rho_{i, \varepsilon}(s_i | \tilde{\tau}_i, \theta) P_\delta(\theta | \tilde{\tau}_i)}{\sum_{\hat{\theta}} \rho_{i, \varepsilon}(s_i | \tilde{\tau}_i, \hat{\theta}) P_\delta(\hat{\theta} | \tilde{\tau}_i)} \neq \frac{\rho_{i, \varepsilon}(s'_i | \bar{\tau}_i, \theta) P_\delta(\theta | \bar{\tau}_i)}{\sum_{\hat{\theta}} \rho_{i, \varepsilon}(s'_i | \bar{\tau}_i, \hat{\theta}) P_\delta(\hat{\theta} | \bar{\tau}_i)}.$$

The prior $\hat{P}_\delta(\tau, s, \theta) = P_\delta(\tau, \theta) \prod_i \rho_{i, \varepsilon}(s_i | \tau, \theta)$ induces information structure \tilde{P}_δ so that $d^{ACK}(P, \tilde{P}_\delta) < \delta + \varepsilon$. Deduce that simple information structures are dense in \mathcal{P} . \square

Arieli et al. (2022) study first-order belief information structures with possibly uncountable support when Θ is binary. In that case, the set of first-order belief information structures can be identified with a subset of $\Delta([0, 1]^I)$. Cichomski and Osekowski (2024); Arieli et al. (2022); Cichomski and Osekowski (2025) study the extreme points of this set. Every general information structure can be represented in \mathcal{P} by recovering its induced distribution on belief hierarchies. Moreover, every information structure in \mathcal{P} can be represented as a first-order belief information structure via the canonical projection $\tau \mapsto \tau^1$. At each step, information is lost: By representing a general information structure in \mathcal{P} redundancies (and thus correlation possibilities) are lost and by representing an information structure in \mathcal{P} as a first-order belief information structure all information about higher order beliefs that is not also encoded in first-order beliefs is lost. The set \mathcal{P} (and by extension the class of general information structures) thus seems to be richer than the set of first-order belief information structures. However, Proposition 7 implies the set of first-order belief information structures are in fact rich enough to generate all equilibrium outcomes in all finite games (at least up to any degree of approximation). The approach taken in the papers mentioned above may thus be used to study the impact of information on strategic behavior.

4.4 General Information Structures and Bayes Nash Equilibrium

Our main approach in this paper is to remove correlating devices from the information structure (and thus work with non-redundant information structures) and put correlating devices in the solution concept (BIBCE). However, to relate our work to the literature and discuss applications it is useful to discuss how our results apply when we allow correlating devices in the information structure (and thus allow for general *redundant information structures*) but remove correlating devices from the solution concept (and work with Bayes Nash equilibrium).

4.4.1 General Information Structures

A (common prior) *general information structure* describes a set of signals for each player and a joint distribution over states and profiles of signals: thus a general information structure is a tuple $\mathcal{S} = ((S_i)_{i \in I}, Q)$, where each S_i is a measurable space of signals⁶ for player i and Q is a probability measure⁷ in $\Delta(S \times \Theta)$ whose marginal on Θ is given by μ .

We dub this object a general information structure. Liu (2015b) describes how one can always decompose a general information structure into a (non-redundant) information structure and a correlation device. While a general information structure is described by a signal space and a probability measure, we will adopt the convention of describing non-redundant information structures as just a prior, leaving it understood that the player's signal space is the universal space of hierarchies.

Every redundant information structure can be naturally mapped to its non-redundant information structure by essentially integrating out redundant types. In particular, we map information structure $\mathcal{S} = ((S_i)_{i \in I}, Q)$ to a (non-redundant) information structure P as follows. For every i , and version of the conditional probability Q_i , first-order beliefs can be obtained for any player i , $\bar{\tau}_i^1(s_i) = \text{marg}_{\Theta}(Q_i(s_i))$ for every $s_i \in S_i$. For every $m > 1$ and $(m - 1)$ -order beliefs representation $\bar{\tau}_j^{m-1}: S_j \rightarrow \mathcal{T}_j^{m-1}$, for every j , obtain the m -order belief representation of player i , $\bar{\tau}_i^m(s_i) = Q_i(s_i) \circ (\bar{\tau}_{-i}^{m-1} \times \text{id})^{-1}$, where id is the identity on Θ and $\bar{\tau}_{-i}^{m-1}: s_{-i} \mapsto (\bar{\tau}_j^{m-1}(s_j))_{j \neq i}$. The representation of s_i in \mathcal{T}_i is then given by $\bar{\tau}_i(s_i) = (\bar{\tau}_i^m(s_i))_m$. For every $s \in S$ let $\bar{\tau}(s) := (\bar{\tau}_i(s_i))_{i \in I}$ and we write $P_{\mathcal{S}}$ for the information structure thus induced by redundant information structure \mathcal{S} .

4.4.2 Solution Concepts

Now we will say that a base game and a redundant information structure $(\mathcal{G}, \mathcal{S})$ together define a *Bayesian game*.

⁶We adopt the convention of referring to “signals” rather than “types” when describing general information structures in this section. We reserve the terminology “type” for hierarchies of beliefs, introduced in the next section.

⁷The product of measurable spaces is always endowed with the product sigma algebra.

Belief-invariant Bayes correlated equilibrium will be defined as before for general information structures and Bayesian games. For completeness, we will spell out the definition allowing for general information structures, and also define Bayes Nash Equilibrium.

For any Bayesian game $(\mathcal{G}, \mathcal{S})$, a decision rule is a measurable map $\sigma : S \times \Theta \rightarrow \Delta(A)$, where $A := \prod_{i \in I} A_i$ and $\Delta(A)$ is endowed with the Euclidean topology. A general information structure \mathcal{S} and decision rule σ jointly induce a measure $\sigma \circ Q \in \Delta(A \times S \times \Theta)$ in the natural way. We will be interested in *outcomes* specifying a joint distribution over actions and states $\nu \in \Delta(A \times \Theta)$. Decision rule σ *induces* outcome ν_σ if ν_σ is the marginal of $\sigma \circ Q$ on $A \times \Theta$. For every player i , a decision rule σ and a version of the conditional probability $Q_i : S_i \rightarrow \Delta(S_{-i} \times \Theta)$ of Q induce a belief for every signal $s_i \in S_i$, $\sigma \circ Q_i : S_i \rightarrow \Delta(A \times S_{-i} \times \Theta)$, which for every measurable set $E \subseteq A \times S_{-i} \times \Theta$ and every signal $s_i \in S_i$ satisfies $\sigma \circ Q_i(E|s_i) = \int_E \sigma(a|\theta, s_{-i}, s_i) dQ_i(\theta, s_{-i}|s_i)$.

Now we have:

Definition 12. A decision rule σ is ε -obedient if, for every player i , there is a version of the conditional probability $Q_i : S_i \rightarrow \Delta(S_{-i} \times \Theta)$ so that every action $a_i \in A_i$ and deviation a'_i ,

$$\int_{S \times \Theta} \sum_{a_{-i} \in A_{-i}} (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) d\sigma \circ Q_i(a_i, a_{-i}, s_{-i}, \theta|s_i) > -\varepsilon, \text{ a.s.}$$

Definition 13. A decision rule σ is belief-invariant if, for every $a_i \in A_i$, the marginal probability $\sigma(a_i \times A_{-i} | (s_i, s_{-i}), \theta) = \sigma_i(a_i | s_i)$ does not depend on (s_{-i}, θ) .

Definition 14. A decision rule σ is an ε -belief-invariant Bayes correlated equilibrium (ε -BIBCE) of $(\mathcal{G}, \mathcal{S})$ if it satisfies ε -obedience and belief invariance.

If we have a (non-redundant) information structure, these definitions reduce to those introduced earlier.

Now a decision rule σ is *conditionally independent* if $\sigma(a|(s_i, s_{-i}), \theta) = \prod_{i \in I} \sigma_i(a_i|s_i)$, for every $(a, s, \theta) \in A \times S \times \Theta$. Conditional independence requires that any randomization in a player's actions depends on their type only and is conditionally independent of others' types and the state. If a decision rule is conditionally independent, it gives a behavioral strategy for each player in the incomplete information game.

Definition 15. A decision rule σ is an ε -Bayes Nash equilibrium (ε -BNE) of $(\mathcal{G}, \mathcal{S})$ if it satisfies obedience (ε -obedience), belief-invariance and conditional independence.

We will say that a decision rule is a BIBCE or BNE if it is a 0-BIBCE or 0-BNE respectively. We will write $\mathcal{B}(\mathcal{G}, \mathcal{S})$ and $\mathcal{B}^{BNE}(\mathcal{G}, \mathcal{S})$ for the set of BIBCE and BNE decision rules, and ν_σ for the outcome in $\Delta(A \times \Theta)$ induced by decision rule σ ; we will write $\mathcal{O}(\mathcal{G}, \mathcal{S})$ for the set of BIBCE outcomes

$$\mathcal{O}(\mathcal{G}, \mathcal{S}) := \{\nu_\sigma \in \Delta(A \times \Theta) : \sigma \in \mathcal{B}(\mathcal{G}, \mathcal{S})\}$$

and $\mathcal{O}^{BNE}(\mathcal{G}, \mathcal{S})$ for the set BNE outcomes:

$$\mathcal{O}^{BNE}(\mathcal{G}, \mathcal{S}) := \{\nu_\sigma \in \Delta(A \times \Theta) : \sigma \in \mathcal{B}^{BNE}(\mathcal{G}, \mathcal{S})\}$$

4.4.3 Existence of Equilibria

We have existence of BIBCE.

Proposition 8. (*Existence of BIBCE*) There exists a BIBCE for every $(\mathcal{G}, \mathcal{S})$.

This was already established for (non-redundant) information structures in Section 4.3. For a redundant information structure, it is enough to find a BIBCE for its non-redundant version and extend the BIBCE to have players ignore redundancies.

However, strong conditions are required to ensure the existence of BNE, and this has been one obstacle to constructing topologies on (general) information structures for BNE. [Milgrom and Weber \(1985\)](#) and [Balder \(1988\)](#) are apparently the best available (even if we restrict attention to finite action games). Two sufficient conditions from [Milgrom and Weber \(1985\)](#) are important. First, existence is guaranteed if information structures have countable support. Second, existence is guaranteed if the measure on signals is absolutely continuous with respect to the product of the marginal measures on individual player's signals. [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#) therefore restricted attention countable information structures. However, existence of BNE or even ε -BNE fails when these properties fail: see [Simon \(2003\)](#), [Hellman \(2014\)](#) and [Simon and Tomkowicz \(2017\)](#) for examples. In particular, existence is not guaranteed on (non-redundant) information structures.⁸ In the Appendix, we report an example from [Hellman \(2014\)](#) where BNE fails to exist and report a BIBCE for that example.

4.4.4 Varying Solution Concepts and Information Structures

We first discuss the connection between BNE and BIBCE outcomes. The next proposition states that the set of BIBCE outcomes depends only on the non-redundant information structure:

Proposition 9. *For any base game \mathcal{G} and general information structure \mathcal{S} , $\mathcal{O}(\mathcal{G}, \mathcal{S}) = \mathcal{O}(\mathcal{G}, P_{\mathcal{S}})$.*

This is true because any need for redundancies / correlating devices is built into the solution concept. This observation parallels the observation of [Dekel et al. \(2007b\)](#) that interim correlated rationalizability depends only on hierarchies of beliefs; [Liu \(2015b\)](#) showed that (a subjective version of) BIBCE is equilibrium analogue of ICR and thus provides a proof.⁹ For completeness, we give a proof in our notation in the Appendix. It is immediate from the definitions that BNE outcomes are BIBCE outcomes for any information structure.

Proposition 10. *For any game \mathcal{G} and general information structure \mathcal{S} , $\mathcal{O}^{BNE}(\mathcal{G}, \mathcal{S}) \subseteq \mathcal{O}(\mathcal{G}, \mathcal{S})$.*

Now define the strategic distance between a pair of general information structures \mathcal{S} and \mathcal{S}' to be the strategic distance between their non-redundant representations, so $d^{**}(\mathcal{S}, \mathcal{S}') := d^*(P_{\mathcal{S}}, P_{\mathcal{S}'})$.

⁸[van Zandt \(2010\)](#) establishes existence of BNE on the universal type space for supermodular games.

⁹See also [Bergemann and Morris \(2017\)](#) for a discussion of these issues.

4.4.5 Correlation and Non-Redundant Information Structures

On general information structures, players' ability to correlate their actions (using redundant correlating devices) matters for the set of BNE. However, non-redundant information structures, as long as there are at least two states, are very rich objects and, intuitively, there will be plenty of opportunity to approximate arbitrary correlating devices within them.

Our use of BIBCE as a solution concept allowed us to focus attention on (non-redundant) information structures and ensured existence of equilibrium. Thus we obtained cleaner results with this solution concept. But suppose one is interested in Bayes Nash equilibrium. In this case, redundant types / correlating devices potentially matter for equilibrium. And we potentially have problems with equilibrium existence. In this section, we will argue that there is a natural approach to dealing with these difficulties (maintaining BNE as the preferred solution concept) and that the same almost common knowledge topology is relevant for continuity of equilibrium outcomes.

Our main observation is that, since the universal type space is rich (as long as there at least two states), any correlating device can be embedded in a (non-redundant) information structure by perturbing types' first-order beliefs. The following proposition establishes that any BIBCE on any finite information structure can be approximated by an approximate BNE of some nearby simple information structure. Let $\mathcal{B}^{\varepsilon BNE}(\mathcal{G}, P)$ denote the collection of ε -Bayes Nash equilibria.

Proposition 11. *Let $|\Theta| \geq 2$. For any finite, general information structure \mathcal{S} , any $\sigma \in \mathcal{B}(\mathcal{G}, \mathcal{S})$ and any $\varepsilon > 0$, there exists (i) a simple information structure \mathcal{S}' such that $d^{**}(\mathcal{S}, \mathcal{S}') \leq \varepsilon$; (ii) a decision rule σ' such that (a) $\sigma' \in \mathcal{B}^{\varepsilon BNE}(\mathcal{G}, \mathcal{S}')$ and (b) the outcome induced by $\mathcal{S}' \circ \sigma'$ is ε -close to the outcome induced by $\mathcal{S} \circ \sigma$.*

Proof. Let $Q \circ \sigma \in \Delta(A \times S \times \Theta)$ be the extended outcome corresponding to $\sigma \in \mathcal{B}(\mathcal{G}, \mathcal{S})$. Consider the general information structure where each player's signal space was $\bar{S}_i = A_i \times S_i$ and the prior was $Q \circ \sigma$. Note that since σ is an arbitrary BIBCE, this information structure will in general have redundancies. But under this information structure, there is a pure strategy BNE σ' where each player sets his action equal to his "recommendation" (i.e., the action component of his signal): $\sigma'(a_i | a_i, s_i) = \mathbf{1}_{a_i}$ and so

$$\sum_{\theta, s} \Delta u_i(a, a'_i, \theta) \prod_i \sigma'_i(a_i | a_i, s_i) \sigma(a | s, \theta) Q(s_{-i}, \theta | s_i) = \sum_{\theta, s} \Delta u_i(a, a'_i, \theta) \sigma(a | s, \theta) Q(s_{-i}, \theta | s_i) \geq 0.$$

We now apply the same perturbation as in the proof of Proposition 6: Let $\Omega_{i, Q \circ \sigma}(a_i, s_i) := \{(\hat{a}_i, \hat{s}_i) : P \circ \sigma(\hat{a}_i, \hat{s}_i) > 0, \bar{\tau}_i^1(\hat{s}_i) = \bar{\tau}_i^1(s_i)\}$. Let $\rho_{i, \varepsilon} : \bar{S}_i \times \Theta \rightarrow \Delta(\{0, 1\})$ have the property that $\rho_{i, \varepsilon}(\hat{a}_i, \hat{s}_i, \theta) = \rho_{i, \varepsilon}(\hat{a}_i, \hat{s}_i, \theta')$ for all $\theta, \theta' \in \Theta$ and all \hat{a}_i, \hat{s}_i in the support of $Q \circ \sigma$ satisfying $|\Omega_{i, P \circ \sigma}(\hat{a}_i, \hat{s}_i)| = 1$. For any \hat{a}_i, \hat{s}_i in the support of $Q \circ \sigma$ satisfying $|\Omega_{i, Q \circ \sigma}(\hat{a}_i, \hat{s}_i)| > 1$ let $\|\rho_{i, \varepsilon}(\hat{a}_i, \hat{s}_i, \theta) - \rho_{i, \varepsilon}(\hat{a}_i, \hat{s}_i, \theta')\|_2 < \varepsilon$ so that for all distinct $(\tilde{a}_i, \tilde{s}_i), (\bar{a}_i, \bar{s}_i) \in \Omega_{i, P_\delta}(a_i, s_i)$ and $x_i, x'_i \in \{0, 1\}$,

$$\frac{\rho_{i, \varepsilon}(x_i | \tilde{a}_i, \tilde{s}_i, \theta) Q \circ \sigma(\theta | \tilde{a}_i, \tilde{s}_i)}{\sum_{\hat{\theta}} \rho_{i, \varepsilon}(x_i | \tilde{a}_i, \tilde{s}_i, \hat{\theta}) Q \circ \sigma(\hat{\theta} | \tilde{a}_i, \tilde{s}_i)} \neq \frac{\rho_{i, \varepsilon}(x'_i | \bar{a}_i, \bar{s}_i, \theta) Q \circ \sigma(\theta | \bar{a}_i, \bar{s}_i)}{\sum_{\hat{\theta}} \rho_{i, \varepsilon}(x'_i | \bar{a}_i, \bar{s}_i, \hat{\theta}) Q \circ \sigma(\hat{\theta} | \bar{a}_i, \bar{s}_i)}.$$

The prior $\hat{Q}(a, s, \theta, x) = Q \circ \sigma(a, s, \theta) \prod_i \rho_{i, \varepsilon}(x_i | a, s, \theta)$ induces (non-redundant) prior $Q' \in \mathcal{P}$ so that $d^{ACK}(P, P') < \varepsilon$ and σ' induces an outcome equivalent ε -BNE on P' . \square

Proposition 11 shows that that any correlation device required for a BIBCE can be embedded in the universal state space at the expense of ε -slack in the obedience constraint. Notice that under complete information (i.e., if there was a single state), it would not be possible to do so. Now the denseness of simple information structures implies the immediate corollary that this is true for all general information structures.

Corollary 1. *Let $|\Theta| \geq 2$. For any information structure $P \in \mathcal{P}$, any BIBCE $\sigma \in \mathcal{B}(\mathcal{G}, P)$ and any $\varepsilon > 0$, there exists (i) a simple information structure $P' \in \mathcal{P}^*$ such that $d^*(P, P') \leq \varepsilon$; and (ii) a decision rule σ' such that (a) $\sigma' \in \mathcal{B}^{\varepsilon BNE}(\mathcal{G}, P')$ and (b) the outcome induced by $P' \circ \sigma'$ is ε -close to the outcome induced by $P \circ \sigma$.*

Proof. The denseness of simple information structures implies that there exists a simple information structure P'' with $d^{ACK}(P, P'') \leq \varepsilon$. Now the corollary follows from applying Proposition 11 to P'' . \square

Corollary 1 implies that if we extend the notion of approximate BNE to allow not only slack in the obedience constraints but also to allow nearby (in the ACK topology) information structures, we can first establish the existence of approximate BNE and then establish continuity of (approximate) BNE outcomes with respect to the ACK topology.

We first define an extended notion of ε -BNE outcomes:

$$\mathcal{O}_\varepsilon^{BNE}(\mathcal{G}, P) = \left\{ \nu \in \Delta(A \times \Theta) : \begin{array}{l} \exists P' \text{ with } d^{ACK}(P, P') \leq \varepsilon, \\ \sigma \in \mathcal{B}^{\varepsilon BNE}(\mathcal{G}, P) \text{ s.t. } \|\nu - \nu_\sigma\| < \varepsilon \end{array} \right\}$$

Now we have existence:

Corollary 2. *For every game \mathcal{G} , prior $P \in \mathcal{P}$ and $\varepsilon > 0$, $\mathcal{O}_\varepsilon^{BNE}(\mathcal{G}, P) \neq \emptyset$.*

Proof. This result follows from Corollary 1 and the fact that $\mathcal{O}^{BNE}(\mathcal{G}, P) \neq \emptyset$ for all finite $P \in \mathcal{P}$. \square

We now introduce a “richness” property of a base game.

Definition 16. (*Richness*) *A base game \mathcal{G} is rich if for every action profile $a \in A$, there exists $\theta_a \in \Theta$ such that for every player i , every $a_i \in A_i$ and every $a'_i \neq a_i$,*

$$u_i(a_i, a_{-i}, \theta_a) - u_i(a'_i, a_{-i}, \theta_a) > 0.$$

With richness, we have continuity of approximate BNE outcomes:

Proposition 12. *Let $|\Theta| \geq 2$. Then for every rich base game \mathcal{G} and any information structure $P \in \mathcal{P}$,*

$$\lim_{\varepsilon \downarrow 0} \bigcup_{\hat{P} \in \mathcal{P}^*: d^{ACK}(P, \hat{P}) \leq \varepsilon} \mathcal{O}_\varepsilon^{BNE}(\mathcal{G}, \hat{P}) = \mathcal{O}(\mathcal{G}, P).$$

Proof. Let $\hat{P} \in \mathcal{P}^*$ satisfy $d^{ACK}(P, \hat{P}) \leq \varepsilon$. Then by Corollary 1, for every $\nu \in \mathcal{O}(\mathcal{G}, P)$ there exists a pure strategy ε -BNE, σ , of (\mathcal{G}, \hat{P}) , so that $\|\nu_\sigma - \nu\|_2 \leq \varepsilon$. For every τ denote the

associated action recommendation by $\alpha(\tau, \theta) = (\alpha_i(\tau_i))_i$, where $\sigma(\alpha(\tau, \theta)|\tau, \theta) = 1$. For any choice $\delta > 0$, define the stochastic map $\rho_\delta: \mathcal{T} \times \Theta \rightarrow \Delta(A)$

$$\rho_\delta(a|\tau, \theta) = \begin{cases} 1 - \delta & \text{if } a = \alpha(\tau, \theta) \\ \delta & \text{if } \theta = \theta_a \\ 0 & \text{otherwise.} \end{cases}$$

For every $(a, \tau, \theta) \in A \times \mathcal{T} \times \Theta$, let $P_\delta(a, \tau, \theta) := \hat{P}(\tau, \theta)\rho_\delta(a|\tau, \theta)$ and note that P_δ has a non-redundant representation $\hat{P}_\delta \in \mathcal{P}$. Let $J_{\mathcal{G}} := \min_{i, a, a'} u_i(a_i, a_{-i}, \theta_a) - u_i(a'_i, a_{-i}, \theta_a)$. Then for every i and type τ_i in the support of \hat{P} , and any deviation a'_i ,

$$\begin{aligned} \sum_{\theta, \tau, a} \Delta u_i(\alpha_i(\tau_i), \alpha_{-i}(\tau_{-i}), a'_i, \theta) \hat{P}_\delta(\tau_{-i}, \theta|\tau_i) &> -\varepsilon(1 - \delta) + \delta \sum_{\theta, \tau, a} \Delta u_i(\alpha_i(\tau_i), \alpha_{-i}(\tau_{-i}), a'_i, \theta_{\alpha(\theta, \tau)}) \hat{P}_\delta(\tau_{-i}, \theta|\tau_i) \\ &\geq -\varepsilon(1 - \delta) + \delta J_{\mathcal{G}} \end{aligned}$$

Letting $\delta = \frac{\varepsilon}{J_{\mathcal{G}} + \varepsilon}$ implies that σ is a ε -BNE under $\mathcal{G}, \hat{P}_\delta$. Moreover, $d^{ACK}(\hat{P}, \hat{P}_\delta) \leq \delta$ and so

$$\lim_{\varepsilon \downarrow 0} \bigcup_{\hat{P} \in \mathcal{P}^*: d^{ACK}(P, \hat{P}) \leq \varepsilon} \mathcal{O}_\varepsilon^{BNE}(\mathcal{G}, \hat{P}) \supseteq \mathcal{O}(\mathcal{G}, P).$$

By the property of upper hemi-continuity established in Proposition 19 in the appendix we also have that

$$\lim_{\varepsilon \downarrow 0} \bigcup_{\hat{P} \in \mathcal{P}^*: d^{ACK}(P, \hat{P}) \leq \varepsilon} \mathcal{O}_\varepsilon^{BNE}(\mathcal{G}, \hat{P}) \subseteq \lim_{\varepsilon \downarrow 0} \bigcup_{\hat{P} \in \mathcal{P}^*: d^{ACK}(P, \hat{P}) \leq \varepsilon} \mathcal{O}(\mathcal{G}, \hat{P}) \subseteq \mathcal{O}(\mathcal{G}, P),$$

and so the result follows. \square

4.4.6 Discussion

We conclude this section by discussing related literature about correlating devices and BNE.

A number of papers have highlighted the importance of redundant types, or correlating devices, for BNE, see for example Liu (2009) and Sadzik (2019). Our approach in this section is to observe that such correlation devices can be embedded in the universal state space in an almost payoff-irrelevant way, so it is natural to work with (non-redundant) information structures even if one is interested in BNE. Ely and Peski (2011) propose an alternative to the standard universal type space that embeds some correlation devices.

In the context of complete information games, Brandenburger and Friedenberg (2008) (see also Du (2012)) asked if correlation devices (supporting correlated equilibrium) could reflect higher-order strategic uncertainty (in this case, they said there is intrinsic correlation) or if extrinsic correlation is required. Their answer is that most correlated equilibria can be explained by intrinsic correlation alone. An analogous question (in an incomplete information context) is to ask is which BIBCE could reflect higher-order uncertainty about strategic uncertainty and payoffs. The spirit of our results is that most BIBCE can be justified this way.

Gossner (2000b) provided a partial order on correlating devices (for complete information games) capturing which correlating devices would support a larger set of correlated equilibria in all games. A natural exercise would be to define a topology on correlating devices (generating continuity of the set of correlated equilibria) although as far as we know that has not been done. We could imagine decomposing a general information structure into a non-redundant information structure and a correlating device and defining a topology on information structure / correlating device pairs that was sufficient for continuity of BNE outcomes. We have not pursued this approach.

4.5 Information Design

When studying information design problems, there will typically be many equilibria. In formulating information design problems, one must decide which equilibrium will be played. Two standard choices are to assume (i) the best equilibrium for the designer is played; or (ii) the worst equilibrium for the designer is played. In this section, we propose a formulation of information design problems that includes both those cases, but also allows for any continuous selection of equilibrium.

We will consider the following class of information design problems. A designer's evaluates outcomes with a utility function $u: A \times \Theta \rightarrow \mathbb{R}$ and takes as given an arbitrary continuous selection from BIBCE, $f: 2^{\Delta(A \times \Theta)} \setminus \emptyset \rightarrow \Delta(A \times \Theta)$ with the following property: for every $O' \subseteq O \subseteq \Delta(A \times \Theta)$,

$$f(O) \in O' \implies f(O') = f(O). \quad (4.3)$$

The designer then has the following objective function

$$V(O) = \sum_{a, \theta} f(a, \theta | O) u(a, \theta).$$

Recall that $\mathcal{P}^* \subseteq \mathcal{P}$ denotes the collection of simple (i.e., finite and first-order) information structures. Now we have:

Proposition 13. *For any rich \mathcal{G} , any open set $\mathcal{P}' \subseteq \mathcal{P}$ and any set \mathcal{S} of general information structures satisfying $\{P_{\mathcal{S}} : \mathcal{S} \in \mathcal{S}\} = \mathcal{P}'$,*

$$\sup_{\mathcal{S} \in \mathcal{S}} V(\mathcal{O}^{BNE}(\mathcal{G}, \mathcal{S})) = \sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}^{BNE}(\mathcal{G}, P)) = \sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}(\mathcal{G}, P)).$$

Thus to choose the optimal information structure within an open set, it is enough to focus on either BNE or BIBCE with simple information structures.

We will first describe how the information design problems described above fit within this class, and then describe how the result follows from denseness reported in Section 4.3.4 and the BNE results reported in Section 4.4.

4.5.1 Applications

Our assumption is that the designer cares about a set of outcomes. We are thinking that the designer cares about the outcomes consistent with a solution concept. This sub-section spells out the leading examples of designer objectives where the designer is interested in the best, or the worst, or some other continuous selection from the equilibrium outcomes.

The usual approach in information design is to assume that the designer can choose which equilibrium is played. In this case, if $O \subseteq \Delta(A \times \Theta)$ is the set of equilibrium outcomes, the designer's utility will be:

$$V_{MAX}(O) = \sup_{\nu \in O} \sum_{a, \theta} \nu(a, \theta) u(a, \theta).$$

This objective is continuous and Proposition 13 applies. However, the revelation principle applies to this problem, so we already know that we can restrict attention to finite information structures, without appeal to Proposition 13 and the machinery behind it.

An alternative assumption in information design is that there is “adversarial equilibrium selection,” i.e., the designer expects the worst possible equilibrium (for her) to be played. In this case the designer's utility over sets of outcomes will be

$$V_{MIN}(O) = \inf_{\nu \in O} \sum_{a, \theta} \nu(a, \theta) u(a, \theta).$$

A few papers have studied this problem in recent years (Mathevet et al. (2020b), Inostroza and Pavan (2023), Morris et al. (2024) and Li et al. (2023)). It is well known that the maximum is typically not attained in this design problem. However, while it is clear in the particular settings of these papers that the supremum can be approached using simple information structures, there is no existing general statement of this property. Thus Proposition 13 is a useful tool for this literature.

4.5.2 Proof of Proposition 13

Our denseness result implies:

Lemma 3. *For every rich game \mathcal{G} and every information structure P , there exists $\varepsilon > 0$ and a simple information structure $P^\varepsilon \in \mathcal{P}^*$ so that $d^{ACK}(P, P^\varepsilon) < \varepsilon$ and*

$$|V(\mathcal{O}(P, \mathcal{G})) - V(\mathcal{O}^{BNE}(P^\varepsilon, \mathcal{G}))| < \varepsilon.$$

Proof. Suppose $\sigma \in \mathcal{B}(\mathcal{G}, P)$ satisfies $V(\mathcal{O}(\mathcal{G}, P)) = \sum_{a, \theta} \nu_\sigma(a, \theta) u(a, \theta)$. Then there are sequences $(P^k, \varepsilon^k)_k$ with $P^k \in \mathcal{P}^*$ for all k and $\varepsilon^k \downarrow 0$ so that $d^{ACK}(P, P^k) < \varepsilon^k$ for all k . Moreover, by upper hemi-continuity established in Proposition 19, there is a subset $\mathcal{O}^\infty(\mathcal{G}, P) \subseteq \mathcal{O}(P, \mathcal{G})$ so that the sequence satisfies $\lim_{k \uparrow \infty} d_{\mathcal{H}}(\mathcal{O}^{BNE}(\mathcal{G}, P^k), \mathcal{O}^\infty(\mathcal{G}, P)) = 0$, where $d_{\mathcal{H}}$ denotes the Hausdorff distance. Moreover, by the arguments in Proposition 12 we can pick the sequence so that $\nu_\sigma \in \mathcal{O}^\infty(\mathcal{G}, P)$ and so by property (4.3) and continuity of V we have that $\lim_{k \uparrow \infty} V(\mathcal{O}^{BNE}(P^k, \mathcal{G})) = V(\mathcal{O}(P, \mathcal{G}))$ and so the result follows. \square

Now the proof of Proposition 13 is completed as follows. By Lemma 3, we have that

$$\sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}^{BNE}(\mathcal{G}, P)) \geq \sup_{P \in \mathcal{P}'} V(\mathcal{O}(\mathcal{G}, P)).$$

Since $\mathcal{P}^* \cap \mathcal{P}' \subseteq \mathcal{P}'$, we have that

$$\sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}^{BNE}(\mathcal{G}, P)) \leq \sup_{P \in \mathcal{P}'} V(\mathcal{O}^{BNE}(\mathcal{G}, P)).$$

Moreover, by property (4.3) and the fact that $\mathcal{O}^{BNE}(\mathcal{G}, P) \subseteq \mathcal{O}(\mathcal{G}, P)$, we have that

$$\sup_{P \in \mathcal{P}'} V(\mathcal{O}^{BNE}(\mathcal{G}, P)) \leq \sup_{P \in \mathcal{P}'} V(\mathcal{O}(\mathcal{G}, P))$$

and so

$$\begin{aligned} \sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}^{BNE}(\mathcal{G}, P)) &\leq \sup_{P \in \mathcal{P}'} V(\mathcal{O}(\mathcal{G}, P)) \\ &= \sup_{P \in \mathcal{P}' \cap \mathcal{P}^*} V(\mathcal{O}(\mathcal{G}, P)), \end{aligned}$$

where the last line follows from denseness.

4.6 Alternative Formulations and Related Literature

In this section, we will discuss a number of alternative topologies characterizing convergence of strategic outcomes. One purpose in doing so is that it will allow us to formally relate our work to the relevant related literatures.

4.6.1 Interim Topologies

We have defined and characterized an (ex ante) strategic topology on (common prior) information structures under an equilibrium solution concept (BIBCE). By contrast, Dekel et al. (2006) defined an interim strategic topology on hierarchies of beliefs under the solution concept of interim correlated rationalizability (ICR). Two belief hierarchies were said to be close in the interim strategic topology if, in any game, an action that was ICR at one hierarchy was approximately ICR at the other hierarchy. Chen et al. (2017) provide a characterization of the interim strategic topology in terms of belief hierarchies. Crucially, the interim strategic topology imposes restrictions on the tails of hierarchies of beliefs, unlike the product topology. We provide a formal statement of the characterization of Chen et al. (2017) in the Appendix.

Our definition of the almost common knowledge topology used the product topology in defining the event that interim beliefs were close. But we noted that the use of the product topology was not essential. A first purpose of this section is to record what properties the interim topology must satisfy in order to induce our almost common knowledge topology: it is enough that it is induced by a “nice” metric that (1) refines the product topolog while preserving the Borel sigma algebra of the product topology and (2) induces a locally compact, complete, separable metric space.

Naturally, the product topology satisfies these properties. But we also establish that if we had used the interim strategic topology as our initial interim topology, we could have

dispensed with the requirement of approximate common knowledge in our definition of the ex ante topology and induced the same strategic topology on information structures. Intuitively, this is because interim strategic topology is induced by a “nice” metric and (unlike the product topology) imposes restrictions on infinite tails of hierarchies of beliefs which has implications for approximate common knowledge. In particular, the interim strategic topology has the following property.

Definition 17. (*Common Belief Invariance*) A metric d_{CB} on Ω satisfies common belief invariance if for all events $E, E' \subseteq \Omega$ and every $\varepsilon > 0$,

$$C^{1-\varepsilon}(\mathcal{N}_{d_{CB},\varepsilon}(E) \cap \mathcal{N}_{d_{CB},\varepsilon}(E')) = \mathcal{N}_{d_{CB},\varepsilon}(E) \cap \mathcal{N}_{d_{CB},\varepsilon}(E'),$$

where $\mathcal{N}_{d_{CB},\varepsilon}(E)$ is the union of ε -neighborhoods around the points in E using metric d_{CB} .

For any metric $d : \Omega \times \Omega \rightarrow [0, \infty)$, $\varepsilon > 0$ and $P \in \mathcal{P}$, let $\text{supp}_{d,\varepsilon}(P) := \bigcup_{\omega \in \Omega: P(\mathcal{N}_{d,\varepsilon}(\omega)) > 0} \mathcal{N}_{d,\varepsilon}(\omega)$. Now we can consider the simplest and weakest natural way of translating an interim distance into an ex ante distance

Definition 18. The weak ex-ante distance induced by an interim distance d is defined as

$$d'(P, P') = \inf \left\{ \varepsilon \geq 0 : \begin{array}{l} P(\hat{T}_{d,\varepsilon}(P, P')) > 1 - \varepsilon \\ P'(\hat{T}_{d,\varepsilon}(P, P')) > 1 - \varepsilon \end{array} \right\},$$

where $\hat{T}_{d,\varepsilon}(P, P') := \text{supp}_{d,\varepsilon}(P) \cap \text{supp}_{d,\varepsilon}(P')$.

Definition 19. (*Nice Interim Metric*) A metric d on Ω is nice if (1) its induced topology refines the product topology while preserving the Borel sigma algebra of the product topology \mathcal{B} ; (2) the metric space (d, Ω) is complete, separable, locally compact and Hausdorff.

A topology on Ω is nice if it is induced by a nice metric. Our ACK topology would be the same if we replace the product topology with any nice interim topology in the definition that satisfies common belief invariance. We used the product topology. We could also have substituted the interim strategic topology as defined by [Dekel et al. \(2006\)](#) and characterized by [Chen et al. \(2017\)](#).

Proposition 14. The weak ex ante distance induced by any nice interim metric satisfying common belief invariance induces the ACK topology.

Proof. Appendix. □

Here, we have the special property that approximate common knowledge is for free.

Proposition 15. The interim strategic topology is nice and common belief invariant.

Proof. Follows from Propositions [18](#) and [17](#) in the Appendix. □

For example, [Kajii and Morris \(1998\)](#) define a topology on ex ante information structures (discussed below) but use total variation as a metric on interim beliefs.

4.6.2 Value-Based Topology

We could alternatively define our topology in terms of convergence of the ex ante expected equilibrium payoffs rather than equilibrium outcomes. This was the approach of [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#) and also the recent work of [Gensbittel et al. \(2022\)](#) for zero sum games. This distinction is not important for our results.

Let $V(G, P)$ be the set of ex ante utilities of players (in \mathbb{R}^I) that can arise from some BIBCE of (G, P) , and let $V_\varepsilon(G, P)$ be the set of ex ante utilities of players (in \mathbb{R}^I) that are within ε of some ε -BIBCE of (G, P) . We can say that P, P' are ε -value close in game \mathcal{G} if $V(\mathcal{G}, P) \subseteq V_\varepsilon(G, P')$ and $V(\mathcal{G}, P') \subseteq V_\varepsilon(G, P)$.

Definition 20. Let $d^V(P, P'|\mathcal{G})$ be the infimum of the set of ε such that P and P' are ε -value close in game \mathcal{G} .

Lemma 4. Now $d^{ACK}(P^k, P|\mathcal{G}) \rightarrow P$ if and only if $d^V(P^k, P|\mathcal{G}) \rightarrow 0$ for all \mathcal{G} .

Proof. One direction is immediate, because convergence of outcomes implies convergence of values. In the other direction, it is enough to change payoffs so differences in outcomes translate into large differences in payoffs. Consider two outcomes ν, ν' of P and P' respectively so that

$$V_i(\nu) := \sum_{a, \theta} \nu(a, \theta) u_i(a, \theta) = \sum_{a, \theta} \nu'(a, \theta) u_i(a, \theta) =: V_i(\nu')$$

but $\nu \neq \nu'$. Consider the augmented payoffs $u_i(a, \theta) + h_i(a_{-i}, \theta)$ and associated values $V_i^h(\nu) = V_i(\nu) + \sum_{a, \theta} \nu(a, \theta) h_i(a_{-i}, \theta)$ and $V_i^h(\nu') = V_i(\nu') + \sum_{a, \theta} \nu'(a, \theta) h_i(a_{-i}, \theta)$. Hence

$$V_i^h(\nu) - V_i^h(\nu') = \sum_{a, \theta} h_i(a_{-i}, \theta) (\nu(a, \theta) - \nu'(a, \theta)).$$

We now argue that there exists a choice of h_i so that $V_i^h(\nu) \neq V_i^h(\nu')$. Indeed, let $\zeta_i(a_{-i}, \theta) := \sum_{a_i} \nu(a_i, a_{-i}, \theta) - \nu'(a_i, a_{-i}, \theta)$ and consider

$$h_i(a_{-i}, \theta) := \begin{cases} +1 & \text{if } \zeta_i(a_{-i}, \theta) > 0 \\ 0 & \text{if } \zeta_i(a_{-i}, \theta) = 0 \\ -1 & \text{if } \zeta_i(a_{-i}, \theta) < 0. \end{cases}$$

Then we have that

$$V_i^h(\nu) - V_i^h(\nu') = \sum_{a_{-i}, \theta} h_i(a_{-i}, \theta) \zeta_i(a_{-i}, \theta) > 0.$$

□

Thus there is little difference working with outcome-based strategic topologies and value-based strategic topologies.

However, in the case of zero-sum games, the value is uniquely defined although many outcomes might give rise to the same value. So it is convenient and natural to work with value-based strategic topologies in that case. [Peski \(2008\)](#) and [Gossner and Mertens \(2020\)](#) characterize changes in information structure that increase one player's payoff in all zero sum games. [Gensbittel et al. \(2022\)](#) study essentially the value-based strategic topology defined above but restricted to zero-sum games.

4.6.3 Join Measurability

Our approach in this paper has been to fix a set of payoff-relevant states Θ , and look at common prior information structures that describe beliefs and higher-order beliefs about those states. Then we characterize the coarsest topology under which equilibrium outcomes converge for any game where payoffs depend on Θ . In particular, we do not allow games to depend in an arbitrary way on players' types (or signals).

An alternative approach would be to allow any game where payoffs were measurable with respect to the join of players' types. Equivalently, we could restrict attention to information structures where each payoff-relevant state could arise under only one profile of types; we call this a "join measurability" restriction on information structures. It was implicitly maintained in the works of [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#). This restriction greatly simplifies the arguments. In particular, in the proofs of sufficiency analogous to Proposition 4 the join measurability approach allows a straightforward mapping of a strategy profile on one information structure to another. We were not able to do that, and required a continuous extension exploiting the structure of the universal type space. In the proofs of necessity analogous to Proposition 5, the join measurability approach requires only an email game like component where an infection argument operates and not also an iterated scoring rule to reveal finite levels of beliefs, as in this paper.

This is the most important difference between the work of [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#), and this work. There are number of other differences that are less important. First, the earlier papers focused on BNE as a solution concept, while we focus on BIBCE. Second, they focused on countable information structures (to ensure existence of BNE), whereas we do not impose that restriction. Third, their topologies were value-based whereas our topology is outcome-based (a difference that we argued was not important in the previous section) . Fourth, we restrict ourselves to a finite set of payoff-relevant states, but these papers must allow for countable payoff relevant states.

The set of information structures considered in [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1998\)](#), while both satisfying join measurability, were modelled differently. [Monderer and Samet \(1996\)](#) fixed a state space and prior probability. An information structure was then a profile of (countable) partitions of the state space. And payoffs could depend in arbitrary ways on the state space. On the other hand, [Kajii and Morris \(1998\)](#) fixed a countable set of "types" (or labels) for each player and allowed arbitrary probability distributions on the types space. The exact connection between the similar topologies defined on different classes of information structures was not known until recently, when [Kambhampati \(2023\)](#) showed an equivalence between the results.

4.6.4 Improper Priors and Completeness

In this paper, we have focused on common prior information structures. Our results imply that equilibrium outcomes converge along Cauchy sequences in the ACK topology. However, in general, Cauchy sequences may not have well-defined limits within the space of information structures. In this section, we show that if we enrich the class of information structures to include "improper" common prior information structures, and extend the ACK topology to this class of information structures, then all Cauchy sequences do converge to a well defined

limit. This result is of independent interest, in the light of the importance of improper common prior limits in the literature on higher-order beliefs in games (discussed below).

An improper prior on the universal type space is simply a measure with perhaps infinite mass such that there is a conditional probability consistent with the interim beliefs on the universal type space. Formally, we have:

Definition 21. (*Improper Prior*) A measure on Ω , $Q: \mathcal{B} \rightarrow [0, \infty]$, is an improper prior if for every player i there is a conditional measure of Q , $Q_i: \mathcal{T}_i \rightarrow \Delta(\Omega)$ so that

$$\tau_i^* = \text{marg}_{\mathcal{T}_i \times \Theta} (Q_i(\tau_i)), \quad Q\text{-a.e.}$$

Let $\bar{\mathcal{P}}$ denote the set of improper priors and note that $\mathcal{P} \subseteq \bar{\mathcal{P}}$. The approximate common knowledge topology extends in a natural way to $\bar{\mathcal{P}}$:

Definition 22. (*Approximate Common Knowledge Distance*) For every $P, P' \in \bar{\mathcal{P}}$, let

$$\bar{d}^{ACK}(P, P') := \inf \left\{ \varepsilon \geq 0 : \begin{array}{l} P \left(\Omega \setminus C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) < \varepsilon P \left(C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) \\ P' \left(\Omega \setminus C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) < \varepsilon P' \left(C^{1-\varepsilon} \left(\hat{T}_\varepsilon(P, P') \right) \right) \end{array} \right\}.$$

Notice that this Definition coincides with the earlier definition 8 when applied to proper priors.

A decision rule σ is a BIBCE of an improper prior $Q \in \bar{\mathcal{P}}$, if σ is belief invariant and obedience holds almost everywhere. For any $\varepsilon > 0$, let the collection of ε -BIBCE on an improper prior Q be denoted by $\bar{\mathcal{B}}^\varepsilon(\mathcal{G}, Q)$.

Proposition 16. *The extended approximate common knowledge topology on $\bar{\mathcal{P}}$ is complete.*

Proof. Let $\mathcal{S} \subseteq \mathcal{B}$ be a semi-ring so that the sigma algebra it generates equals \mathcal{B} . Let $(P^k)_k$ be a Cauchy sequence in \mathcal{P} and $(\varepsilon_k)_k$ a sequence so that for every k ,

$$\varepsilon_k \leq \sup_{h>k} \bar{d}^{ACK}(P^k, P^h).$$

Define a limit pre-measure on \mathcal{S} , which on any event $E \in \mathcal{S}$ takes the form

$$\xi_0(E) := \lim_{k \uparrow \infty} P^k(\mathcal{N}_{\varepsilon_k}(E)),$$

where we set $\mathcal{N}_\varepsilon(\emptyset) = \emptyset$. We now show that ξ_0 is a pre-measure on \mathcal{S} . Indeed, $\xi_0(\emptyset) = 0$. Since each P^k is a measure, they are all finitely additive and so for any finite, disjoint $(E_n)_{n \leq N}$, the limit and sum can be exchanged so that

$$\begin{aligned} \xi_0 \left(\bigcup_{n \leq N} E_n \right) &= \lim_{k \uparrow \infty} \sum_{n \leq N} P^k(\mathcal{N}_{\varepsilon_k}(E_n)) \\ &= \sum_{n \leq N} \xi_0(E_n) \end{aligned}$$

Finally, countable monotonicity follows from Reverse Fatou's lemma (each P^k is bounded and non-negative):

$$\begin{aligned}\xi_0 \left(\bigcup_{n \in \mathbb{N}} E_n \right) &= \lim_{k \uparrow \infty} \sum_{n \in \mathbb{N}} P^k(\mathcal{N}_{\varepsilon^k}(E_n)) \\ &\leq \sum_{n \in \mathbb{N}} \xi_0(E_n).\end{aligned}$$

By Caratheodory-Hahn's Extension Theorem, ξ_0 extends to a measure ξ on the sigma algebra generated by \mathcal{S} . It remains to show that ξ_0 is a non-redundant improper prior. We need to prove that ξ_0 satisfies the consistency condition. For every event $E \in \mathcal{B}$, we have that

$$\begin{aligned}|P^k(E) - \xi(E)| &= \left| \int \tau_i(\mathcal{N}_{\varepsilon^k}(E)) P^k(d\tau_i) - \lim_{k \uparrow \infty} \int \tau_i(\mathcal{N}_{\varepsilon^k}(E)) P^k(d\tau_i) \right| \\ &\leq \varepsilon_k.\end{aligned}$$

Hence $\xi_0 \in \overline{\mathcal{P}}$. Deduce that $\overline{\mathcal{P}}$ is a complete metric space. □

Chapter 5

Information Aggregation Mechanisms

5.1 Introduction

In this short note, we describe an information aggregation mechanism that can be used by players before playing a game of strategic complementarities under incomplete information. In such a game, players may have an incentive to share overly optimistic information with other players, thus inducing them to play higher actions. In this mechanism, players trade a token before playing the game. Players who want to communicate good news must purchase this worthless token and burn resources. The note shows that players only need to observe the market clearing price that arises from the token trades to aggregate their private information. Each element in a player's private information set is encoded as a prime in the prime factorization of the market clearing price. The element that is contained in every player's information set is identified as the prime with the highest multiplicity.

For early-stage start-ups, incomplete information about the future returns of the start-up is an important friction in the fundraising process. Different investors have different beliefs about the quality of the start-up, different beliefs about how investors will evaluate the quality of the start-up, and possibly even different beliefs about those beliefs and so on. Moreover, if the future profits of the project exhibit some form of strategic complementarities, agents have the incentive to misreport what they know: If an agent believes the project to be just weakly above average quality, she may exaggerate her beliefs about the quality of the project so that other investors over-invest into the project. Indeed, over-investment of other agents weakly raises her returns. We see that information aggregation in this context requires a more complicated mechanism, where optimistic reports that are meant to increase the aggregate level of investment become costly. The simplest information aggregation mechanism is a public message board, where posting information is free and agents take turns reporting what they know. This cannot be used to incentivize information aggregation. In this note we propose a token trading protocol that achieve this.

5.1.1 Model

Consider a set I of $n \in \mathbb{N}$ agents, where each $i \in I$ is endowed with a cash endowment $b_i \in \mathbb{R}$ that enters her preferences in a quasi-linear way. Agents decide how much to invest into a

project whose payoffs exhibit strategic complementarities. Returns of this investment game depend on an unknown quality parameter $\theta \in \Theta := \mathbb{Z}$. The investment game is described as follows: Each player chooses a level of investment $y_i \geq 0$. Given a profile of actions $y = (y_i)_i$ and a quality parameter $\theta \in \Theta$, agent i 's payoffs are given by

$$U(y, \theta, b_i) = u(y, \theta) + b_i. \quad (5.1)$$

where u is smooth, strictly concave and supermodular with bounded derivatives. We assume that U is such that for every $\theta, (b_i)_i$, the complete information game with payoffs $U(\cdot)$ and actions $y = (y_i)_i$ has a unique pure strategy Nash equilibrium where at least one player invests a positive amount, denoted $y^*(\theta, (b_j)_j) = (y_i^*(\theta, (b_j)_j))_i$. We also assume that U is such that y^* is strictly increasing and strictly concave with bounded derivatives. In addition, zero investment is always assumed to be a Nash equilibrium. We will later require that these functions be sufficiently concave.

Information Players' private information is given by an improper, uniform prior ϕ on Θ and private signals $s = (s_i)_i$, where for every player i , the signal s_i is obtained as follows:

$$s_i = \theta + \epsilon_i, \quad (5.2)$$

where for every i , the noise term ϵ_i is drawn independently of θ and independently across players from a uniform distribution $\nu_i \in \Delta(\mathbb{Z})$ with finite support and mean zero. The tuple of $(\nu_i)_i$ thus describes a Bayesian game, where players' beliefs about θ are uniformly distributed on some interval. We denote the associated posterior beliefs of player i about θ , conditional on observing signal s_i by $\mu_i(\cdot | s_i)$. The improper distribution on tuples $(\theta, (s_i)_i)$ is denoted by μ .

We say that $(\nu_i)_i$ satisfy the *aggregate consensus* property if the following condition holds: For every $(\theta, (s_i)_i)$ in the support of μ ,

$$\bigcap_{i \in I} \{\theta' \in \Theta : \mu_i(\theta' | s_i) > 0\} = \{\theta\}. \quad (5.3)$$

Suppose the noise term satisfies $\epsilon_i \in \{-4, \dots, 0, \dots, 4\}$. With some abuse of notation, we will denote the first-order beliefs of agent i upon observing s_i with $\mu(\cdot | s_i)$. The distribution on signals s_i conditional on θ for the three agents is summarized below:

		Signals given θ								
Signals:		$\theta - 4$	$\theta - 3$	$\theta - 2$	$\theta - 1$	θ	$\theta + 1$	$\theta + 2$	$\theta + 3$	$\theta + 4$
Signal distribution	i_1 :	0	0	0	$1/3$	$1/3$	$1/3$	0	0	0
Signal distribution	i_2 :	0	0	$1/3$	0	$1/3$	0	$1/3$	0	0
Signal distribution	i_3 :	$1/5$	$1/5$	0	0	$1/5$	0	0	$1/5$	$1/5$

Figure 5.1: Signal distributions for agents i_1, i_2 and i_3 conditional on the true state being θ .

So if agent i_1 observes signal s_1 , she knows $\theta \in \{s_1 - 1, s_1, s_1 + 1\}$. She also knows that if $\theta = s_1 - 1$, say, then i_2 ' signal is contained in $\{s_1 - 3, s_1 - 1, s_1 + 1\}$. If $s_2 = s_1 - 1$, then i_2 knows that $\theta \in \{s_1 - 3, s_1 - 1, s_1 + 1\}$. In this case, agents i_1 and i_2 would not learn the state if they were to reveal the support of their beliefs. Both agents can only narrow the state down to either $s_1 - 1$ or $s_1 + 1$. Note that if $\theta = s_1 - 1$, then i_3 knows that $\theta \neq s_1 + 1$ no matter which of her signals she observed. So all three agents can learn the state if they manage to credibly communicate the support of their first-order beliefs. It can be verified that this is true for any state and any draw of the signals. So the beliefs described above satisfy the consensus condition.

Information Aggregation Mechanism Suppose that agents participate in a mechanism after they received their private signals but before they have to make their investment choices. The mechanism considered in this note is a special kind of mechanism. Suppose all agents are endowed with a large number of worthless tokens. An information aggregation mechanism consists of a token market. Every agent i reports a demand or supply schedule d_i for tokens $p \mapsto d_i(p)$, where $d_i(\cdot)$ takes as argument a positive price per token and maps it to a number of tokens. If the amount is positive, $d_i(p)$ represents a demand for tokens. If the amount is negative it represents a supply of tokens. For every profile of demand/supply schedules, let the public signal be given by a market clearing price¹, i.e. a number $p^* > 0$ so that

$$\sum_{i \in I} d_i(p^*) = 0. \quad (5.4)$$

At every round, agents effectively submit demand/supply schedules for tokens and observe the resulting market price of a token. Then agents pay or receive payment for tokens in units of endowment as dictated by their submitted demand/supply schedules and the market price. Hence transfers given $d = (d_i)_i, p$ take the form

$$\tau(d_i, d) := -p^* d_i(p^*). \quad (5.5)$$

¹When multiple such prices exist we consider any selection rule defined on the aggregate schedules $\sum_{i \in I} d_i(\cdot)$. If no market clearing price exists, set p to zero.

Agents choose their reported demand/supply schedule d_i as a function of their private signal s_i . A strategy in the aggregation mechanism is then denoted $s_i \mapsto d_i(\cdot|s_i)$. We say that information has been aggregated if for every $(\theta, s = (s_i)_i)$ in the support of μ ,

$$Pr(\theta|s_i, p^*) = 1, \forall i. \quad (5.6)$$

In the above expression, p^* is the value of the market clearing price and strategies are common knowledge. Note that no agent needs to condition on their own strategy $d_i(\cdot|s_i)$ to deduce the true state, the price and her private signal are enough.

Agents choose their strategy in the extended, two-period game, which consists of a strategy in the aggregation mechanism $s_i \mapsto d_i(\cdot|s_i)$ and an investment strategy $(s_i, p^*) \mapsto y_i(s_i, p^*) \geq 0$. Agent i 's expected payoff before the trades, given a strategy $(d_i, y_i)_i$, is given by

$$U_i(s_i, y, d) := \mathbb{E}(u(y(s))\theta + b_i - p^*d_i(p^*|s_i)|s_i) \quad (5.7)$$

and expected payoffs after trade, but before the investment game:

$$U_i^*(s_i, y, d, p^*) := \mathbb{E}(u(y(s))\theta + b_i - p^*d_i(p^*|s_i)|s_i, p^*). \quad (5.8)$$

A Bayes-Nash equilibrium of this game consists of a strategy profile $(d_i, y_i)_i$ so that for every player i and every signal s_i , (d_i, y_i) maximizes U_i^* above given (d_{-i}, y_{-i}) .

5.2 The Mechanism

We now provide an explicit strategy for every agent i which is a Bayes Nash equilibrium of the game and aggregates information. For every $m \in \mathbb{N}_{>0}$, let P_m denote the m -th prime number. That is $P_1 = 2, P_2 = 3, P_3 = 5$ etc... Fix any increasing function $f: \mathbb{N} \rightarrow \mathbb{N}$ and let $\rho_f: \Theta \rightarrow \mathbb{N}$ be defined as follows

$$\rho_f(\theta) := \begin{cases} P_1, & \text{if } \theta \leq 1 \\ P_{f(\theta)}, & \text{otherwise.} \end{cases} \quad (5.9)$$

We then propose the following strategy $(d_i, y_i)_i$ for any given $\alpha > 0$:

- (i) For every s_i , define the associated demand/supply schedule for player i under signal s_i

$$d_i(p|s_i) = \frac{\alpha n}{p} \sum_{\theta: \mu_i(\theta|s_i) > 0} \log(\rho_f(\theta)) - 1. \quad (5.10)$$

The market clearing price p^* associated to a signal profile $s = (s_i)_i$ satisfies

$$\begin{aligned} \frac{\alpha n}{p^*} \sum_{i \in I} \sum_{\theta: \mu_i(\theta|s_i) > 0} \log(\rho_f(\theta)) - n &= 0 \\ \alpha \sum_{i \in I} \sum_{\theta: \mu_i(\theta|s_i) > 0} \log(\rho_f(\theta)) &= p^* \end{aligned} \quad (5.11)$$

We now verify that $d = (d_i)_i$ aggregates information: Indeed, upon seeing p^* , each agent can perform the following computation:

$$\begin{aligned} \exp\left(\frac{p^*}{\alpha}\right) &= \exp\left(\sum_{i \in I} \sum_{\theta: \mu_i(\theta|s_i) > 0} \log(\rho_f(\theta))\right) \\ &= \prod_{\theta: \exists i \text{ s.t. } \mu_i(\theta|s_i) > 0} \rho_f(\theta)^{\eta(\theta)}, \end{aligned} \quad (5.12)$$

where for every $\theta \in \Theta$, $\eta(\theta)$ is the number of agents whose information set contains θ ,

$$\eta(\theta) := |\{i \in I : \mu_i(\theta|s_i) > 0\}|. \quad (5.13)$$

Assuming that every agent can compute the prime factorization of $\exp(p^*)$, the aggregate consensus property implies that for any drawn tuple $(\theta^*, s = (s_i)_i)$ we have that

$$\eta(\theta^*) > \eta(\theta), \quad \forall \theta \neq \theta^*. \quad (5.14)$$

Hence information is aggregated.

- (ii) Let $d = (d_i)_i$, as defined above, be played during the information aggregation mechanism stage, and let the associated market clearing price p^* reveal state θ^* . We then consider an investment strategy $y = (y_i: (s_i, p^*) \mapsto y_i(s_i, p^*) = y_i(p^*))_i$, which is independent of each player's private signal s_i and let y be the Nash equilibrium $y^*(\theta^*, (b_i - p^* d_i(p^*))_i)$ of the complete information game with payoffs

$$U(y, \theta^*, b_i - p^* d_i(p^*)). \quad (5.15)$$

Moreover, we assume that if a prime factorization that is inconsistent with the information structure is observed, each agent invests zero.

We now check that no agent has an incentive to deviate during the information aggregation stage. Suppose agent i observes signal s_i . Reporting a higher information set means buying more tokens and may increase the consensus state that results from the market clearing price. Conversely, reporting a lower information set is cheaper during the information aggregation stage but might result in a more pessimistic consensus.

A demand schedule $d'_i(\cdot)$ is more optimistic than $d_i(\cdot|s_i)$ if there exists $\Theta' \subseteq \Theta$ containing some $\theta > 2$, which is also greater than $\Theta(s_i) := \{\theta : \mu_i(\theta|s_i) > 0\}$ in the weak set order: 1) For every $\theta \in \Theta(s_i)$ there exists $\theta' \in \Theta'$ so that $\theta \leq \theta'$ and 2) For every $\theta' \in \Theta'$ there exists $\theta \in \Theta(s_i)$ so that $\theta \leq \theta'$. Symmetrically, a demand schedule $d'_i(\cdot)$ is more pessimistic than $d_i(\cdot|s_i)$ if there exists $\Theta' \subseteq \Theta$, so that $\Theta(s_i)$ is greater than Θ' in the weak set order.

Suppose first that an agent reports a demand schedule $d'_i(\cdot)$ that is strictly more optimistic than the prescribed demand schedule $d_i(\cdot|s_i)$. The additional cost is bounded from below by

$$p'^* d'_i(p'^*) - p^* d_i(p^*|s_i) > c_f(s_i) := \alpha n \min_{\theta \in \Theta(s_i)} (\log(\rho_f(\theta + 1)) - \log(\rho_f(\theta))), \quad (5.16)$$

where p^* is the price that results from the demand profile d and p'^* is the price resulting from player i 's deviation. The Bertrand-Chebyshev Theorem, [Tchebichef \(1852\)](#), states that for any

$l \in \mathbb{N}$ there is a prime number contained in $[l, 2l]$. Hence, we can find a strictly increasing function $f: \mathbb{N} \rightarrow \mathbb{N}$ so that for every $\theta \in \mathbb{N}$,

$$\log(\rho_f(\theta)) \in ((\theta - 1) \log(2), \theta \log(2)]. \quad (5.17)$$

Hence log primes can be used to span an approximately even grid. Hence, there is a choice of f so that

$$c_f(s_i) \geq \alpha n \log(2). \quad (5.18)$$

Let $\bar{\Theta}(s_i)$ denote the set of possible states that could arise as consensus with positive probability:

$$\bar{\Theta}(s_i) := \bigcup_{s_{-i}: \mu(s_i, s_{-i}) > 0} \bigcup_{j \neq i} \Theta_j(s_j). \quad (5.19)$$

The set $\bar{\Theta}(s_i)$ is finite and its cardinality is bounded uniformly by M for all signals s_i . Since agents invest zero if they see a prime factorization that is inconsistent with the information structure and information sets have finite size independent of the realized signals, no deviation encoding elements outside of $\bar{\Theta}(s_i)$ could ever be profitable. In the best case scenario, information aggregation after a deviation leads agents to the (possibly incorrect) consensus $\bar{\theta} := \max \bar{\Theta}(s_i)$. Let $\underline{\theta} := \min \Theta(s_i)$. Then the payoff gains from deviating to a more optimistic demand schedule under state θ^* is bounded from above by

$$\left(\max_{y_i \geq 0} u(y_i, y_{-i}^*(\bar{\theta}, (\bar{b}_j^* + \alpha \log(2))_{j \neq i}), \theta^*) - u(y^*(\underline{\theta}), (\underline{b}_j^*)_j, \theta^*) \right) - \alpha n \log(2), \quad (5.20)$$

where \bar{b}_j^* and \underline{b}_j^* are the highest and (resp.) lowest possible post-trade endowment for agent j given s_i if d is followed:

$$\begin{aligned} \bar{b}_j^* &:= \max_{s_{-i}: \mu(s_i, s_{-i}) > 0} (b_j - \tau(d_j(s_j), d_{-j}(s_{-j}))). \\ \underline{b}_j^* &:= \min_{s_{-i}: \mu(s_i, s_{-i}) > 0} (b_j - \tau(d_j(s_j), d_{-j}(s_{-j}))). \end{aligned} \quad (5.21)$$

By the bounds above we have that there exists M so that for all signals s_i ,

$$\begin{aligned} |\bar{b}_j^* - \underline{b}_j^*| &\leq M, \\ |\bar{\theta} - \underline{\theta}| &\leq M. \end{aligned} \quad (5.22)$$

We will thus assume that u and y^* be sufficiently concave with bounded derivatives, so that there exists a constant $L \in (0, 1)$ so that for all endowments b all $\theta \in \Theta$

$$|u(y^*(\theta + M, (b_j + M + \log(2))_j), b_i, \theta + M) - u(y^*(\theta, b), b_i, \theta)| \leq Ln \log(2). \quad (5.23)$$

We conclude that a deviation towards a more optimistic demand schedule is never profitable. Similarly, a symmetric argument shows that deviating to a less optimistic demand schedule is not profitable. Finally, by concavity, no mixture of a optimistic and a pessimistic deviation would ever be optimal.

5.3 Discussion

5.3.1 Natural Encryption

Note that the equilibrium described above naturally encrypts the state whenever factoring a number into its basis elements is costly. For an outside observer without any private information about the state, the factorization can be made arbitrarily costly when using the protocol we introduced earlier by varying f . Agents who participate in the mechanism know which primes they have used in the encoding. Under a factorization cost, this mechanism seems more suitable for information aggregation among well informed experts. That is agents who already know a lot but also care a lot about aggregating their information without revealing their information to ignorant outsiders.

5.3.2 Related Literature

Information aggregation has been studied mostly in the context of auctions and competitive markets to provide micro-foundations for rational expectation equilibria [Milgrom \(1981\)](#); [Reny and Perry \(2006\)](#); [Grossman and Stiglitz \(1976\)](#); [Siga and Mihm \(2021\)](#). By contrast, in this paper, we view information aggregation from a mechanism design perspective. We want to design a mechanism whose sole purpose is to eliminate or reduce the friction of incomplete information when playing an investment game. Our mechanism should ideally be an add-on to investment projects in a broad class of different environments.

There is also a literature [Kyle \(1985\)](#); [Vives \(2014\)](#); [Lambert et al. \(2014\)](#) on information aggregation where traders submit monotone demand (or supply) schedules. This strand is close in spirit to our paper as it studies markets with strategic traders who receive private information and a market maker who determines prices. Like in our mechanism, trading is dynamic and information revelation occurs over time. In our mechanism, the message space is also given by demand/supply schedules. However, the asset being traded here has no intrinsic value.

Since we are focusing on supermodular investment games, we also relate to a larger literature on information frictions in coordination games. Examples like global games [Carlsson and Van Damme \(1993a\)](#) and the Email game [Rubinstein \(1989a\)](#) have shown that incomplete information can have dramatic consequences on the outcomes of such coordination games. Applications include currency attacks, and bank runs but also simple investment problems [Morris and Shin \(2003b\)](#). Our mechanism can be viewed as a general tool to reduce these problems through mechanism design. However, we want those tools to find applications beyond the scope of two-player binary action supermodular games with global game information structures.

From this angle, a related paper is [Angeletos and Werning \(2006\)](#), where a financial market is used to prevent the global game selection that is introduced by perturbing private information of agents in a large economy. In their set-up, a public but noisy signal that results from the market price of an asset is used as a device to facilitate coordination - which is eliminated in a global games information structure. Their mechanism is tailored to the class of information structures that yield the global game selection, which applies mainly to

two-player binary-action super-modular games². In this paper, we take the role of information aggregation more literally and require the trading mechanism to yield a consensus. An important feature with many-action supermodular games is that agents have an incentive to exaggerate moderately good news: If agent i knows that the state is good but not great, revealing this to other agents only results in moderate investment levels. If instead, she could convince others that the state is great, this would incentivize other agents to provide much higher investment levels, which in turn benefits i . When there are only two actions, The state is either too low to invest (even if all other players invest) or it is high enough, in which case full investment is an equilibrium anyway.

²While many interpretations of global games feature a continuum of states and players, there is usually a binary regime state which switches with a cut-off rule. This makes it effectively a two-player, binary action supermodular game with a continuum of types per player.

Appendix A

Proofs Chapter 1

A.1 Appendix: Omitted Proofs

Claim 8. 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

$$\sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i})(s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}}) \geq 0, \forall a'_i \in A_i. \quad (\text{A.1})$$

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

$$B_i = \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^\sigma} \text{br}_i \left(\int_{S \times S} \kappa_{-i}(\cdot|y, s') d\nu_\sigma(y, s'|s) \right), \quad (\text{A.2})$$

where $\mathcal{S}_{-i}^\sigma := \{\kappa_{-i}: S \times S \rightarrow \Delta(A_{-i}) : \forall y, s', \text{supp}(\kappa_{-i}(y, s')) \subseteq \mathcal{R}_{-i}^\sigma(s')\}$. For every $\kappa_{-i} \in \mathcal{S}_{-i}^\sigma$, define $\hat{\kappa}_{-i}: \mathcal{A}_{-i} \rightarrow \Delta(A_{-i})$, which for every $B_{-i} \in \mathcal{A}_{-i}$ satisfies

$$\hat{\kappa}_{-i}(a_{-i}|B_{-i}) := \int_S \int_{\mathcal{R}_{-i}^\sigma(B_{-i})} \kappa_{-i}(a_{-i}|y, s') d\nu_\sigma(y, s'|s), \forall a_{-i} \in B_{-i}. \quad (\text{A.3})$$

We conclude that

$$\begin{aligned} \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^\sigma} \text{br}_i \left(\int_{S \times S} \kappa_{-i}(\cdot|y, s') d\nu_\sigma(y, s'|s) \right) &= \bigcup_{\kappa_{-i} \in \mathcal{S}_{-i}^\sigma} \text{br}_i \left(\sum_{B_{-i} \in \mathcal{A}_{-i}} P_i^\sigma(B_{-i}|s) \hat{\kappa}_{-i}(a_{-i}|B_{-i}) \right) \\ &= \bigcup_{\kappa \in \Sigma_{-i}} \text{br}_i \left(\sum_{B_{-i} \in \mathcal{A}_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i}) \right) \end{aligned} \quad (\text{A.4})$$

□

Proof of Lemma 1

Proof. Let $(a_i, a'_i) \in B_i \times B'_i$. Claim 8 above shows that for every $s \in \partial\mathcal{R}^\sigma(B, B')$, there exists a selection rule $\kappa \in \Sigma_{-i}$ so that

$$\sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i}) (s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}}) = 0. \quad (\text{A.5})$$

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

$$\begin{aligned} \sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i}) (s_i^{a_i, a_{-i}} - s_i^{\hat{a}_i, a_{-i}}) &\geq 0, \quad \forall \hat{a}_i \in A_i, \\ \sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa'(a_{-i}|B_{-i}) (s_i^{a'_i, a_{-i}} - s_i^{\hat{a}_i, a_{-i}}) &\geq 0, \quad \forall \hat{a}_i \in A_i. \end{aligned} \quad (\text{A.6})$$

But then we have that

$$\begin{aligned} \sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa'(a_{-i}|B_{-i}) (s_i^{a'_i, a_{-i}} - s_i^{a_i, a_{-i}}) &\geq 0 \\ \sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i}) (s_i^{a'_i, a_{-i}} - s_i^{a_i, a_{-i}}) &\leq 0. \end{aligned} \quad (\text{A.7})$$

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

$$\sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) (\alpha \kappa'(a_{-i}|B_{-i}) + (1 - \alpha) \kappa(a_{-i}|B_{-i})) (s_i^{a'_i, a_{-i}} - s_i^{a_i, a_{-i}}) = 0. \quad (\text{A.8})$$

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

$$\sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|\tilde{s}_\epsilon) \tilde{\kappa}(a_{-i}|B_{-i}) (\tilde{s}_{\epsilon i}^{a_i, a_{-i}} - \tilde{s}_{\epsilon i}^{a'_i, a_{-i}}) < 0, \quad (\text{A.9})$$

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

$$\lim_{n \uparrow \infty} \max_{\tilde{\kappa} \in \Sigma_{-i}} \sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|\tilde{s}_{\epsilon^n}) \tilde{\kappa}(a_{-i}|B_{-i}) (\tilde{s}_{\epsilon^n i}^{a_i, a_{-i}} - \tilde{s}_{\epsilon^n i}^{a'_i, a_{-i}}) = 0, \quad (\text{A.10})$$

which concludes the proof. \square

Proof of Proposition 9

Proof. We have shown in Claim 8 that every $s \in \mathcal{R}^\sigma(B)$ satisfies (A.1). 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

$$\sum_{B_{-i}, a_{-i}} P_i^\sigma(B_{-i}|s) \kappa(a_{-i}|B_{-i}) (s_i^{a_i, a_{-i}} - s_i^{a'_i, a_{-i}}) \geq 0, \quad \forall a'_i \in A_i. \quad (\text{A.11})$$

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 9 then follows from the continuity of $s \mapsto P_i^\sigma(B_{-i}|s)$. \square

Proof of Lemma 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 $s \in S$

$$\eta(\mathbb{R}^{m,\sigma}(s)) = \mathbb{R}^{m,\sigma}(Xs). \quad (\text{A.12})$$

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

$$\begin{aligned} \bar{P}^{m,\sigma}(b|s) &:= \nu \times \nu_0 \left(\left\{ (E, E', y) : B = \mathbb{R}^{m,\sigma}(e^{\sigma E'} y) \right\} \mid e^{\sigma E} y = s \right) \\ &= \nu \left(\left\{ (E, E') : B = \mathbb{R}^{m,\sigma}(e^{\sigma(E'-E)} s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = \mathbb{R}^{m,\sigma}(X e^{\sigma(E'-E)} s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = \mathbb{R}^{m,\sigma}(X e^{\sigma(E'-E)} X^\top X s) \right\} \right) \\ &= \nu \left(\left\{ (E, E') : \eta(B) = \mathbb{R}^{m,\sigma}(e^{\sigma(E'-E)} X s) \right\} \right) \\ &= \bar{P}^{m,\sigma}(\eta(B)|Xs). \end{aligned} \quad (\text{A.13})$$

Suppose s, s' satisfy

$$\mathbb{R}^{m+1,\sigma}(s) = \mathbb{R}^{m+1,\sigma}(s'). \quad (\text{A.14})$$

Then by property (1.47), we must have that

$$\mathbb{R}^{m+1,\sigma}(Xs) = \mathbb{R}^{m+1,\sigma}(Xs'). \quad (\text{A.15})$$

□

Claim 9. 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 10. For every $\tilde{X} \in \mathbb{P}^{2n}$ there exists permutation $\tilde{\eta}: \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \mathcal{A}_1 \times \mathcal{A}_2$ so that $(\tilde{\eta}, \tilde{X})$ preserves best-replies.

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

$$\eta_\zeta(b_1, b_2) := (\{\zeta_1(a_1, a_2) : (a_1, a_2) \in B_1 \times B_2\}, \{\zeta_2(a_1, a_2) : (a_1, a_2) \in B_1 \times B_2\}). \quad (\text{A.16})$$

□

Proof of Lemma 3

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

$$\mathcal{X}(\eta) := \{\tilde{X} \in \mathbb{P}^{2n} : (\eta, \tilde{X}) \in \mathcal{I}\}. \quad (\text{A.17})$$

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

$$Xs^* \neq s^*. \quad (\text{A.18})$$

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

$$\hat{X}^{(i,a),(i',a')} = \begin{cases} \pm 1, & \text{if } (i, a) \in \mathcal{K} \text{ and } \omega_{\mathcal{K}}(i, a) = (i', a'), \\ 1, & \text{if } (i, a) \notin \mathcal{K} \text{ and } (i, a) = (i', a'), \\ 0, & \text{otherwise,} \end{cases} \quad (\text{A.19})$$

and

$$\hat{X}s^* = Xs^*. \quad (\text{A.20})$$

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

$$\{\hat{X}s : s \in S^+\} = S^-, \quad (\text{A.21})$$

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 $\gamma^+ : [0, 1] \rightarrow \overline{S^+}$ and $\gamma^- : [0, 1] \rightarrow \overline{S^-}$ satisfying the following conditions

- (i) $\mathcal{R}^\sigma(\gamma^+(x)) = \mathcal{R}^\sigma(\gamma^-(x'))$, 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 $\mathcal{R}^\sigma(\mathcal{R}(s^*)) \cap S^+ \neq \emptyset$ and $\mathcal{R}^\sigma(\mathcal{R}(s^*)) \cap S^- \neq \emptyset$. We show in Claim 9 that each rationalizable set is path-connected and so the existence of γ^+, γ^- satisfying properties (i)-(iii) is guaranteed. Suppose now that

$$\mathcal{R}^\sigma(s^*) \neq \mathcal{R}^\sigma(\hat{X}s^*). \quad (\text{A.22})$$

We show in Claim 10 that for every matrix $\tilde{X} \in \mathbb{P}^{2n}$ there is a permutation of action set labels $\tilde{\eta}$ so that $(\tilde{\eta}, \tilde{X}) \in \mathcal{I}$. Since $\tilde{X} \in \mathbb{P}^{2n}$ is itself part of an invariance, it must be that

$$\mathcal{R}^\sigma(\gamma^+(x)) \neq \mathcal{R}^\sigma(\tilde{X}\gamma^+(x)), \quad \forall x \in [0, 1]. \quad (\text{A.23})$$

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

$$\mathcal{R}^\sigma(\gamma^+(x)) = \mathcal{R}^\sigma(\hat{X}\gamma^+(x)), \quad \forall x \in [0, 1]. \quad (\text{A.24})$$

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

$$\hat{X}^\top Xs^* = s^*. \quad (\text{A.25})$$

□

Proof of Lemma 4

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

$$p_{i,s}^\sigma(B) := \nu \times \nu_0 \left(\{(E_i, E_{-i}, y) : B = R^\sigma(e^{\sigma E_{-i}} y)\} \mid e^{\sigma E_i} y = s \right). \quad (\text{A.26})$$

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

$$p_{i,s}(B) := \lim_{\sigma \downarrow 0} p_{i,s_\sigma^k}^\sigma(B), \quad \forall B \in \mathcal{A}_1 \times \mathcal{A}_2. \quad (\text{A.27})$$

Since every collision is full, by varying the permutation, we deduce from Claims 2 and 3 that

$$\sum_{k=1}^{m_i} p_{i,s^k}(B) = 1. \quad (\text{A.28})$$

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

$$\tilde{p}_{i,k,s}(a, B) := \frac{p_{i,s}(B)}{\prod_{j=1,2} \prod_{\hat{k} \neq k} d_{s_j}^{a_j^{\hat{k}}, a_j^{\hat{k}}, a_{-j}}}. \quad (\text{A.29})$$

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

$$\lim_{\sigma \downarrow 0} g_i^{a_i^k, a_i^{k^*}, \sigma}(s) = \sum_{B \in \mathcal{A}_1 \times \mathcal{A}_2} \sum_{a \in B_1 \times B_2} \tilde{p}_{i,k,s}(a, B) \underbrace{\prod_{j=1,2} \prod_{\hat{k}=1}^{m_j} d_{s_j}^{a_j^{\hat{k}}, a_j^{\hat{k}}, a_{-j}}}_{:= \pi_{a,s}} = 0. \quad (\text{A.30})$$

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

$$\bar{p}_{i,s}(a, B) = \sum_{k=1}^{m_i} \tilde{p}_{i,k,s}(a, B) \prod_{j=1,2} \prod_{\hat{k}=1}^{m_j} d_{s_j}^{a_j^{\hat{k}}, a_j^{\hat{k}}, a_{-j}}. \quad (\text{A.31})$$

By condition (A.28), \bar{p}_s is a vector with entries in $\{0, 1\}$. Hence the limit of every IC takes the same form:

$$\delta_C(d_s^1, \dots, d_s^K) = \sum_{a \in A_1 \times A_2} \chi_{a,C}(s) \prod_{i=1,2} \prod_{k=1}^{m_i} d_{s_i}^{a_i^k, a_i^{k^*}, a_{-i}}, \quad (\text{A.32})$$

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 . \square

Proof of Proposition 5

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 $((a_1^1, \dots, a_1^m), (a_2^1, \dots, a_2^m)) \in G_{C,1} \times G_{C,2}$ 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 $\xi := (0, \dots, 1, \dots, -1, \dots, 0)$, 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 $n \leq m$,

$$a_i^n \in R_i^\sigma(s + \epsilon\xi). \quad (\text{A.33})$$

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

$$\frac{\partial p_i^{\sigma, a_i^2, a_i^1}(a_{-i}^2 | s)}{\partial d_i^{a_i^2, a_i^1, a_{-i}^1}} > 0. \quad (\text{A.34})$$

where

$$p_i^{\sigma, a_i^2, a_i^1}(a_{-i}^2 | s) := \sum_{b_{-i} \in \mathcal{A}_{-i}} P^\sigma(B_{-i} | s) \kappa(a_{-i}^2 | B_{-i}), \quad (\text{A.35})$$

where $\kappa \in \Sigma_{-i}$ solves (1.37) 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

$$\frac{\partial g^{a_i^2, a_i^1, \sigma}}{\partial p_i^{\sigma, a_i^2, a_i^1}(a_{-i}^2 | s)} > 0. \quad (\text{A.36})$$

Expression (A.38) 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 $d_{s_i}^{a_i, a_i', a_{-i}}$, $|R(s + \epsilon\xi)| = 1$, 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 $n \leq m$,

$$\frac{\partial g^{a_i^*, a_i^n, \sigma}}{\partial d_{s_i}^{a_i, a_i', a_{-i}}} > 0. \quad (\text{A.37})$$

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

$$a_i^n \notin R_i^\sigma(s + \epsilon\xi). \quad (\text{A.38})$$

And so the result follows. \square

Lemma 19 (AI). *Suppose s is strictly supermodular and strictly concave. Then for all actions $(a_i, a_i') \in C_i$, $(a_{-i}, a_{-i}') \in C_{-i}$,*

$$\begin{aligned} & (a_{-i} \in I^{a_i, a_i'}, a_i \in I^{a_{-i}, a_{-i}'}) \\ & \text{or} \\ & (a_{-i}' \in I^{a_i', a_i}, a_i' \in I^{a_{-i}', a_{-i}}) \end{aligned} \iff \frac{d_{s_i}^{a_i, a_i', a_{-i}}}{d_{s_{-i}}^{a_{-i}, a_{-i}', a_i}} > 0 \quad (\text{A.39})$$

Proof. Recall that

$$\delta_C(d_s^1, \dots, d_s^K) = \sum_{a \in \bar{b}} \chi_{a,C} \prod_{i=1,2} \prod_{k=1}^{m_i} d_{s_i}^{a_i^k, a_i'^k, a_{-i}}. \quad (\text{A.40})$$

Then we have that

$$\frac{\partial \delta_C}{\partial d_{s_i}^{a_i, a_i', a_{-i}}} = \prod_{k \neq k_{\substack{a_i, a_i', a_{-i} \\ d_{s_i}}} d_{s_i}^{a_i^k, a_i'^k, a_{-i}} \sum_{\tilde{a} \in A_1 \times A_2 : \tilde{a}_{-i} = a_{-i}} \chi_{\tilde{a}, C} \prod_{k=1}^{m_{-i}} d_{s_{-i}}^{a_{-i}^k, a_{-i}'^k, \tilde{a}_i}, \quad (\text{A.41})$$

where

$$k_{\substack{a_i, a_i', a_{-i} \\ d_{s_i}}} := \{k \leq m_i : a_i^k = a_i, a_i'^k = a_i'\}. \quad (\text{A.42})$$

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

$$\text{sgn} \left(\frac{\partial \delta_C}{\partial d_{s_i}^{a_i, a_i', a_{-i}}} \right) = \text{sgn} \left(\prod_{k \neq k_{\substack{a_i, a_i', a_{-i} \\ d_{s_i}}} d_{s_i}^{a_i^k, a_i'^k, a_{-i}} \prod_{k=1}^{m_{-i}} d_{s_{-i}}^{a_{-i}^k, a_{-i}'^k, a_i} \right). \quad (\text{A.43})$$

where a_i satisfies

$$a_i \in \max_{\prec_i} \{ \tilde{a}_i \in A_i : \chi_{a_{-i}, \tilde{a}_i, C} = 1 \}. \quad (\text{A.44})$$

Hence

$$\frac{\frac{\partial \delta_C}{\partial d_{s_i}^{a_i, a_i', a_{-i}}}}{\frac{\partial \delta_C}{\partial d_{s_{-i}}^{a_{-i}, a_{-i}', a_i}}} = \frac{d_{s_{-i}}^{a_{-i}, a_{-i}', a_i}}{d_{s_i}^{a_i, a_i', a_{-i}}}. \quad (\text{A.45})$$

□

Proof of Lemma 6

Proof. This result then follows from Lemma 19 and Proposition 4. □

Proof of Theorem 1:

Proof. The argument in the main text establishes that if every $s \in O$ has no risk-dominant better response cycles in any BRC set, then limit uniqueness holds. We thus focus on the converse. Let $O \subseteq S$ satisfy limit uniqueness and suppose $s \in S$ admits a BRC with a risk-dominant better response cycle $\{(a_1^1, \dots, a_1^m), (a_2^1, \dots, a_2^m)\}$ and a player i so that for every $n \in \{1, \dots, m-1\}$,

$$\begin{aligned} u_i(a_i^n, a_{-i}^{n+1}) &< u_i(a_i^{n+1}, a_{-i}^{n+1}), \\ u_{-i}(a_{-i}^n, a_i^n) &< u_{-i}(a_{-i}^{n+1}, a_i^n), \end{aligned} \quad (\text{A.46})$$

and

$$\begin{aligned} u_i(a_i^m, a_{-i}^1) &< u_i(a_i^1, a_{-i}^1), \\ u_{-i}(a_{-i}^m, a_i^m) &< u_{-i}(a_{-i}^1, a_i^m). \end{aligned} \quad (\text{A.47})$$

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 $\mathcal{P}_{s, \bar{s}} \subseteq S$,

$$\mathcal{P}_{s, \bar{s}} := \{s_\alpha : \alpha \in [0, 1]\}, \quad (\text{A.48})$$

where $s_\alpha := \frac{\alpha s + (1-\alpha)\bar{s}}{\|\alpha s + (1-\alpha)\bar{s}\|_2}$, for every $\alpha \in [0, 1]$. By Lemma 7 we conclude that \bar{s} satisfies limit multiplicity, i.e. for every player i ,

$$\{a_i^1, \dots, a_i^m\} \subseteq R_i(\bar{s}). \quad (\text{A.49})$$

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 ,

$$\{a_i^1, \dots, a_i^m\} \subseteq R_i(s_{\alpha^l}). \quad (\text{A.50})$$

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 3, 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 $\{a_i^1, \dots, a_i^m\} \subseteq R_i(s_{\alpha^1})$. The same argument then implies that (A.50) holds for all $l \leq L$. \square

Proof of Proposition 5

Proof. The proof is by construction.

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

$$\text{span}(L_{u, u'}u, L_{u, u'}u') = \text{span}(e_1, e_2), \quad (\text{A.51})$$

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

$$\varphi(u) = \exp \left(\underbrace{\theta L_{u, u'}^\top \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & & \\ \vdots & \vdots & & \ddots & \\ 0 & 0 & & & 0 \end{pmatrix} L_{u, u'}}_{:= E_{u, u'}} \right) \varphi(u') \quad (\text{A.52})$$

where

$$\theta := \arccos(\langle \varphi(u), \varphi(u') \rangle). \quad (\text{A.53})$$

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

$$\varphi(y + \sigma\varepsilon) = e^{f_{y, \varepsilon}(\sigma) E_{y+\varepsilon, y}} \varphi(y), \quad (\text{A.54})$$

where

$$f_{y, \varepsilon}(\sigma) := \frac{\arccos(\langle \varphi(y), \varphi(y + \sigma\varepsilon) \rangle)}{\arccos(\langle \varphi(y), \varphi(\varepsilon) \rangle)}. \quad (\text{A.55})$$

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 ,

$$\Gamma_{s,i}: (\varepsilon_1, \varepsilon_2) \mapsto (E_{s,s-\varepsilon_i}, E_{s+(\varepsilon_{-i}-\varepsilon_i),s-\varepsilon_i}) \quad (\text{A.56})$$

Define the induced distribution over pairs of rotations

$$\tilde{\nu}_s := \mu_1 \times \mu_2 \circ \Gamma_s^{-1}. \quad (\text{A.57})$$

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

$$\hat{\mathcal{R}}_\epsilon^\sigma(B) := \{u \in \mathcal{N}_\epsilon(S) : \bar{\mathbf{R}}^\sigma(u) = B\}, \quad (\text{A.58})$$

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

$$\hat{\mathcal{R}}_\epsilon^\sigma(B|S) := \{\varphi(u) : u \in \hat{\mathcal{R}}_\epsilon^\sigma(B)\}. \quad (\text{A.59})$$

Then we have that

$$\begin{aligned} & \lim_{\epsilon \downarrow 0} \lim_{\sigma \downarrow 0} \tilde{\nu}_s(\{(E_1, E_2) : \exp(f_{e^{-E_i s, s-e^{-E_i s}}(\sigma)}(E_{-i} - E_i))s \in \hat{\mathcal{R}}_\epsilon^\sigma(B|S)\}) \\ &= \lim_{\epsilon \downarrow 0} \lim_{\sigma \downarrow 0} \mu_1 \times \mu_2(\{(\varepsilon_1, \varepsilon_2) : s + \sigma(\varepsilon_{-i} - \varepsilon_i) \in \hat{\mathcal{R}}_\epsilon^\sigma(B)\}), \end{aligned} \quad (\text{A.60})$$

and so the result follows. □

A.2 Risk-Dominance and Limit Multiplicity

In Subsection A.2.2, 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 A.2.3 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.

A.2.1 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).

		H		L	
		a	b	a	b
H	a	1,-1	-1,1	0	
	b	-1,1	1,-1		
L	a	0		$\epsilon,-\epsilon$	$-\epsilon,\epsilon$
	b			$-\epsilon,\epsilon$	$\epsilon,-\epsilon$

Table A.1: 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 R, 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
		1,1	0,0
H	0,0	ϵ, ϵ	

Table A.2: 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 A.1, s , we would expect that for ϵ small, the limit selection is of the form

$$R_i(s) = \{(H, a), (H, b)\}, \forall i = 1, 2. \tag{A.61}$$

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

A.2.2 Risk-Dominance

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 20. *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 20 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 $\mathcal{R}(s)$ 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 21 (Upper Bound on Limit Selection). *For every $s \in S$ and every player i ,*

$$R_i(s) \subseteq \bigcup_{(B_1, B_2) \in \mathcal{R}(s)} B_i. \quad (\text{A.62})$$

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.

Proof of Lemma 20

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 sets fail the risk-dominance criterion (??), i.e.

$$p_1^B(s) + p_2^B(s) > 1. \quad (\text{A.63})$$

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

$$\text{br}_i(p_i^* | s_i) \cap B_i = \emptyset, \quad (\text{A.64})$$

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

$$\text{br}_i(p_i^{**} | s_i) \subseteq \hat{B}_i, \quad (\text{A.65})$$

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

Proof of Lemma 21

Proof. Fix $\sigma > 0$ and suppose that for all $B \in \mathcal{R}(s)$, $a_i \notin B_i$. We show that $a_i \notin R(s)$. 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

$$B_i^* := \bigcup_{(B_1, B_2) \in \mathcal{R}(s)} B_i. \quad (\text{A.66})$$

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 $a_i \in \text{br}_i(p_i)$ for each $i = 1, 2$,

$$p_1(B_2) + p_2(B_1) > 1, \quad \forall (B_1, B_2) \in \mathcal{R}(s). \quad (\text{A.67})$$

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

$$\beta_i(a_{-i}) \in \arg \max_{a_i \in A_i} s_i^{a_i, a_{-i}}. \quad (\text{A.76})$$

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 $l < n_c$,

$$\beta_1(a_1^l) = a_2^l, \quad \beta_2(a_2^l) = a_1^{l+1}, \quad (\text{A.77})$$

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 (1.35). Say that s has a *risk-dominant cyclic component* if the cycle is the union of all risk-dominant BRC sets, i.e.

$$C_i(s) = \bigcup_{(B_1, B_2) \in \mathcal{R}(s)} B_i. \quad (\text{A.78})$$

Proposition 7 below shows that a risk-dominant cyclic component is always contained in the selection.

Proposition 7 (Sufficient Condition for Limit Multiplicity). *Let $s \in S$ have a risk-dominant cyclic component $C(s)$ then for every player i ,*

$$C_i(s) \subseteq R_i(s). \quad (\text{A.79})$$

Proof. Since $R_i(s) \neq \emptyset$, we conclude from Lemma 21 that $R_i(s) \subseteq C_i(s)$. Since $R(s)$ is itself a BRC set it cannot be smaller than $C(s)$. \square

Coming back to the example game in Table A.1: 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

$$p_i^{B^H} = \frac{\epsilon}{1 + \epsilon}. \quad (\text{A.80})$$

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 .

A.3 Topological Properties of Limit Selection

Proposition 8 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 7.

Proposition 8 (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 open¹ set $O \subseteq S$ so that*

$$O \subseteq \lim_{\sigma \downarrow 0} \mathcal{R}^\sigma(B). \quad (\text{A.81})$$

Proposition 9 below establishes upper-hemi continuity of \mathcal{R}^σ . Upper-hemi continuity was established in Dekel et al. (2007a) 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 9 (Upper-Hemi Continuity). *For every $\sigma > 0$ and any pair $(B, B') \in (\mathcal{A}_1 \times \mathcal{A}_2)^2$ so that $\partial \mathcal{R}^\sigma(B, B') \neq \emptyset$,*

$$\partial \mathcal{R}^\sigma(B, B') = \partial \mathcal{R}^\sigma(B \cup B', B'). \quad (\text{A.82})$$

Define the collection of nested action set profiles,

$$\mathcal{A} := \{(B, B') \in (\mathcal{A}_1 \times \mathcal{A}_2)^2 : \forall i, B'_i \subseteq B_i\}. \quad (\text{A.83})$$

We now study the structure of collisions. By Proposition 9 it is without loss of generality to restrict attention to collisions of action set-pairs in \mathcal{A} .

Proof of Proposition 8

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,

$$s_1^{a_1, a_2} = \begin{cases} 1, & \text{if } \exists k < |B_i| \text{ s.t. } a_1 = a_1^k, a_2 = a_2^{k+1} \\ 1, & \text{if } a_1 = a_1^{|B_i|}, a_2 = a_2^1 \\ -1, & \text{if } \exists k \text{ s.t. } a_1 = a_1^k, a_2 \neq a_2^{k+1} \\ -1, & \text{if } a_1 = a_1^{|B_i|}, a_2 \neq a_2^1 \\ \epsilon, & \text{otherwise.} \end{cases} \quad (\text{A.84})$$

For player 2,

$$s_2^{a_1, a_2} = \begin{cases} 1, & \text{if } \exists k > 1 \text{ s.t. } a_2 = a_2^k, a_1 = a_1^{k-1} \\ 1, & \text{if } a_2 = a_2^1, a_1 = a_1^{|B_i|} \\ -1, & \text{if } \exists k > 1 \text{ s.t. } a_2 = a_2^k, a_1 \neq a_1^{k-1} \\ -1, & \text{if } a_2 = a_2^1, a_1 \neq a_1^{|B_i|} \\ \epsilon, & \text{otherwise.} \end{cases} \quad (\text{A.85})$$

¹A subset $O \subseteq S$ is open if it is the intersection of an open set in U and S .

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

□

Appendix B

Proofs Chapter 2

B.1 Appendix

B.1.1 Preliminaries

We introduce some additional notation. For any mapping $f : X \rightarrow Y$ we denote the image of f by $\text{Im}(f)$. The following lemma is key for our results:

Lemma 22 (BR Factorization of ICR). *Let $(T_i, \pi_i)_i$ be a Harsanyi type space. Then for every m and every i , \mathbb{R}_i^m admits the following factorization through BR_i ,*

$$\begin{array}{ccc} T_i & \xrightarrow{\pi_i} & \Delta_{K \times T_{-i}} \\ \downarrow \mathbb{R}_i^m & & \text{id} \downarrow \quad \downarrow \mathbb{R}_{-i}^{m-1} \\ \mathcal{B}_i & \xleftarrow{\text{BR}_i} & \Delta_{K \times \mathcal{B}_{-i}} \end{array}$$

Proof. Let $\sigma : K \times T_{-i} \rightarrow \Delta(A_{-i})$ be a $\pi_i(t_i)$ -measurable conjecture. Write the t_i mixture of σ as

$$\langle \sigma, \pi_i(t_i) \rangle_{T_i}(k, a_{-i}) := \int_{T_{-i}} \sigma(k, t_{-i})(a_{-i}) \pi_i(t_i)(k, dt_{-i}), \quad \forall k, a_{-i}$$

Then by definition of \mathbb{R} we have that

$$\mathbb{R}_i^m(t_i) = \{ \mathbf{B}(\langle \sigma, t_i \rangle_{T_i}) : \sigma \text{ is } \pi_i(t_i)\text{-meas.}, \text{supp } \sigma(k, t_{-i}) \subseteq \mathbb{R}_{-i}^{m-1}(t_{-i}) \}$$

where $\mathbf{B}(p) := \text{argmax}_{a_i} \sum_{k, a_{-i}} u_i(k, a_i, a_{-i}) p(k, a_{-i})$ for every $p \in \Delta(K \times A_{-i})$. We now show that for every $\pi_i(t_i)$ -measurable conjecture $\sigma : K \times T_{-i} \rightarrow \Delta(A_{-i})$ so that $\text{supp } \sigma(k, t_{-i}) \subseteq \mathbb{R}_{-i}^{m-1}(t_{-i})$ we can construct a conjecture $\tilde{\sigma} : K \times \mathcal{A}_{-i} \rightarrow \Delta(A_{-i})$ so that $\text{supp } \sigma(k, b_{-i}) \subseteq b_{-i}$ and

$$\langle \tilde{\sigma}, p_i^{m-1}(t_i) \rangle(k, a_{-i}) = \langle \sigma, \pi_i(t_i) \rangle_{T_i}(k, a_{-i}), \quad \forall k, a_{-i} \quad (\text{B.1})$$

where $p_i^{m-1}(t_i) := \pi_i(t_i) \circ (\text{id} \times \mathbb{R}_{-i}^{m-1})^{-1}$ is the push forward probability. Define the required conjecture for every k, a_{-i} and $b_{-i} \in \text{Im}(\mathbb{R}_{-i}^{m-1})$,

$$\tilde{\sigma}(a_{-i} | k, b_{-i}) := \int_{(\mathbb{R}_{-i}^{m-1})^{-1}(b_{-i})} \sigma(k, t_{-i})(a_{-i}) \pi_i(t_i)(k, dt_{-i})$$

and for $b_{-i} \notin \text{Im}(R_{-i}^{m-1})$ let $\tilde{\sigma}$ be arbitrary. This construction clearly satisfies (B.1). Conversely, $\tilde{\sigma}$ can be written as a conjecture $K \times T_{-i} \rightarrow \Delta(A_{-i})$ which is constant on the pre-image of R_{-i}^{m-1} and so the result follows. \square

A *monotone stochastic transformation* for player i is a map $\rho_i: K \times B_{-i} \rightarrow \Delta(B_{-i})$ so that for every $b \in B$ and $k \in K$,

$$b'_{-i} \subseteq b_{-i}, \quad \forall b'_{-i} \in \text{supp}(\rho_i(k, b_{-i})). \quad (\text{B.2})$$

Lemma 23 (Monotonicity of BR). *For any monotone stochastic transformation $\rho_i: K \times \mathcal{B}_{-i} \rightarrow \Delta(\mathcal{B}_{-i})$ and for any $p_i \in \Delta(K \times \mathcal{B}_{-i})$,*

$$\text{BR}_i(p_i \circ \rho_i) \subseteq \text{BR}_i(p_i), \quad (\text{B.3})$$

where for all $k \in K$ and $b_{-i} \in B_{-i}$,

$$p_i \circ \rho_i(k, b_{-i}) := \sum_{b'_{-i} \in B_{-i}} \rho_i(b_{-i}|k, b'_{-i}) p_i(k, b'_{-i}). \quad (\text{B.4})$$

Proof. Consider any conjecture $\sigma_i: K \times B_{-i} \rightarrow \Delta(A_{-i})$ so that $\text{supp}(\sigma(\cdot|k, b_{-i})) \subseteq b_{-i}$ for all $k \in K, b_{-i} \in B_{-i}$. Now define the conjecture $\sigma_i \circ \rho_i$, which for every $a_{-i} \in A_{-i}, k \in K, b'_{-i} \in B_{-i}$ is given by

$$\sigma_i \circ \rho_i(a_{-i}|k, b'_{-i}) := \sum_{b_{-i}} \sigma_i(a_{-i}|k, b_{-i}) \rho_i(b_{-i}|k, b'_{-i}). \quad (\text{B.5})$$

Since ρ_i is monotone, the conjecture $\sigma_i \circ \rho_i$ also satisfies the support constraint of σ_i . Hence

$$\begin{aligned} \langle \sigma_i, p_i \circ \rho_i \rangle(k, a_{-i}) &= \sum_{b'_{-i} \in B_{-i}} \left(\sum_{b_{-i} \in B_{-i}} \sigma_i(a_{-i}|k, b_{-i}) \rho_i(b_{-i}|k, b'_{-i}) \right) p_i(k, b'_{-i}) \\ &= \sum_{b'_{-i} \in B_{-i}} \sigma_i \circ \rho_i(a_{-i}|k, b'_{-i}) p_i(k, b'_{-i}) \\ &= \langle \sigma_i \circ \rho_i, p_i \rangle(k, a_{-i}). \end{aligned} \quad (\text{B.6})$$

Now the result is immediate from the definition of BR_i . \square

B.1.2 Strategic Type Spaces

Lemma 7 *For every strategic type space $(\mathcal{S}_i, \psi_i)_i$ and every $m \in \mathbb{N}$, there exists measurable $\sigma_i^m: \mathcal{S}_i \rightarrow \mathcal{A}_i$ so that for every Harsanyi type space $(T_i, \pi_i)_i$ and associated maps $(\eta_i)_i$ satisfying the diagram of Definition 1,*

$$R_i^m(t_i) = \sigma_i^m \circ \eta_i(t_i), \quad \forall t_i \in T_i, \quad \forall i \in N \quad (\text{B.7})$$

Proof. Let $(\mathcal{S}_i, \psi_i)_i$ be a STS and let $(\Sigma_i)_i$ be a measurable family of strategically closed behaviors. Proceed by induction. Base case: The constant map $\sigma_i^0: s_i \mapsto A_i$ is in Σ_i and for every Harsanyi type space $(T_i, \pi_i)_i$, $R_i^0(t_i) = A_i = \sigma_i^0 \circ \eta_i(t_i)$ for all i and $t_i \in T_i$.

Inductive hypothesis: Suppose $(\sigma_i^l)_{l \leq m, i}$ satisfy (B.7). Then for every player i and type t_i , $\pi_i(t_i) \circ (id \times R_{-i}^{m-1})^{-1} = \pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1}$ and so by Lemma 22 $R_i(t_i) = BR_i(\pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1})$. By the *quotient property* in Definition 1 $\psi_i(\pi_i(t_i) \circ (id \times \eta_{-i})^{-1}) = \eta_i(t_i)$ and by the *strategic closure property* in Definition 2 there exists σ_i^m so that $\sigma_i^m \circ \eta_i(t_i) = BR_i(\pi_i(t_i) \circ (id \times \sigma_{-i}^{m-1} \circ \eta_{-i})^{-1})$. \square

Lemma 8

- (i) Let $s^m \in \mathcal{B}^m$, then $s^m \in \mathcal{S}^m$ if and only if there exists a Harsanyi type space (T, π) and a type profile $t \in T$ so that $s^m = (R^l(t))_{l \leq m}$.
- (ii) Let $s \in \mathcal{B}^{\mathbb{N}}$, then $s \in \mathcal{S}$ if and only if there exists a Harsanyi type space (T, π) and a type profile $t \in T$ so that $s = (R^l(t))_{l \geq 0}$.

Proof. We start with (i): We prove both directions inductively using the following base case: For every Harsanyi type space $(T_i, \pi_i)_i$ and player i we have $\mathcal{S}_i^0 = \{A_i\} = R_i^0(t_i)$. For the “if” direction: Inductive hypothesis: Fix a Harsanyi type space $(T_i, \pi)_i$, and suppose that for any player i and type $t_i \in T_i$, $(R_i^l(t_i))_{l \leq m-1} \in \mathcal{S}_i^{m-1}$. Then by Lemma 22, $R_i^m(t_i) = BR_i(\pi_i(t_i) \circ (id \times R_{-i}^{m-1})^{-1})$ and so $\pi_i(t_i) \circ (id \times \prod_{l \leq m-1} R_{-i}^l)^{-1} \in \Delta(K \times \mathcal{S}_{-i}^{m-1})$ implies $(R_i^l(t_i))_{l \leq m} \in \mathcal{S}_i^m$, as required.

Prove the “only if” by constructing a type space for every $s^m \in \mathcal{S}^m$ containing a type t so that $\prod_{l \leq m} R^l(t) = s^m$. For every i , $s_i^m \in \mathcal{S}_i^m$ implies that there exists $p_i^m \in \Delta(K \times \mathcal{S}_{-i}^{m-1})$ so that $s_i^m = (BR_i(\text{marg}_{K, \mathcal{S}_{-i}^l}(p_i^m)))_{l \leq m-1}$. We proceed by induction on m . Suppose we have picked a selection $\iota_i^m: \mathcal{S}_i^m \rightarrow \Delta(K \times \mathcal{S}_{-i}^{m-1})$ from $(\psi_i^m)^{-1}(\cdot)$. For every s_i^{m+1} , we may construct a selection ι_i^{m+1} so that for every $s_i^{m+1} \in \mathcal{S}_i^{m+1}$, $\text{marg}_{K, \mathcal{S}_{-i}^m}(\iota_i^{m+1}(s_i^{m+1})) = \iota_i^m(s_i^m)$. Finally, by Lemma 22 the m -th component $s_{i,m}^m$ of s_i^m is given by $s_{i,m}^m = R_i^m(\iota_i^m(s_i^m))$, for all m , as required.

Finally, (ii) follows from the properties of inverse limits of the constructions in the proof of (i). \square

Lemma 9 (\mathcal{S}, ψ) obtained in our construction of hierarchies of best replies are an STS.

Proof. Lemmas 22 and 8 (i) ensure that (\mathcal{S}, ψ) is a type space quotient. The coordinate projection then yields a best-reply closed family of measurable strategic behaviors. \square

Theorem 2

- (i) (\mathcal{S}, ψ) is a minimal STS.
- (ii) If (\mathcal{S}', ψ') and (\mathcal{S}'', ψ'') are minimal STS then \mathcal{S}'' and \mathcal{S}' are isomorphic.

Proof. By Lemma 7 any finite order ICR hierarchy can be recovered continuously from any STS. By Lemma 8 the set \mathcal{S} coincides with all ICR hierarchies. Then by Lemma 9, (\mathcal{S}, ψ) is a STS which can be recovered from all STS. The product sigma algebra then ensures that (\mathcal{S}, ψ) is a STS which is minimal. As the property of minimality is universal, every minimal STS is isomorphic to \mathcal{S} . \square

Claim 1 $X \ll X' \implies B(X) \ll B(X')$.

Proof of Claim 1. Since $X \ll X'$ there exists a mapping $\rho: X \rightarrow \Delta(X')$ so that for all $x \in X$,

$$\rho(\{x' \in X' : \forall n, x_n \subseteq x'_n\} | x) = 1. \quad (\text{B.8})$$

For every player i and $x_i \in X_i$, we thus derive a stochastic transformation $\rho_{-i, x_i}: K \times X_{-i} \rightarrow \Delta(X_{-i})$, which for every k, x_{-i} satisfies

$$\rho_{-i, x_i}(x'_{-i} | k, x_{-i}) := \sum_{x'_i \in X'_i} \rho(x'_i, x'_{-i} | k, x_{-i}, x_i). \quad (\text{B.9})$$

Then the result follows by applying Lemma 11 at every coordinate of x_i . \square

Claim 2 *The sequence $s \in \mathcal{S}$ is maximal after round m if and only if it is maximal at round n for all $n > m$.*

Proof. One direction is immediate: If s is maximal at round n for all $n > m$ then it must be maximal. Suppose now that s is maximal after round m . For any $n \in \mathbb{N}$, let $\bar{\mathcal{S}}^n \subseteq \mathcal{S}$ denote the set of sequences which are maximal after round n . Moreover, let $\hat{\mathcal{S}}^n \subseteq \mathcal{S}$ denote the set of sequences which are maximal at round n . For every $\hat{s} \in \hat{\mathcal{S}}^n$, define the set of sequences obtained by making the n -th entry maximal and preserving the rest of the sequence:

$$\hat{\mathcal{S}}^n(\hat{s}) := \{(\hat{s}^{n-1}, \tilde{s}_n, \tilde{s}_{n+1}, \dots) : \tilde{s} \in \hat{\mathcal{S}}^{n+1} \text{ s.t. } \tilde{s}^n = \hat{s}^n\}. \quad (\text{B.10})$$

Note that for every $\bar{s} \in \mathcal{S}$ and any $\hat{s} \in \hat{\mathcal{S}}^n(\bar{s})$ we have that $\hat{s}^n \in \mathcal{S}^n$. By Claim 1 we conclude that

$$B(\mathcal{S}) \ll B(\bar{\mathcal{S}}^n) \ll B(\cup_{\bar{s} \in \bar{\mathcal{S}}^n} \hat{\mathcal{S}}^n(\bar{s})) \quad (\text{B.11})$$

Since $B(\mathcal{S}) = \mathcal{S}$, we conclude that $\mathcal{S} \ll B(\cup_{\bar{s} \in \bar{\mathcal{S}}^n} \hat{\mathcal{S}}^n(\bar{s}))$. Moreover, by construction of \mathcal{S} we must have that for all $\hat{s} \in B(\cup_{\bar{s} \in \bar{\mathcal{S}}^n} \hat{\mathcal{S}}^n(\bar{s}))$, $\hat{s}^{n+1} \in \mathcal{S}^{n+1}$. Applying this argument for all $n' \geq n$ implies that $\hat{\mathcal{S}}^n(\hat{s}) = \bar{\mathcal{S}}^n(\hat{s})$ and so $\bar{\mathcal{S}}^n = \bigcap_{l > n} \hat{\mathcal{S}}^l$, which concludes the proof. \square

Claim 3 *There exists $L \in \mathbb{N}$ so that for every $m \in \mathbb{N}$ and every $s \in \mathcal{S}$,*

$$|\bar{\mathcal{S}}^m(s^m)| < L. \quad (\text{B.12})$$

Proof. Since the sequences in \mathcal{S} are weakly decreasing in the set inclusion order, we conclude from the sequential maximality property established in Claim 2 that for any sequence $s \in \bar{\mathcal{S}}^m$ and $\bar{s} \in \bar{\mathcal{S}}^m(s^m)$,

$$\{s_n : n \in \mathbb{N}\} = \{\bar{s}_n : n \in \mathbb{N}\} \implies s = \bar{s}. \quad (\text{B.13})$$

Hence the result. \square

Claim 4 *There exists M so that for every $m \in \mathbb{N}$ and $s \in \mathcal{S}$, the set $\bar{\mathcal{S}}^m(s^m)$ converged at round $m + M$.*

Proof. We proceed inductively on $m \in \mathbb{N}$. We start with the base case: $\bar{\mathcal{S}}^0(s^0)$. By Claim 3 there is some finite M so that $\bar{\mathcal{S}}^0(s^0)$ has converged at round M . Suppose now that $\bar{\mathcal{S}}^{m-1}(s^{m-1})$ has converged at round $m - 1 + M$ for every s^{m-1} . It follows from the definition of B that $B(\bar{\mathcal{S}}^{m-1})$ has converged at round $m + M$. By maximality of $\bar{\mathcal{S}}^m$ and monotonicity of B , we deduce that $\bar{\mathcal{S}}^m(s^m) \subseteq B(\bar{\mathcal{S}}^{m-1})$, and so the result follows. \square

Claim 5 For every $m \geq 0, n > 0$, $\bar{T}^{m,n}$ is the set of sequences s that are maximal after round n , admit at most $m - 1$ non-maximal transitions before round n , and such that:

$$s \in B(\bar{T}^{m,n-1}).$$

Proof. We already established that $\bar{\mathcal{S}}^m(s^m) \subseteq B(\bar{\mathcal{S}}^{m-1})$. We deduce from the monotonicity of B that every sequence that makes $m - 1$ non-maximal transitions before round n is a best-reply to beliefs supported on sequences that make at most $m - 1$ non-maximal transitions before round $n - 1$. Given $\bar{T}^{m,0}$ and $\bar{T}^{m,n-1}$, we thus have that

$$\bar{T}^{m,n} := \bigcup_{s \in B(\bar{T}^{m,n-1}); s^{n-1} \in (\bar{T}^{m,0})^{n-1}} \bar{\mathcal{S}}^n(s^n), \quad (\text{B.14})$$

where $(\bar{T}^{m,0})^{n-1} := \{s^{n-1} : s \in \bar{T}^{m,0}\}$. □

Claim 6 There exists $N \in \mathbb{N}$ so that

$$\bigcup_{m=1}^N \bar{T}^{m,0} = \mathcal{S}. \quad (\text{B.15})$$

Proof. We will argue that for the infinite sequence of sets $(\bar{T}^{m,0})_{m \in \mathbb{N}}$ obtained from the recursive construction above, there is $N \leq |\mathcal{B}|$ so that

$$\mathcal{T}^N := \bigcup_{m=1}^N \bar{T}^{m,0} = \mathcal{S}. \quad (\text{B.16})$$

Since $\bar{T}^{1,0} \subseteq \mathcal{S}$, we conclude that $\mathcal{T}^w \subseteq \mathcal{S}$ for all w . Fix any $s \in \mathcal{S}$. Note that s makes at most $|\mathcal{B}|$ many non-maximal transitions: The set of rounds $m \in \mathbb{N}$ so that $s \notin \bar{\mathcal{S}}^m(s^m)$ is at most $|\mathcal{B}|$, where $\bar{\mathcal{S}}^m(s^m)$ is the set of sequences $\tilde{s} \in \mathcal{S}$ which are maximal at round m and satisfy $\tilde{s}^{m-1} = s^{m-1}$. We now show by induction on n that

$$s^n \in \{\bar{s}^n : \bar{s} \in \mathcal{T}^w\}, \quad (\text{B.17})$$

for every $w \geq |\mathcal{B}|$. Note that by construction of \mathcal{S} , condition (B.17) holds for $n = 0$. Suppose then that we have shown condition (B.17) for all $s \in \mathcal{S}$ and some $n = n' - 1$. Then there exists $m \leq |\mathcal{B}|$ so that

$$s^{n'-1} \in \{\bar{s}^{n'-1} : \bar{s} \in \bar{T}^{m,n'-1}\}. \quad (\text{B.18})$$

By the construction of \mathcal{S} , we conclude that

$$s^{n'} \in \{\bar{s}^{n'} : \bar{s} \in B(\bar{T}^{m,n'-1})\}, \quad (\text{B.19})$$

and so there is $m' \leq |\mathcal{B}|$ so that

$$s^{n'} \in \{\bar{s}^{n'} : \bar{s} \in \bar{T}^{m',n'}\}. \quad (\text{B.20})$$

We deduce that for all $w \geq |\mathcal{B}|$ we must have that $s \in \mathcal{T}^w$ and so the result follows. □

Claim 7 Let N satisfy (B.15). For every $m < N$ there exists $z_m \in \mathbb{N}$ and $n_m \in \mathbb{N}$ so that for all $s \in \bar{T}^{m+1,0}$

$$O(s^n) = O(s^{n+z_m}), \quad \forall n \geq n_m. \quad (\text{B.21})$$

Proof. By monotonicity of B we have that for every $n \in \mathbb{N}$,

$$\bar{S}^n \subseteq B(\bar{S}^{n-1}). \quad (\text{B.22})$$

Hence $\bar{T}^{m,n} \subseteq B(\bar{T}^{m,n-1})$. For any two rounds $n > \tilde{n}$ so that

$$\{O(s^l) : s \in \bar{T}^{m,l}\} = \{O(s^{l+n-\tilde{n}}) : s \in \bar{T}^{m,l+n-\tilde{n}}\}, \quad (\text{B.23})$$

for every $l \in \{\tilde{n} - M_m, \dots, \tilde{n}\}$, we also have that

$$\{O(s^{\tilde{n}+1}) : s \in \bar{T}^{m,\tilde{n}+1}\} = \{O(s^{n+1}) : s \in \bar{T}^{m,n+1}\}. \quad (\text{B.24})$$

We conclude from both corollaries that for every $m \in \mathbb{N}$, there exists $n_m \in \mathbb{N}$ and $\tilde{n}_m < n_m$ so that (B.23) holds. Hence the result follows. \square

Theorem 3 Ω is a finite set.

Proof. Follows from Claims 6 and 7. \square

Appendix C

Proofs Chapter 3

C.1 Appendix

The set of nonnegative integers is denoted \mathbb{N} and the set of positive integers is denoted $\mathbb{N}_{>0}$. For any family of sets $(X_i)_i$, we let $X = \prod_i X_i$ and $X_{-i} = \prod_{j \neq i} X_j$, for every i . For a family of maps $f_i: X_i \rightarrow Y_i$, we let $f: X \rightarrow Y$ be given by $f(x) = (f_i(x_i))_i$ for $x \in X$ and $f_{-i}: X_{-i} \rightarrow Y_{-i}$ by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$ for $x_{-i} \in X_{-i}$. Given a measurable set X , $\Delta(X)$ denotes the set of probability distributions on X . A marginal on coordinates x_1, \dots, x_n of a distribution $P \in \Delta(\prod_\ell X_\ell)$ is denoted $\text{marg}_{x_1, \dots, x_n}(P)$.

We use the following notation. For every player i let $\mathcal{A}_i = 2^{A_i} \setminus \{\emptyset\}$ denote the set of non-empty action sets of player i .

For any subset of action set profiles $B' \subseteq B$, let $\max_{\subseteq} B' \subseteq B'$ denote the collection of maximal profiles in B' with respect to set inclusion.

C.1.1 Theorem 5: Revelation Principle

Theorem 5 For $P \in \Delta(K \times B^{\mathbb{N}})$ the three conditions are equivalent:

1. $P \in \mathcal{P}$,
2. $P = P_S$, i.e. P is the image of itself viewed as a common prior,
3. P satisfies $P(s^0 = A) = 1$ and the family of obedience constraints:

$$s_i^m = \text{BR}_i(\text{marg}_{k, s_{-i}^{m-1}}(P(\cdot, \cdot | s_i))) \text{ for a.e. } s_i = (s_i^m)_m. \quad (\text{C.1})$$

Proof. (1. \implies 3) Let $P \in \Delta(K \times T)$ be a common prior. The induced profile of conditional probabilities $(P_i: T_i \rightarrow \Delta(K \times T_{-i}))_i$ is a Harsanyi type space and so the profile of maps $((R_i^m)_m: T_i \rightarrow S_i)_i$ satisfies: for every player i , $t_i \in T_i$ and $m \in \mathbb{N}$,

$$R_i^m(t_i) = \text{BR}_i(P_i(t_i) \circ (\text{id} \times R_{-i}^{m-1})^{-1}). \quad (\text{C.2})$$

Then note that for every $B_i \in \mathcal{A}_i$, $\text{BR}_i(B_i)^{-1} \subseteq \Delta(K \times \mathcal{A}_{-i})$ is convex. Write $\tilde{P} := P \circ (\text{id} \times R)^{-1}$ with conditional probabilities $(\tilde{P}_i: S_i \rightarrow \Delta(K \times S_{-i}))_i$ and so for all $k \in K$,

$m \in \mathbb{N}$ and $s \in S$ so that $P(R_i^{-1}(s_i)) > 0$,

$$\begin{aligned} s_i^m &= \text{BR}_i \left(\int_{\{t_i: R_i^m(t_i) = s_i^m\}} P_i(t_i) \circ (\text{id} \times R_{-i}^{m-1})^{-1} dP(t_i | R_i^m = s_i^m) \right) \\ &= \text{BR}_i \left(\tilde{P}_i(s_i) \right). \end{aligned} \quad (\text{C.3})$$

So \tilde{P} satisfies (C.1), as required. (3. \implies 2) If $P \in \Delta(K \times B^{\mathbb{N}})$ satisfies (C.1) then, in particular it is a common prior with type profiles given by S . Since ICR is characterized exactly by (3.3), we deduce that $P = P_S$. (2. \implies 1) $P = P_S$ implies that $P \in \mathcal{P}$, which concludes the proof. \square

C.1.2 Theorem 6: Sufficiency of SCAMP

Monotonicity Property of BR

A *monotone stochastic transformation* for player i is a map $\rho_i: K \times \mathcal{A}_{-i} \rightarrow \Delta(\mathcal{A}_{-i})$ so that for every $b \in B$ and $k \in K$,

$$b'_{-i} \subseteq B_{-i}, \quad \forall B'_{-i} \in \text{supp}(\rho_i(k, B_{-i})). \quad (\text{C.4})$$

Claim 11 (Monotonicity of BR). *For any monotone stochastic transformation $\rho_i: K \times \mathcal{A}_{-i} \rightarrow \Delta(\mathcal{A}_{-i})$ and for any $p_i \in \Delta(K \times \mathcal{A}_{-i})$,*

$$\text{BR}_i(p_i \circ \rho_i) \subseteq \text{BR}_i(p_i), \quad (\text{C.5})$$

where for all $k \in K$ and $B_{-i} \in \mathcal{A}_{-i}$,

$$p_i \circ \rho_i(k, B_{-i}) := \sum_{B'_{-i} \in \mathcal{A}_{-i}} \rho_i(B_{-i} | k, B'_{-i}) p_i(k, B'_{-i}). \quad (\text{C.6})$$

Proof. Consider any conjecture $\sigma_i: K \times \mathcal{A}_{-i} \rightarrow \Delta(\mathcal{A}_{-i})$ so that $\text{supp}(\sigma_i(\cdot | k, B_{-i})) \subseteq B_{-i}$ for all $k \in K, B_{-i} \in \mathcal{A}_{-i}$. Now define the conjecture $\sigma_i \circ \rho_i$, which for every $a_{-i} \in \mathcal{A}_{-i}, k \in K, B'_{-i} \in \mathcal{A}_{-i}$ is given by

$$\sigma_i \circ \rho_i(a_{-i} | k, B'_{-i}) := \sum_{B_{-i}} \sigma_i(a_{-i} | k, B_{-i}) \rho_i(B_{-i} | k, B'_{-i}). \quad (\text{C.7})$$

Since ρ_i is monotone, the conjecture $\sigma_i \circ \rho_i$ also satisfies the support constraint of σ_i . Hence

$$\begin{aligned} \langle \sigma_i, p_i \circ \rho_i \rangle(k, a_{-i}) &= \sum_{B'_{-i} \in \mathcal{A}_{-i}} \left(\sum_{B_{-i} \in \mathcal{A}_{-i}} \sigma_i(a_{-i} | k, B_{-i}) \rho_i(B_{-i} | k, B'_{-i}) \right) p_i(k, B'_{-i}) \\ &= \sum_{B'_{-i} \in \mathcal{A}_{-i}} \sigma_i \circ \rho_i(a_{-i} | k, B'_{-i}) p_i(k, B'_{-i}) \\ &= \langle \sigma_i \circ \rho_i, p_i \rangle(k, a_{-i}). \end{aligned} \quad (\text{C.8})$$

Now the result is immediate from the definition of BR_i in expression (3.2). \square

The set of conjectures on automaton states Σ_B is the set of random selections $\sigma: K \times \mathcal{A}_{-i} \rightarrow \Delta(A_{-i})$, where for every k, b , $\sigma(B_{-i}|k, B_{-i}) = 1$. For every $(k, s) \in K \times S$, player i and round m , define the minimal and maximal conjectured payoff increments between $a_i, a'_i \in A_i$ respectively as,

$$\begin{aligned}\underline{u}_{i,a_i,a'_i}^m(k, s) &:= \min_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s_{-i}^m} \sigma(a_{-i}|k, s^m)(u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})) \\ \bar{u}_{i,a_i,a'_i}^m(k, s) &:= \max_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s_{-i}^m} \sigma(a_{-i}|k, s^m)(u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})).\end{aligned}$$

For every player i and $s \in S_P$ define i 's information set

$$S_{i,P}(v_i) := \{\tilde{s} \in S_P : \tilde{s}_i = s_i\}. \quad (\text{C.9})$$

The expected payoff increments are denoted

$$\begin{aligned}\underline{U}_{i,P}^m(s_i, a_i, a'_i) &:= \sum_{(k, \hat{s}) \in K \times S_{i,P}(s_i)} \underline{u}_{i,a_i,a'_i}^m(k, \hat{s})P(k, \hat{s}), \\ \bar{U}_{i,P}^m(s_i, a_i, a'_i) &:= \sum_{(k, \hat{s}) \in K \times S_{i,P}(s_i)} \bar{u}_{i,a_i,a'_i}^m(k, \hat{s})P(k, \hat{s}).\end{aligned} \quad (\text{C.10})$$

OCs can be grouped into two sets of constraints:

- (i) *Sub-obedience*: $\forall a'_i \in s_i^m \setminus s_i^{m+1}, \exists a_i \in s_i^{m+1}$ s.t. $0 < \underline{U}_i^m(s_i, a_i, a'_i)$,
- (ii) *Super-obedience*: $\forall a_i \in s_i^{m+1}, \forall a'_i \in A_i, 0 \leq \bar{U}_i^m(s_i, a_i, a'_i)$.

Lemma 10 (Semi-lattice Priors)

For any subset $V \subseteq K \times B^{\mathbb{N}}$ and any player i , define player i 's sequences

$$V_i := \{s_i : \exists k \in K \text{ s.t. } (k, s) \in V\}. \quad (\text{C.11})$$

Define the information set of a sequence $s_i \in V_i$,

$$V_i(s_i) := \{(k, \hat{s}) \in V : \hat{s}_i = s_i\}. \quad (\text{C.12})$$

Let $k \in K, s, s' \in S$ satisfy $(k, s), (k, s') \in V$ and $\{s_i^m : m \in \mathbb{N}\} = \{s_i'^m : m \in \mathbb{N}\}$ for every player i . Then for every players i, j , every $m \in \mathbb{N}$ and $(\hat{k}, \hat{s}) \in V_i(s_i) \cup V_i(s'_i)$, define

$$\xi_j^m(\hat{k}, \hat{s}|s_i, s'_i) := \left(\hat{k}, \min_{\subseteq} \{\tilde{s}_j^m : (\hat{k}, \tilde{s}) \in V_i(s_i) \cup V_i(s'_i), \{\tilde{s}_j^n : n \in \mathbb{N}\} = \{\hat{s}_j^n : n \in \mathbb{N}\}\} \right). \quad (\text{C.13})$$

Then say that V is *closed from below*, if

$$\bigcup_i \left\{ \left(\xi_j^m(\hat{k}, \hat{s}|s_i, s'_i) \right)_{j \in I, m \in \mathbb{N}} : (\hat{k}, \hat{s}) \in V_i(s_i) \cup V_i(s'_i) \right\} \subseteq V. \quad (\text{C.14})$$

Note that if V is closed from below then in particular:

- (i) for every $(k, v), (k', v') \in V$ satisfying $k = k'$ and $\{v_i^m : m \in \mathbb{N}\} = \{v_i'^m : m \in \mathbb{N}\}$ for every player i , there exists $(k'', v'') \in V$ so that $k'' = k' = k$ and

$$v'' \subseteq_S v \text{ and } v'' \subseteq_S v', \quad (\text{C.15})$$

- (ii) for every $(k, v), (k', v') \in V$ satisfying $k = k'$ and $\{v^m : m \in \mathbb{N}\} = \{v'^m : m \in \mathbb{N}\}$, there exists $(k'', v'') \in V$ so that $k'' = k' = k$ and

$$v'' \subseteq_S v \text{ and } v'' \subseteq_S v'. \quad (\text{C.16})$$

Fix a subset $X_0 \subseteq K \times B$ and let $\mathcal{V}(X_0)$ denote the collection of subsets $V \subseteq K \times S$ which are (1) closed from below and (2) whose sequences have limits that are consistent with X_0 in the following sense

$$\mathcal{V}(X_0) := \left\{ V \subseteq K \times S : \left\{ \left(k, \lim_{m \uparrow \infty} s^m \right) : (k, s) \in V \right\} = X_0 \right\}. \quad (\text{C.17})$$

For every $V, V' \in \mathcal{V}(X_0)$, say that a profile of mappings $(\phi_i : V_i \rightarrow V'_i)_i$ is *monotonic* if for every $v' \in V'$, every player i and every $m \in \mathbb{N}$, there exists $v_i \in \phi_i^{-1}(v'_i)$ so that

$$v_i'^m \subseteq v_i^m. \quad (\text{C.18})$$

We define a partial order \ll on $\mathcal{V}(X_0)$, where for every $V, V' \in \mathcal{V}(X_0)$,

$$V \ll V' \iff \exists \text{ monotonic } (\phi_i : V_i \rightarrow V'_i)_i. \quad (\text{C.19})$$

For any X_0 , $(\ll, \mathcal{V}(X_0))$ is a semi-lattice: For any $V, V' \in \mathcal{V}(X_0)$, we have that:

- (i) $V \cup V' \in \mathcal{V}(X_0)$,
- (ii) $V \cup V'' \ll V$ and $V \cup V' \ll V'$.

The set of outcome distributions that are consistent with X_0 :

$$\mathcal{O}(X_0) := \bigcup_{V \in \mathcal{V}(X_0)} \{\nu_P : P \in \mathcal{P} \cap \Delta(V)\}. \quad (\text{C.20})$$

Now, fix a distribution $\nu \in \mathcal{O}(X_0)$ and a set $V \in \mathcal{V}(X_0)$. Define the collection $\mathcal{P}_\nu^1 \subseteq \Delta(V)$ of distributions

$$\mathcal{P}_\nu^1(V) := \left\{ P \in \Delta(V) : \forall i, \text{BR}_i(\text{marg}_{k, s_{-i}^0}(P(\cdot, \cdot | s_i))) = s_i^1, \text{ a.s.} \right\}. \quad (\text{C.21})$$

Moreover, define the following best-response operator, on sets $V \in \mathcal{V}$,

$$\mathbb{B}_\nu(V) := \bigcup_{P \in \mathcal{P}_\nu^1(V)} \left\{ (\hat{k}, (\text{BR}_i(\text{marg}_{k, s_{-i}^{m-1}}(P(\cdot, \cdot | \hat{v}_i))))_i)_{m \in \mathbb{N}} : P(\hat{k}, \hat{v}) > 0 \right\}. \quad (\text{C.22})$$

Claim 12 (Properties of Best-Response Operator). *For every outcome distribution $\nu \in \mathcal{O}(X_0)$,*

- (i) $\mathbb{B}_\nu(V) \in \mathcal{V}(X_0)$, $\forall V \in \mathcal{V}(X_0)$,
- (ii) $\mathbb{B}_\nu(V \cup V') = \mathbb{B}_\nu(V) \cup \mathbb{B}_\nu(V')$, $\forall V, V' \in \mathcal{V}(X_0)$,
- (iii) $V \ll V' \implies \mathbb{B}_\nu(V) \ll \mathbb{B}_\nu(V')$, $\forall V, V' \in \mathcal{V}(X_0)$.

Proof. (i) This fact follows from the monotonicity properties of BR_i .

(ii) Note that for every $\bar{P} \in \Delta(V \cup V')$ there exists an associated $P \in \Delta(V)$, $P' \in \Delta(V')$ and $\alpha \in [0, 1]$ so that $\bar{P}(v) = \alpha P(v) + (1 - \alpha)P'(v)$ for all $v \in V$. Conversely, for every $P \in \Delta(V)$, $P' \in \Delta(V')$ and $\alpha \in [0, 1]$, $P := \alpha P + (1 - \alpha)P' \in \Delta(V \cup V')$. Hence the result follows.

(iii) Fix a distribution $P \in \mathcal{P}_\nu^1$ and consider any $v \in V$ and let $(\phi_i)_i$ be the profile of maps that verifies the statement $V \ll V'$. Define the induced distribution

$$\hat{P} := P \circ \left(\text{id} \times \prod_i \phi_i \right)^{-1}. \quad (\text{C.23})$$

Then for every player i ,

$$\hat{P}(k|v'_i) = \sum_{v_i \in \phi_i^{-1}(v'_i)} P(k|v_i)P(v_i|\phi_i^{-1}(v'_i)), \quad \forall k \in K. \quad (\text{C.24})$$

We conclude that $\hat{P} \in \mathcal{P}_\nu^1$. Moreover, for every $m \in \mathbb{N}$,

$$\begin{aligned} & \min_{\subseteq} \left\{ \text{BR}_i \left(\text{marg}_{k, s_{-i}^m} (P(\cdot, \cdot|v_i)) \right) : v_i \in \phi_i^{-1}(v'_i) \right\} \\ & \subseteq \text{BR}_i \left(\text{marg}_{k, s_{-i}^m} (P(\cdot, \cdot|\phi_i^{-1}(v'_i))) \right) \\ & \subseteq \max_{\subseteq} \left\{ \text{BR}_i \left(\text{marg}_{k, s_{-i}^m} (P(\cdot, \cdot|v_i)) \right) : v_i \in \phi_i^{-1}(v'_i) \right\}. \end{aligned} \quad (\text{C.25})$$

□

Claim 13 (Existence of Fixed Point). *For every $\nu \in \mathcal{O}(X_0)$, $\mathbb{B}_\nu : \mathcal{V}(X_0) \rightarrow \mathcal{V}(X_0)$ has a fixed point.*

Proof. Let $\hat{V} := \bigcup_{V \in \mathcal{V}(X_0)} V \in \mathcal{V}(X_0)$. Moreover, we have that

$$\mathbb{B}_\nu(\hat{V}) \subseteq \hat{V}. \quad (\text{C.26})$$

Hence, $\hat{V} \ll \mathbb{B}_\nu(\hat{V})$. Then by monotonicity we have that for every $m \in \mathbb{N}$,

$$\hat{V} \ll \mathbb{B}_\nu(\hat{V}) \ll \dots \ll \mathbb{B}_\nu^m(\hat{V}) \ll \dots \quad (\text{C.27})$$

Let $\bar{V} := \bigcup_{m=0}^{\infty} \mathbb{B}_\nu^m(\hat{V})$, where $\mathbb{B}_\nu^0(\hat{V}) := \hat{V}$. Then by additivity, we have that for every $n \in \mathbb{N}$,

$$\mathbb{B}_\nu^n \left(\bigcup_{m=0}^{\infty} \mathbb{B}_\nu^m(\hat{V}) \right) = \bigcup_{m=n}^{\infty} \mathbb{B}_\nu^m(\hat{V}) = \bigcap_{h=0}^n \bigcup_{m=h}^{\infty} \mathbb{B}_\nu^m(\hat{V}). \quad (\text{C.28})$$

Since $\mathbb{B}_\nu(V) \neq \emptyset$, for all $V \in \mathcal{V}(X_0)$, we have

$$V^* := \bigcap_{n=0}^{\infty} \bigcup_{m=n}^{\infty} \mathbb{B}_\nu^m(\hat{V}) \in \mathcal{V}(X_0). \quad (\text{C.29})$$

But then we must have that

$$\mathbb{B}_\nu(V^*) = V^*. \quad (\text{C.30})$$

□

The following then is a corollary of our revelation principle (Theorem 5) and establishes the result in Lemma 10:

Corollary 6. *If $V^* \in \mathcal{V}(X_0)$ is a fixed point of \mathbb{B}_ν for some $\nu \in \mathcal{O}(X_0)$, then there exists $P \in \mathcal{P}^*$ satisfying $P(V^*) = 1$ and $\nu_P = \nu$.*

Outcome Distributions

Lemma 16. *The relative closure of the set \mathcal{O} is a convex polyhedron.*

Proof. Clearly, the relative closure of \mathcal{O}^{m^*} and that of \mathcal{O}^∞ are convex polyhedra. Consider the un-scaled limit measure \bar{p}^* ,

$$\bar{p}^*(k, b) := \sum_{(\omega^0, \dots, \omega^{m^*}) \in X^* : \beta(\omega^{m^*}) = b} p^*(k, \omega). \quad (\text{C.31})$$

Then we have that the relative closure of $\{\bar{p}^* : p \in \mathcal{O}^{m^*}\}$ is a convex polyhedron. Finally, note that the collection $\{p^{m^*}(X^*) : p \in \mathcal{O}^{m^*}\} \subseteq [0, 1]$ is convex and thus equal to an interval $[\underline{x}, \bar{x}]$. Then we have that

$$\{\bar{p} : p \in \mathcal{O}^{m^*}\} = \left\{ \frac{1}{x} \bar{p}^* : p \in \mathcal{O}^{m^*}, x \in [\underline{x}, \bar{x}] \right\} \cap \Delta(X^*). \quad (\text{C.32})$$

Indeed, each \bar{p}^* has a unique $x \in [\underline{x}, \bar{x}]$ so that $\frac{1}{x} \bar{p}^* \in \Delta(X^*)$, we obtain \mathcal{O} as the intersection of a cone and a simplex, making it a convex polyhedron. □

Theorem 7. *\mathcal{O} coincides with the set of all Rationalizable distributions, $\mathcal{O} = \mathcal{O}^*$. Its relative closure is a convex polyhedron.*

Proof. We now show that \mathcal{O} coincides with the set of all outcome distributions. For every $p \in \mathcal{O}$ there is a distribution p^* on $K \times \Omega^{m^*}$ and a subset $X^* \subseteq \Omega^{m^*}$ so that $p = p^*|_{X^*}$. Let p^* be represented on the strategic automaton verifying Theorem 6. Then $p^*|_{X^*}$ induces a distribution on the branches of the automaton and thus a transition probability $\mu_{p^*} : K \times \Omega \rightarrow \Delta(\Omega)$. Using the arguments used to establish Theorem 6 we find a cycling probability ζ so that the Markov process induced by $\mu_{p^*}, \zeta, \text{marg}_K(p)$ is SCAMP with outcome distribution p . We thus obtain our characterization of rationalizable outcomes in Theorem 7. □

Appendix D

Proofs Chapter 4

D.1 Appendix

D.1.1 Proof of Proposition 1

The ACK topology is metrizable by the metric

$$d^{f,g}(P, P') = \inf \left\{ \varepsilon \geq 0 : \begin{array}{l} P \left(C^{1-g(\varepsilon)} \left(\hat{T}_{g(\varepsilon)}(P, P') \right) \right) > f(\varepsilon) \\ P' \left(C^{1-g(\varepsilon)} \left(\hat{T}_{g(\varepsilon)}(P, P') \right) \right) > f(\varepsilon) \end{array} \right\}.$$

where $g(\varepsilon) := \varepsilon^2$ and $f(\varepsilon) := 1 - \frac{1}{2}\varepsilon(\varepsilon + 1)$.

Proof. First note that $d^{f,g}(P, P) = 0$ and $d^{f,g}(P, P') = d^{f,g}(P', P)$ are both immediate. Suppose now that $d^{f,g}(P, P') = 0$. Then we have that $\text{supp}(P) = \text{supp}(P')$, which under the natural quotient of \mathcal{P} means that $P = P'$. It remains to show that $d^{f,g}$ satisfies the triangle inequality. Let $P^1, P^2, P^3 \in \mathcal{P}$ satisfy $d^{f,g}(P^1, P^2) < \varepsilon_{1,2}$ and $d^{f,g}(P^3, P^2) < \varepsilon_{3,2}$, for $\varepsilon_{1,2}, \varepsilon_{3,2} \in [0, 1]$. Note that for every $\omega \in \hat{T}_{g(\varepsilon_{1,2})}(P^1, P^2) \cap \hat{T}_{g(\varepsilon_{3,2})}(P^3, P^2) \subseteq \text{supp}_{g(\varepsilon_{1,2})}(P^1) \cap \text{supp}_{g(\varepsilon_{3,2})}(P^3)$ there is $\omega' \in \text{supp}(P^2)$ so that

$$d_{\Pi}(\omega, \omega') \leq \underline{\varepsilon} := \min \{g(\varepsilon_{1,2}), g(\varepsilon_{3,2})\}.$$

Define the event F as follows:

$$F := C^{1-g(\varepsilon_{1,2})} \left(\hat{T}_{g(\varepsilon_{1,2})}(P^1, P^2) \right) \cap C^{1-g(\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{3,2})}(P^3, P^2) \right).$$

which can be rewritten as

$$F = \left(\left(C^{1-g(\varepsilon_{1,2})} \left(\hat{T}_{g(\varepsilon_{1,2})}(P^1, P^2) \right) \right)^{\complement} \cup \left(C^{1-g(\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{3,2})}(P^3, P^2) \right) \right)^{\complement} \right)^{\complement}$$

Hence

$$\begin{aligned} P^2(F) &= 1 - P^2 \left(\left(C^{1-g(\varepsilon_{1,2})} \left(\hat{T}_{g(\varepsilon_{1,2})}(P^1, P^2) \right) \right)^{\complement} \cup \left(C^{1-g(\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{3,2})}(P^3, P^2) \right) \right)^{\complement} \right) \\ &\geq 1 - \left(1 - P^2 \left(C^{1-g(\varepsilon_{1,2})} \left(\hat{T}_{g(\varepsilon_{1,2})}(P^1, P^2) \right) \right) + 1 - P^2 \left(C^{1-g(\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{3,2})}(P^3, P^2) \right) \right) \right) \\ &> 1 - (1 - f(\varepsilon_{1,2}) + 1 - f(\varepsilon_{3,2})). \end{aligned}$$

Using the fact that $g(\varepsilon_{1,2}) + g(\varepsilon_{3,2}) = \varepsilon_{1,2}^2 + \varepsilon_{3,2}^2 \leq (\varepsilon_{1,2} + \varepsilon_{3,2})^2 = g(\varepsilon_{1,2} + \varepsilon_{3,2})$, we also have that $F \subseteq C^{1-g(\varepsilon_{1,2}+\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{1,2}+\varepsilon_{3,2})}(P^1, P^3) \right)$. Letting $F_i := \text{proj}_{\mathcal{T}_i}(F)$, it can be shown that

$$\begin{aligned} |P^k(F) - P^2(F)| &= \left| \int_{F_i} \tau_i(F) P^k \circ \varphi_i^{-1}(d\tau_i) - \int_{F_i} \tau_i(F) P^2 \circ \varphi_i^{-1}(d\tau_i) \right| \\ &\leq |\text{marg}_{\mathcal{T}_i}(P^k)(F_i) - \text{marg}_{\mathcal{T}_i}(P^2)(F_i)| \\ &\leq \underline{\varepsilon}, \end{aligned}$$

where $\varphi_i: F \rightarrow F_i$ is measurable and satisfies $F = \{(\varphi_i(\omega), \text{proj}_{\mathcal{T}_{-i} \times \Theta}(\omega)) : \omega \in F\}$. Hence $P^k(F) \geq P^2(F) - \underline{\varepsilon} > f(\varepsilon_{1,2}) + f(\varepsilon_{3,2}) - 1 - \underline{\varepsilon}$. Then for any $\varepsilon_{1,2}, \varepsilon_{3,2} \in [0, 1]$,

$$P^k \left(C^{1-g(\varepsilon_{1,2}+\varepsilon_{3,2})} \left(\hat{T}_{g(\varepsilon_{1,2}+\varepsilon_{3,2})}(P^1, P^3) \right) \right) > f(\varepsilon_{1,2}) + f(\varepsilon_{3,2}) - 1 - \underline{\varepsilon} \quad ,$$

It remains to show that $f(\varepsilon_{1,2}) + f(\varepsilon_{3,2}) - \underline{\varepsilon} - 1 \geq f(\varepsilon_{1,2} + \varepsilon_{3,2})$. To show this, we start out with the following observation:

$$\underline{\varepsilon} = \min_{k \in \{1,3\}} \varepsilon_{k,2}^2 \leq \varepsilon_{1,2} \varepsilon_{3,2}.$$

Then we can do the following transformations to obtain the desired inequality:

$$\begin{aligned} 2\underline{\varepsilon} &\leq 2\varepsilon_{1,2}\varepsilon_{3,2} \\ (\varepsilon_{1,2}^2 + \varepsilon_{3,2}^2 + \varepsilon_{1,2} + \varepsilon_{3,2}) + 2\underline{\varepsilon} &\leq (\varepsilon_{1,2}^2 + \varepsilon_{3,2}^2 + \varepsilon_{1,2} + \varepsilon_{3,2}) + 2\varepsilon_{1,2}\varepsilon_{3,2} \\ (\varepsilon_{1,2}^2 + \varepsilon_{3,2}^2 + \varepsilon_{1,2} + \varepsilon_{3,2}) + 2\underline{\varepsilon} &\leq ((\varepsilon_{1,2} + \varepsilon_{3,2})^2 + \varepsilon_{1,2} + \varepsilon_{3,2}) \\ \varepsilon_{1,2}(\varepsilon_{1,2} + 1) + \varepsilon_{3,2}(\varepsilon_{3,2} + 1) + 2\underline{\varepsilon} &\leq (\varepsilon_{1,2} + \varepsilon_{3,2})(\varepsilon_{1,2} + \varepsilon_{3,2} + 1) \\ -\frac{1}{2}\varepsilon_{1,2}(\varepsilon_{1,2} + 1) - \frac{1}{2}\varepsilon_{3,2}(\varepsilon_{3,2} + 1) - \underline{\varepsilon} &\geq -\frac{1}{2}(\varepsilon_{1,2} + \varepsilon_{3,2})(\varepsilon_{1,2} + \varepsilon_{3,2} + 1) \\ 1 - \frac{1}{2}\varepsilon_{1,2}(\varepsilon_{1,2} + 1) + 1 - \frac{1}{2}\varepsilon_{3,2}(\varepsilon_{3,2} + 1) - \underline{\varepsilon} - 1 &\geq 1 - \frac{1}{2}(\varepsilon_{1,2} + \varepsilon_{3,2})(\varepsilon_{1,2} + \varepsilon_{3,2} + 1) \\ f(\varepsilon_{1,2}) + f(\varepsilon_{3,2}) - \underline{\varepsilon} - 1 &\geq f(\varepsilon_{1,2} + \varepsilon_{3,2}). \end{aligned}$$

Hence $d^{f,g}(P^1, P^3) < \varepsilon_{1,2} + \varepsilon_{3,2}$ and so $d^{f,g}$ satisfies the triangle inequality. \square

D.1.2 Proof of Proposition 2

In the set-up of [Stinchcombe \(2011\)](#) a prior is a countably additive probability $P \in \Delta(\prod_i \Omega_i)$, where each Ω_i is endowed with a sigma algebra \mathcal{F}_i and P is defined on a sigma algebra containing the product sigma algebra. Payoffs are given by $U_i: \prod_i \Omega_i \rightarrow \mathbb{R}^A$. It is assumed that $\int_{\Omega} \|U_i(\omega)\|_{\infty} dP(\omega) < \infty$. A behavior strategy is a \mathcal{F}_i -measurable map $b_i: \Omega_i \rightarrow \Delta(A_i)$. Let \mathbb{B}_i be the set of behavior strategies and let $\mathbb{B} := \prod_i \mathbb{B}_i$. A measure $\nu \in \Delta(\prod_i \mathbb{B}_i)$ is a correlated equilibrium¹ if for every i and measurable deviation $\gamma_i: \Omega_i \times \Delta(A_i) \rightarrow \Delta(A_i)$,

$$\int_{\mathbb{B}} \int_{\Omega} \left(\sum_{a \in A} U_i(\omega)(a) \prod_i b_i(a|\omega) - \sum_{a \in A} U_i(\omega)(a) \cdot \gamma_i(\omega, b_i(a|\omega)) \cdot \prod_{j \neq i} b_j(a|\omega) \right) dP(\omega) d\nu(b) \geq 0.$$

¹In [Stinchcombe \(2011\)](#) this corresponds to a variation Strategy Correlated Equilibrium where deviation strategies depend action recommendations at realized types only.

We show that for every correlated equilibrium (P, ν) there is an outcome equivalent equilibrium (P, σ) . We now translate this set-up into ours for any choice of general information structure \mathcal{S} : Ω_i must correspond to $\Theta \times T_i$, where \mathcal{F}_i is generated by the projection of $\Theta \times T_i$ onto T_i . Payoffs translate into our set-up by setting $U_i(\theta, \tau)(a) = u_i(\theta, a)$. The condition that $\int_{\Omega} \|U_i(\omega)\|_{\infty} dP(\omega) < \infty$ is then satisfied since Θ is finite and payoffs depend Θ and not T . A behavior strategy corresponds to a marginal decision rule, $\sigma(a_i|t_i)$. It remains to show that a distribution ν over profiles of marginal decision rules as in [Stinchcombe \(2011\)](#) is equivalent to a belief invariant decision rule. First, note that every measure $\nu \in \Delta(\prod_i \mathbb{B}_i)$ induces a belief invariant decision rule σ_{ν} where for every measurable event $E \subseteq A$

$$\sigma_{\nu}(a|\omega) = \int_{\mathbb{B}} \prod_i b_i(\tau_i)(a_i) d\nu(b).$$

Indeed, the resulting marginal probability $\sigma_{\nu}(a_i|\omega) = \int_{\mathbb{B}_i} b_i(\tau_i)(a_i) d\nu(b_i)$ verifies belief invariance.

D.1.3 Proof of Proposition 9

Let (\mathcal{G}, P) be base game and an information structure, i.e. so that $P \in \mathcal{P}$. Consider an information structure $\hat{\mathcal{S}} = ((S_i, \hat{P}_i)_{i \in I}, \hat{P})$ satisfying $P = P_{\hat{\mathcal{S}}}$ and let σ' be a BIBCE of $(\mathcal{G}, \hat{\mathcal{S}})$. We construct an outcome equivalent BIBCE σ of $(\mathcal{G}, \hat{\mathcal{S}})$. For every $\theta \in \Theta$ and P -almost every $\tau \in \mathcal{T}$, let

$$\sigma(a|\theta, \tau) = \int_{\bar{\tau}^{-1}(\tau)} \sigma'(a|\theta, s') d\hat{P}(s'|\theta, \bar{\tau}^{-1}(\tau)),$$

where $\tau \mapsto \hat{P}(\cdot|\theta, \bar{\tau}^{-1}(\tau)) \in \Delta(\bar{\tau}^{-1}(\tau))$ is a conditional probability on S . Then obedience constraints for any P -almost every type $\tau_i \in \mathcal{T}_i$ are satisfied

$$\begin{aligned} \sum_{\theta, a_{-i}} (u_i(a, \theta) - u_i(a'_i, a_{-i}, \theta)) \int_{S_{-i}} \sigma(a|\theta, \tau) dP(\tau_{-i}, \theta|\tau_i) = \\ \sum_{\theta, a_{-i}} (u_i(a, \theta) - u_i(a'_i, a_{-i}, \theta)) \int_{S_{-i}} \int_{\bar{\tau}^{-1}(\tau)} \sigma'(a|\theta, s') d\hat{P}(s'|\theta, \bar{\tau}^{-1}(\tau)) dP(\tau_{-i}, \theta|\tau_i). \end{aligned}$$

Noting that

$$\int_{\mathcal{T}_{-i}} \int_{\bar{\tau}^{-1}(\tau)} \sigma'(a|\theta, s') d\hat{P}(s'|\theta, \bar{\tau}^{-1}(\tau)) dP(\tau_{-i}, \theta|\tau_i) = \int_{S_{-i}} \sigma'(a|\theta, s') d\hat{P}(s'_{-i}, \theta|s_i),$$

we conclude that σ is a BIBCE of (\mathcal{G}, P) . Conversely, note that every BIBCE σ of (\mathcal{G}, P) induces a $id \times \bar{\tau}$ -measurable decision rule σ' , where for every action profile $a \in A$, state $\theta \in \Theta$ and \hat{P} -almost every $s' \in S$,

$$\sigma'(a|\theta, s') = \sigma(a|\theta, \bar{\tau}(s')).$$

Performing a change of variables

$$\begin{aligned} \int_{S_{-i}} \sigma'(a|\theta, s') \, d\hat{P}(s'_{-i}, \theta|s_i) &= \int_{\mathcal{T}_{-i}} \int_{\bar{\tau}^{-1}(\tau)} \sigma'(a|\theta, s') \, d\hat{P}(s'|\theta, \bar{\tau}^{-1}(\tau)) \, dP(\tau_{-i}, \theta|\tau_i) \\ &= \int_{\mathcal{T}_{-i}} \sigma(a|\theta, \tau) \, dP(\tau_{-i}, \theta|\tau_i). \end{aligned}$$

and so obedience follows.

D.1.4 Proof of Proposition 10

Let $(\sigma_i: S_i \rightarrow \Delta(A_i))_{i \in I}$ be a BNE under $(\mathcal{G}, \mathcal{S})$. Consider the following representation of \mathcal{S} : Define the graph of $\bar{\tau}$, $\phi_{\bar{\tau}}: s \mapsto (\bar{\tau}(s), s)$ and consider the push forward probability $P^* := P \circ (id \times \phi_{\bar{\tau}})^{-1}$. Let $\sigma^R: \Theta \times \mathcal{T} \rightarrow \Delta(S)$ denote a $(id \times \bar{\tau})$ -conditional probability of P^* so that for any measurable $E \subseteq S$ and state $\theta \in \Theta$ in the support of P ,

$$P(E|\theta) = \int_{\mathcal{T}} \sigma^R(E|\theta, \tau) \, dP^*(\tau|\theta).$$

We now construct a BIBCE σ^* on $\text{marg}_{\Omega} P^*$ as follows:

$$\sigma^*(a|\theta, \tau) = \int_S \prod_{i \in I} \sigma_i(a|s_i) \, d\sigma^R(s|\theta, \tau).$$

Since σ satisfies obedience, so does σ^* :

$$\begin{aligned} \sum_{\theta, a_{-i}} (u_i(a, \theta) - u_i(a'_i, a_{-i}, \theta)) \int_{\mathcal{T}_{-i}} \sigma^*(a|\theta, \tau) \, d\text{marg}_{\Omega} P^*(\theta, \tau_{-i}|\tau_i) \\ = \sum_{\theta, a_{-i}} (u_i(a, \theta) - u_i(a'_i, a_{-i}, \theta)) \int_S \prod_{i \in I} \sigma_i(a|s_i) \, dP(\theta, s_{-i}|s_i). \end{aligned}$$

D.1.5 Proof of Lemma 1

Proof. For every precision $z \in \mathbb{N}$ define a finite grid approximation recursively: Consider the grid on first order beliefs, given by

$$A_i^{1,z} := \left\{ a \in \mathbb{R}^{\Theta} : a_{\theta} \in \left\{ \frac{n}{z} : 0 \leq n \leq z \right\}, \sum_{\theta} a_{\theta} = 1 \right\}$$

for any player i . Given a finite set $A_i^{m-1,z} \subseteq \mathcal{T}_i^{m-1}$ for every i , define

$$A_i^{m,z} := \left\{ a \in \mathbb{R}^{\prod_{n < m} A_{-i}^{n,z} \times \Theta} : \begin{array}{l} a_{a^1, \dots, a^{m-1}, \theta} \in \left\{ \frac{n}{z} : 0 \leq n \leq z \right\}, \\ \sum_{a^1, \dots, a^{m-1}, \theta} a_{a^1, \dots, a^{m-1}, \theta} = 1 \end{array} \right\}. \quad (\text{D.1})$$

Then for any $m \in \mathbb{N}$, and any $t^m \in \mathcal{T}^m$, there is $a^m \in \prod_{i \in I} A_i^{m,z}$ so that $d_{W_{eak}}^m(t^m, a^m) < 1/z$. Then there is m, z large enough so that for every, $\theta \in \Theta$, every $g \in \mathcal{T}$ satisfying $(g^1, \dots, g^m) = a^m$ and $\tau \in \mathcal{T}$ satisfying $(\tau^1, \dots, \tau^m) = t^m$,

$$\begin{aligned} d_{\Pi}((\theta, \tau), (\theta, g)) &= \sum_{n=1}^{\infty} \eta^n d_{W_{eak}}^n(\tau^1, \dots, \tau^n, g^1, \dots, g^n) \\ &= \frac{\eta}{z} \frac{1 - \eta^{m+1}}{1 - \eta} + \eta^m \frac{1}{1 - \eta} \\ &< \varepsilon. \end{aligned}$$

Let $G_\varepsilon := \prod_{i \in I} A_i^{m,z} \cap (\Omega \setminus \mathcal{N}_\varepsilon(E))$. Suppose that $d_{\Pi}(\omega, \omega_g) < \varepsilon$ for some $\omega \in \Omega$ and $\omega_g \in G_\varepsilon$. Then clearly, it must be that $\omega \notin E$, since the opposite would imply that $\omega_g \in \mathcal{N}_\varepsilon(E) \cap G_\varepsilon$. Suppose now that $\omega \in \Omega \setminus \mathcal{N}_\varepsilon(E)$ and so by the above there is $\omega_g \in G_\varepsilon$ so that $d_{\Pi}(\omega', \omega_g) < \varepsilon$. \square

D.1.6 Proof of Lemma 2 (Iterative Scoring Rule)

Proof. Fix m, z . and let $A^{m,z}$ be defined as in D.1. Following Dekel et al. (2006), for every profile $(a^1, \dots, a^m, \theta) \in A^{m,z} \times \Theta$ we define

$$\begin{aligned} u_i^m(a^1, \dots, a^m, \theta) &:= 2a_i^1(\theta) - \sum_{\hat{\theta}} \left(a_i^1(\hat{\theta}) \right)^2 \\ &\quad + \sum_{n=2}^m \left(2a_i^n(a^1, \dots, a^{n-1}, \theta) - \sum_{\hat{a}^1, \dots, \hat{a}^{n-1}, \hat{\theta}} \left(a_i^n(\hat{a}^1, \dots, \hat{a}^{n-1}, \hat{\theta}) \right)^2 \right). \end{aligned}$$

Dekel et al. (2006) show that the game where payoffs are given by u_i^m , the uniquely interim (correlated) rationalizable action profile of τ is the profile a^1, \dots, a^m which is $1/z$ -close to τ in terms of distance $d_{W_{eak}}^m$ on \mathcal{T}^m . The result then follows from Claim 1.

□

D.1.7 Common Belief Invariance

Proposition 17. (*Common Belief Invariance*) For every $P, P' \in \mathcal{P}$ and $\varepsilon > 0$,

$$C_{1-\varepsilon} \left(\hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P') \right) = \hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P').$$

Proof. Consider $\omega = (\theta, \tau) \in \Omega_P$ and $\hat{\omega} = (\hat{\theta}, \hat{\tau}) \in \Omega_{P'}$ with $d_{\mathcal{F}}(\omega, \hat{\omega}) < \varepsilon$. Then it must be that $\omega, \hat{\omega} \in \hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P')$. Since $\Omega_P \in \mathcal{F}$, the fact that $d_{\mathcal{F}}(\omega, \hat{\omega}) < \varepsilon$ implies that

$$\tau_i^*(\Omega_P) \leq \hat{\tau}_i^*(\mathcal{N}_{d_{\mathcal{F}}, \varepsilon}(\Omega_P)) + \varepsilon.$$

Since $\tau_i^*(\Omega_P) = 1$ and $\hat{\tau}_i^*(\Omega_{P'}) = 1$ we have

$$\hat{\tau}_i^*(\hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P')) = \hat{\tau}_i^*(\mathcal{N}_{d_{\mathcal{F}}, \varepsilon}(\Omega_P)) \geq 1 - \varepsilon.$$

A symmetric argument implies that $\tau_i^*(\hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P')) \geq 1 - \varepsilon$ and so $\hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P') = B^{1-\varepsilon}(\hat{T}_{d_{\mathcal{F}}, \varepsilon}(P, P'))$, which establishes the result. □

D.1.8 Proof of Proposition 14

Proof. Let d induce a nice, common belief invariant topology on Ω which is a stronger topology than the one induced by d_{Π} and we have that $\hat{T}_{d, \varepsilon}(P, P') \subseteq \hat{T}_{\varepsilon}(P, P')$. By Proposition 17 we deduce that $\hat{T}_{d, \varepsilon}(P, P') = C^{1-\varepsilon} \left(\hat{T}_{d, \varepsilon}(P, P') \right) \subseteq C^{1-\varepsilon} \left(\hat{T}_{\varepsilon}(P, P') \right)$ and so

$$\begin{aligned} \min \left\{ P(\hat{T}_{d, \varepsilon}(P, P')), P(\hat{T}_{d, \varepsilon}(P, P')) \right\} > 1 - \varepsilon &\implies d^{ACK}(P, P') < \varepsilon \\ &\implies \forall \mathcal{G}, \exists N > 0 \text{ s.t. } d^*(P, P'|\mathcal{G}) < N\varepsilon. \end{aligned}$$

We now show that for every $\varepsilon > 0$ there is $\delta \in (0, \varepsilon]$ so that $P(C^{1-\delta} \left(\hat{T}_{\delta}(P, P') \right)) > 1 - \delta \implies P(\hat{T}_{d, \varepsilon}(P, P')) > 1 - \varepsilon$. Write $\Omega_P := \text{supp}(P)$ and $\Omega_{P'} := \text{supp}(P')$. Next, suppose that $\Omega_P \cap \Omega_{P'} = \emptyset$, (otherwise the condition is trivially satisfied). For any nice metric d , Ω is a locally compact, countably generated Hausdorff topological space². In Folland (1999) (chapter 7, pages 211–233) it is shown that every Borel probability measure on such a space is Radon and thus regular³. Hence P and P' are both regular probability measures under the topology induced by d and so every (measurable) set admits a compact approximation from below. For any $\varepsilon > 0$, define

$$D_{P, \varepsilon} := \Omega_P \setminus \hat{T}_{d, \varepsilon}(P, P') \text{ and } D_{P', \varepsilon} := \Omega_{P'} \setminus \hat{T}_{d, \varepsilon}(P, P').$$

²Recall that a topological space is locally compact if every point admits a compact neighborhood (i.e. is contained in a compact set with nonempty interior). The space is countably generated if it has a countable topological base (i.e. all open sets can be obtained from a countable number of open sets via arbitrary unions and finite intersections). Finally, a space is Hausdorff if any two points can be separated by open neighborhoods.

³Recall that a Borel probability measure P is regular if for every $\varepsilon > 0$ and every measurable set E there is a compact set $K \subseteq E$ so that $P(E \setminus K) < \varepsilon$ and an open set $O \supseteq E$ so that $P(O \setminus E) < \varepsilon$.

Note that $D_{P,\varepsilon}, D_{P',\varepsilon}$ and $\hat{T}_{d,\varepsilon}(P, P')$ are all disjoint. By regularity, we have that for every $\eta \in (0, \varepsilon)$ there exist compact, measurable events $\hat{D}_{P,\varepsilon}^\eta \subseteq D_{P,\varepsilon}$ and $\hat{D}_{P',\varepsilon}^\eta \subseteq D_{P',\varepsilon}$ so that

$$P(D_{P,\varepsilon} \setminus \hat{D}_{P,\varepsilon}^\eta) < \eta \text{ and } P'(D_{P',\varepsilon} \setminus \hat{D}_{P',\varepsilon}^\eta) < \eta.$$

Moreover, by local compactness, there exists $\gamma > 0$ so that for all $\hat{\gamma} < \gamma$, the closure of $\mathcal{N}_{d,\hat{\gamma}}(\hat{D}_{P,\varepsilon}^\eta)$ and the closure of $\mathcal{N}_{d,\hat{\gamma}}(\hat{D}_{P',\varepsilon}^\eta)$, denoted $C_{\hat{\gamma}}$ and $C'_{\hat{\gamma}}$ respectively, are each compact. Moreover, it can readily be shown that for any $0 < \hat{\gamma}_1 < \hat{\gamma}_2 < \gamma$ the hulls H, H' , defined respectively as the closures of $C_{\hat{\gamma}_2} \setminus C_{\hat{\gamma}_1}$ and $C'_{\hat{\gamma}_2} \setminus C'_{\hat{\gamma}_1}$, are also compact. Since d refines the product topology, all the sets $H, H', \hat{D}_{P,\varepsilon}^\eta$ and $\hat{D}_{P',\varepsilon}^\eta$ are also compact in the product topology. Since the product topology is Hausdorff (see chapter III of [Mertens et al. \(2015\)](#)), there exists $\alpha > 0$ so that

$$\mathcal{N}_\alpha(\hat{D}_{P,\varepsilon}^\eta) \cap \mathcal{N}_\alpha(\Omega_{P'}) = \emptyset \text{ and } \mathcal{N}_\alpha(\hat{D}_{P',\varepsilon}^\eta) \cap \mathcal{N}_\alpha(\Omega_P) = \emptyset.$$

Figure D.1 below illustrates the argument.

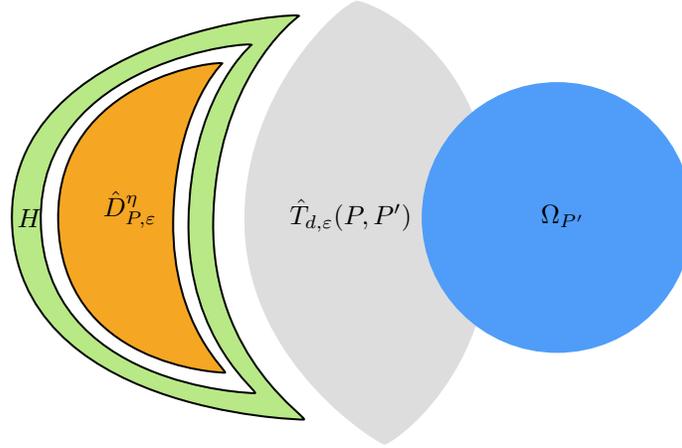


Figure D.1: Separation Argument.

Define $\delta := \varepsilon - \eta$ and suppose that $d^{ACK}(P, P') < \delta$. Then it must be that $P(\mathcal{N}_\alpha(\hat{D}_{P,\varepsilon}^\eta)) < \delta$ and $P'(\mathcal{N}_\alpha(\hat{D}_{P',\varepsilon}^\eta)) < \delta$. Since $\Omega_P \subseteq D_{P,\varepsilon} \cup \hat{T}_{d,\varepsilon}(P, P')$ and $\Omega_{P'} \subseteq D_{P',\varepsilon} \cup \hat{T}_{d,\varepsilon}(P, P')$ we have that $P(D_{P,\varepsilon}) < \delta + \eta$ and $P'(D_{P',\varepsilon}) < \delta + \eta$. Hence

$$P(\hat{T}_{d,\varepsilon}(P, P')) > 1 - \delta - \eta = 1 - \varepsilon.$$

□

D.1.9 Interim Strategic Topology

We follow [Chen et al. \(2017\)](#) in describing the interim strategic topology.

Definition 23. (Frame) A frame is a profile of maps $(\pi_i)_{i \in I}$ where for every $i \in I$, $\pi_i: \mathcal{T}_i \rightarrow F_i$, F_i is a finite set and for every prior $P \in \mathcal{P}$ and all types $\tau_i, \tau'_i \in \mathcal{T}_i$,

$$P(\cdot, \cdot | \tau_i) \circ (id \times \pi_{-i})^{-1} = P(\cdot, \cdot | \tau'_i) \circ (id \times \pi_{-i})^{-1} \implies \pi_i(\tau_i) = \pi_i(\tau'_i).$$

For any set $E \subseteq \Omega$ and any player i , let $E_i \subseteq \mathcal{T}_i$ denote the projection of E on \mathcal{T}_i . Let \mathcal{F} denote the collection of events that are measurable with respect to a frame.

$$\mathcal{F} = \{E \in \mathcal{B} : \exists \pi \in \Pi, \tau \in E \text{ s.t. } \forall i \in I, \pi_i^{-1}(\pi_i(\tau)) = E_i\}.$$

The formal definition is recursive:

Definition 24. (*Uniform Weak Distance on Frames*) For all pairs $\omega = (\theta, \tau), \hat{\omega} = (\hat{\theta}, \hat{\tau}) \in \Omega$, define the Uniform Prokhorov Distance on Frames by $d^{\mathcal{F}}(\omega, \hat{\omega}) := \sup_m d^m(\omega, \hat{\omega})$, where the sequence of functions $(d^m)_m$ on $\Omega \times \Omega$ is defined recursively from a metric representing the weak* topology on profiles of first order beliefs $d^1(\omega, \hat{\omega}) = d_{\Delta(\Theta)^I}(\tau^1, \hat{\tau}^1)$, and for every $m > 1$,

$$d^m(\omega, \hat{\omega}) := d_{\Theta}(\theta, \hat{\theta}) + \inf\{\delta > 0 : \forall i, \tau_i(E) \leq \hat{\tau}_i(\mathcal{N}_{d^{m-1}, \delta}(E)) + \delta, \forall E \in \mathcal{F}\}.$$

Definition 25. (*Interim Strategic Topology*) Define the uniform weak topology on frames as the topology on Ω generated by the sets $\{\omega' \in \Omega : d^{\mathcal{F}}(\omega, \omega') < \varepsilon\}$ for $\omega \in \Omega$ and $\varepsilon > 0$.

Proposition 18. (*Niceness*) The interim strategic topology is nice.

Proof. Dekel et al. (2006) show that a countable set of finite states is dense in Ω in the interim strategic topology and so Ω is separable under this topology. The interim strategic topology is a stronger topology than the product topology and is metrizable (see Dekel et al. (2006)). Moreover, Dekel et al. (2006) and Chen et al. (2017) provide an explicit metric that induces this topology. Moreover, Mertens et al. (2015) show that Ω is complete. Since the interim strategic topology refines the product topology, any Cauchy sequence under the metric inducing the interim strategic topology will also be Cauchy under a metric inducing the product topology. Hence Ω is a complete separable metric space under the interim strategic topology. Applying Willard (2012) (page 182)⁴ we conclude that every neighborhood of a point is totally bounded. Since a metric space is compact if and only if it is complete and totally bounded, and the closure of a neighborhood will also be its completion and be totally bounded, we conclude that every point has a compact neighborhood. Hence Ω is locally compact under the interim strategic topology, which means the topology is nice. \square

D.1.10 Upper Hemi Continuity of Exact BIBCE

Proposition 19. For every base game \mathcal{G} and any information structure $P \in \mathcal{P}$ there is a sequence $(\hat{P}^k)_k$ of simple information structures ACK-converging to P and a subset $\mathcal{O}^\infty(\mathcal{G}, P) \subseteq \mathcal{O}(\mathcal{G}, P)$ so that $\lim_{k \uparrow \infty} d_{\mathcal{H}}(\mathcal{O}(\mathcal{G}, \hat{P}^k), \mathcal{O}^\infty(\mathcal{G}, P)) = 0$.

Proof. By Proposition 6 there is a sequence of simple priors P^k converging strategically to P and a sequence of positive $(\varepsilon^k)_k$ converging monotonically to 0 so that for every k , $d^*(P^k, P) < \varepsilon^k$. Let $(\sigma^k)_k$ be a sequence of BIBCE so that for every $k \in \mathbb{N}$, $\sigma^k \in \mathcal{B}^0(\mathcal{G}, P^k)$ and the induced sequence of outcomes $(\nu_{\sigma^k})_k$ converges to some outcome ν_∞ , i.e. $\lim_{k \uparrow \infty} \|\nu_{\sigma^k} - \nu_\infty\|_2 = 0$. We have shown that for every k there is a ε^k -BIBCE $\sigma^{\varepsilon^k} \in \mathcal{B}^{\varepsilon^k}(\mathcal{G}, P)$ so that $\|\nu_{\sigma^k} - \nu_{\sigma^{\varepsilon^k}}\|_2 \leq 4M|A \times \Theta|\varepsilon^k$. Our construction in the proof of Proposition 4 has the

⁴This result states that a metric space is separable if and only if it is homeomorphic to a totally bounded metric space.

feature that the sequence $(\sigma^{\varepsilon^k})_k$ converges almost surely to an obedient decision rule (as it was constructed from a conditional expectation) and so there is a BIBCE $\sigma^\infty \in \mathcal{B}(\mathcal{G}, P)$ so that $\nu_{\sigma^\infty} = \nu_\infty$, which establishes upper hemi-continuity: There is a subset $\mathcal{O}^\infty(\mathcal{G}, P) \subseteq \mathcal{O}(\mathcal{G}, P)$ so that

$$\lim_{k \uparrow \infty} d_{\mathcal{H}, \mathcal{G}}(\mathcal{O}(\mathcal{G}, P^k), \mathcal{O}^\infty(\mathcal{G}, P)) = 0.$$

□

D.1.11 Example Existence of BIBCE

We consider the game by [Hellman \(2014\)](#) where BNE fails to exist and yet construct a straightforward BIBCE of this game. There are two players A and B and two actions L, R (see [Hellman \(2014\)](#) for the payoff matrix)

$$\sigma(y, \theta) = \begin{cases} \frac{1}{2}(L, L), \frac{1}{2}(R, R) & , \text{ if } \theta \in \{(A, 1), (B, 1)\} \\ \frac{1}{2}(R, L), \frac{1}{2}(L, R) & , \text{ if } \theta \in \{(A, -1), (B, -1)\} \end{cases}$$

This decision rule is independent of the type space and satisfies belief invariance: For any $a_i \in \{L, R\}$ and any $\theta \in \{(A, 1), (B, 1), (A, -1), (B, -1)\}$

$$\sum_{a_{-i} \in \{L, R\}} \sigma(a_i, a_{-i}|y, \theta) = \sigma(a_i, L|y, \theta) + \sigma(a_i, R|y, \theta) = \frac{1}{2}.$$

We do not need to report the full information structure considered in [Hellman \(2014\)](#), to verify obedience as σ is independent of types. Each player $i \in \{A, B\}$ has one of two first-order beliefs $P_i(\cdot|y)$ indexed by $y \in \{-1, 1\}$

$$P_A(\theta|y) = \begin{cases} \frac{1}{2} & \text{if } \theta = (A, y) \\ \frac{1}{4} & \text{if } \theta = (B, 1) \\ \frac{1}{4} & \text{if } \theta = (B, -1) \end{cases}, \quad P_B(\theta|y) = \begin{cases} \frac{1}{2} & \text{if } \theta = (B, y) \\ \frac{1}{4} & \text{if } \theta = (A, 1) \\ \frac{1}{4} & \text{if } \theta = (A, -1) \end{cases}.$$

So if recommended $a_i \in \{L, R\}$ player i 's payoff increment from playing a different action $a'_i \neq a_i$ is given by when $i = A$:

For the case where $y = 1$ and $a_A = L$

$$\begin{aligned} & \frac{1}{2} (u_A(L, L, (A, 1)) - u_A(R, L, (A, 1))) + \frac{1}{4} (u_A(L, L, (B, 1)) - u_A(R, L, (B, 1))) \\ & + \frac{1}{4} (u_A(L, R, (B, -1)) - u_A(R, R, (B, -1))) = 0.35 > 0. \end{aligned}$$

For the case where $y = 1$ and $a_A = R$

$$\begin{aligned} & \frac{1}{2} (u_A(R, R, (A, 1)) - u_A(L, R, (A, 1))) + \frac{1}{4} (u_A(R, R, (B, 1)) - u_A(L, R, (B, 1))) \\ & + \frac{1}{4} (u_A(R, L, (B, -1)) - u_A(L, L, (B, -1))) = 0.15 > 0. \end{aligned}$$

For the case $y = -1$ and $a_A = L$,

$$\begin{aligned} \frac{1}{2} (u_A(L, R, (A, -1)) - u_A(R, R, (A, -1))) + \frac{1}{4} (u_A(L, R, (B, -1)) - u_A(R, R, (B, -1))) \\ + \frac{1}{4} (u_A(L, L, (B, 1)) - u_A(R, L, (B, 1))) = 0.35 > 0. \end{aligned}$$

For the case $y = -1$ and $a_A = R$,

$$\begin{aligned} \frac{1}{2} (u_A(R, L, (A, -1)) - u_A(L, L, (A, -1))) + \frac{1}{4} (u_A(R, L, (B, -1)) - u_A(L, L, (B, -1))) \\ + \frac{1}{4} (u_A(R, R, (B, 1)) - u_A(L, R, (B, 1))) = 0.15 > 0. \end{aligned}$$

Similarly for player $i = B$:

For the case where $y = 1$ and $a_B = L$

$$\begin{aligned} \frac{1}{2} (u_B(L, L, (B, 1)) - u_B(L, R, (B, 1))) + \frac{1}{4} (u_B(L, L, (A, 1)) - u_B(L, R, (A, 1))) \\ + \frac{1}{4} (u_B(R, L, (A, -1)) - u_B(R, R, (A, -1))) = 0.35 > 0. \end{aligned}$$

For the case where $y = 1$ and $a_B = R$

$$\begin{aligned} \frac{1}{2} (u_B(R, R, (B, 1)) - u_B(R, L, (B, 1))) + \frac{1}{4} (u_B(R, R, (A, 1)) - u_B(R, L, (A, 1))) \\ + \frac{1}{4} (u_B(L, R, (A, -1)) - u_B(L, L, (A, -1))) = 0.15 > 0. \end{aligned}$$

For the case $y = -1$ and $a_B = L$,

$$\begin{aligned} \frac{1}{2} (u_B(R, L, (B, -1)) - u_B(R, R, (B, -1))) + \frac{1}{4} (u_B(R, L, (A, -1)) - u_B(R, R, (A, -1))) \\ + \frac{1}{4} (u_B(L, L, (A, 1)) - u_B(L, R, (A, 1))) = 0.35 > 0. \end{aligned}$$

For the case $y = -1$ and $a_B = R$,

$$\begin{aligned} \frac{1}{2} (u_B(L, R, (B, -1)) - u_B(L, L, (B, -1))) + \frac{1}{4} (u_B(L, R, (A, -1)) - u_B(L, L, (A, -1))) \\ + \frac{1}{4} (u_B(R, R, (B, 1)) - u_A(L, R, (B, 1))) = 0.15 > 0, \end{aligned}$$

and thus all the obedience conditions for a BIBCE are satisfied.

Bibliography

- ABREU, D. AND M. K. BRUNNERMEIER (2003): “Bubbles and crashes,” *Econometrica*, 71, 173–204.
- ANGELETOS, G.-M. AND I. WERNING (2006): “Crises and prices: Information aggregation, multiplicity, and volatility,” *American Economic Review*, 96, 1720–1736.
- ARIELI, I., Y. BABICHENKO, F. SANDOMIRSKY, AND O. TAMUZ (2022): “Feasible Joint Posterior Beliefs,” *Journal of Political Economy*.
- AUMANN, R. J. (1974): “Subjectivity and correlation in randomized strategies,” *Journal of Mathematical Economics*, 1, 67–96.
- (1987): “Correlated equilibrium as an expression of Bayesian rationality,” *Econometrica: Journal of the Econometric Society*, 1–18.
- BALDER, E. (1988): “Generalized Equilibrium Results for Games with Incomplete Information,” *Mathematics of Operations Research*, 1988, 265–276.
- BERGEMANN, D. AND S. MORRIS (2016a): “Bayes correlated equilibrium and the comparison of information structures in games,” *Theoretical Economics*, 11, 487–522.
- (2016b): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- (2017): “Belief-Free Rationalizability and Informational Robustness,” *Games and Economic Behavior*, 104, 744–759.
- BERNHEIM, B. D. (1984): “Rationalizable strategic behavior,” *Econometrica: Journal of the Econometric Society*, 1007–1028.
- BRANDENBURGER, A. AND A. FRIEDENBERG (2008): “Intrinsic Correlation in Games,” *Journal of Economic Theory*, 141, 28–67.
- CARLSSON, H. AND E. VAN DAMME (1993a): “Global games and equilibrium selection,” *Econometrica: Journal of the Econometric Society*, 989–1018.
- (1993b): “Global games and equilibrium selection,” *Econometrica: Journal of the Econometric Society*, 989–1018.

- CHEN, Y.-C., A. DI TILLIO, E. FAINGOLD, AND S. XIONG (2017): “Characterizing the Strategic Impact of Misspecified Beliefs,” *Review of Economic Studies*, 84, 1424–1471.
- CICHOMSKI, S. AND A. OSEKOWSKI (2024): “Coherent distributions on the square—extreme points and asymptotics,” *Journal of Applied Probability*, 61, 1240–1262.
- (2025): “On the existence of extreme coherent distributions with no atoms,” *Journal of Theoretical Probability*, 38, 1–15.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2005): “Topologies on types,” *Harvard Institute of Economic Research Discussion Paper*.
- (2006): “Topologies on Types,” *Theoretical Economics*, 1, 275–309.
- (2007a): “Interim correlated rationalizability,” *Theoretical Economics*.
- (2007b): “Interim Correlated Rationalizability,” *Theoretical Economics*.
- DU, S. (2012): “Correlated Equilibrium and Higher-Order Beliefs about Play,” *Games and Economic Behavior*, 76, 74–87.
- ELY, J. AND M. PESKI (2011): “Critical Types,” *Review of Economic Studies*, 78, 907–937.
- ENGL, G. (1995): “Lower Hemicontinuity of the Nash Equilibrium Correspondence,” *Games and Economic Behavior*, 9, 151–160.
- FOLLAND, G. B. (1999): *Real analysis: modern techniques and their applications*, John Wiley & Sons.
- FORGES, F. (1986): “An approach to communication equilibria,” *Econometrica: Journal of the Econometric Society*, 1375–1385.
- (1993a): “Five legitimate definitions of correlated equilibrium in games with incomplete information,” *Theory and decision*, 35, 277–310.
- (1993b): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 61, 329–344.
- (2006): “Correlated Equilibrium in Games with Incomplete Information Re-Visited,” *Theory and Decision*, 61, 329–324.
- FRANKEL, D. M., S. MORRIS, AND A. PAUZNER (2003): “Equilibrium selection in global games with strategic complementarities,” *Journal of Economic Theory*, 108, 1–44.
- GENSBITTEL, F., M. PESKI, AND J. RENAULT (2022): “Value-Based Distance between Information Structures,” *Theoretical Economics*.
- GOLDSTEIN, I. AND A. PAUZNER (2005): “Demand–deposit contracts and the probability of bank runs,” *the Journal of Finance*, 60, 1293–1327.

- GOSSNER, O. (2000a): “Comparison of information structures,” *Games and Economic Behavior*, 30, 44–63.
- (2000b): “Comparison of Information Structures,” *Games and Economic Behavior*, 30, 44–63.
- GOSSNER, O. AND J.-F. MERTENS (2020): “The Value of Information in Zero-Sum Games,” Tech. rep., LSE.
- GOSSNER, O. AND R. VEIEL (2024): “Strategic Type Spaces,” Tech. rep., CREST, Institut Polytechnique de Paris.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1976): “Information and competitive price systems,” *The American Economic Review*, 66, 246–253.
- HALAC, M., E. LIPNOWSKI, AND D. RAPPOPORT (2021): “Rank uncertainty in organizations,” *American Economic Review*, 111, 757–86.
- HARSANYI, J. (1967-68): “Games with Incomplete Information Played by ‘Bayesian’ Players,” *Management Science*, 14, 159–189, 320–334, 485–502.
- HARSANYI, J. C. (1967): “Games with incomplete information played by “Bayesian” players, Part I. The basic model,” *Management science*, 14, 159–182.
- HARSANYI, J. C. AND R. SELTEN (1988): “A general theory of equilibrium selection in games,” *MIT Press Books*, 1.
- HELLMAN, Z. (2014): “A Game with no Bayesian Approximate Equilibria,” *Journal of Economic Theory*, 153, 138–151.
- INOSTROZA, N. AND A. PAVAN (2023): “Adversarial Coordination and Public Information Design,” Tech. rep., University of Toronto and Northwestern University.
- KAJII, A. AND S. MORRIS (1997): “The robustness of equilibria to incomplete information,” *Econometrica: Journal of the Econometric Society*, 1283–1309.
- (1998): “Payoff Continuity in Incomplete Information Games,” *Journal of Economic Theory*, 82, 267–276.
- KAMBHAMPATI, A. (2023): “Payoff Continuity in Games of Incomplete Information Across Models of Knowledge,” Tech. rep., United States Naval Academy.
- KAMENICA, E. (2019): “Bayesian persuasion and information design,” *Annual Review of Economics*, 11, 249–272.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- KYLE, A. S. (1985): “Continuous auctions and insider trading,” *Econometrica: Journal of the Econometric Society*, 1315–1335.

- LAMBERT, N. S., M. OSTROVSKY, AND M. PANOV (2014): “Strategic trading in informationally complex environments,” in *Proceedings of the fifteenth ACM conference on Economics and computation*, 3–4.
- LEHRER, E., D. ROSENBERG, AND E. SHMAYA (2010): “Signaling and Mediation in Games with Common Interests,” *Games and Economic Behavior*, 68, 670–682.
- LI, F., Y. SONG, AND M. ZHAO (2023): “Global Manipulation by Local Obfuscation,” *Journal of Economic Theory*.
- LIPMAN, B. L. (2003): “Finite order implications of common priors,” *Econometrica*, 71, 1255–1267.
- LIU, Q. (2009): “On Redundant Types and the Bayesian Formulation of Incomplete Information,” *Journal of Economic Theory*, 144, 2115–2145.
- (2015a): “Correlation and common priors in games with incomplete information,” *Journal of Economic Theory*, 157, 49–75.
- (2015b): “Correlation and Common Priors in Games with Incomplete Information,” *Journal of Economic Theory*, 157, 49–75.
- MATHEVET, L., J. PEREGO, AND I. TANEVA (2020a): “On information design in games,” *Journal of Political Economy*, 128, 1370–1404.
- (2020b): “On Information Design in Games,” *Journal of Political Economy*, 128, 1370–1404.
- MERTENS, J.-F., S. SORIN, AND S. ZAMIR (2015): *Repeated games*, vol. 55, Cambridge University Press.
- MERTENS, J.-F. AND S. ZAMIR (1985): “Formulation of Bayesian analysis for games with incomplete information,” *International journal of game theory*, 14, 1–29.
- MILGROM, P. AND R. WEBER (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 619–632.
- MILGROM, P. R. (1981): “Rational expectations, information acquisition, and competitive bidding,” *Econometrica: Journal of the Econometric Society*, 921–943.
- MONDERER, D. AND D. SAMET (1989): “Approximating common knowledge with common beliefs,” *Games and Economic Behavior*, 1, 170–190.
- (1996): “Proximity of Information in Games with Incomplete Information,” *Mathematics of Operations Research*, 21, 707–725.
- MONDERER, D. AND L. S. SHAPLEY (1996): “Potential games,” *Games and economic behavior*, 14, 124–143.

- MORRIS, S. (2014): “Coordination, timing and common knowledge,” *Research in Economics*, 68, 306–314.
- MORRIS, S., D. OYAMA, AND S. TAKAHASHI (2020): “Implementation via information design in binary-action supermodular games,” *Available at SSRN 3697335*.
- (2024): “Implementation via Information Design in Binary-Action Supermodular Games,” *Econometrica*, 93, 693–726.
- MORRIS, S. AND H. S. SHIN (1998): “Unique equilibrium in a model of self-fulfilling currency attacks,” *American Economic Review*, 587–597.
- (2003a): *Global Games: Theory and Applications*, Cambridge University Press, 56–114, Econometric Society Monographs.
- (2003b): “Global games: Theory and applications,” in *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1*, Cambridge University Press, 56–114.
- (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.
- OYAMA, D. AND S. TAKAHASHI (2011): “On the relationship between robustness to incomplete information and noise-independent selection in global games,” *Journal of Mathematical Economics*, 47, 683–688.
- (2020): “Generalized Belief Operator and Robustness in Binary-Action Supermodular Games,” *Econometrica*, 88, 693–726.
- PEARCE, D. G. (1984): “Rationalizable strategic behavior and the problem of perfection,” *Econometrica: Journal of the Econometric Society*, 1029–1050.
- PESKI, M. (2008): “Comparison of Information Structures in Zero-Sum Games,” *Games and Economic Behavior*, 62, 732–735.
- RENY, P. J. AND M. PERRY (2006): “Toward a strategic foundation for rational expectations equilibrium,” *Econometrica*, 74, 1231–1269.
- RUBINSTEIN, A. (1989a): “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge,’” *The American Economic Review*, 385–391.
- (1989b): “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge,’” *American Economic Review*.
- (1989c): “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge,’” *American Economic Review*, 79, 385–391.

- SADZIK, T. (2019): “Beliefs Revealed in Bayesian-Nash Equilibrium,” Tech. rep., UCLA.
- SIGA, L. AND M. MIHM (2021): “Information aggregation in competitive markets,” *Theoretical Economics*, 16, 161–196.
- SIMON, R. S. (2003): “Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria,” *Israel Journal of Mathematics*, 138, 73–92.
- SIMON, R. S. AND G. TOMKOWICZ (2017): “A Bayesian Game without e-Equilibria,” Tech. rep., LSE.
- STINCHCOMBE, M. B. (2011): “Correlated Equilibrium Existence for Infinite Games with Type-Dependent Strategies,” *Journal of Economic Theory*, 146, 638–655.
- TCHEBICHEF, M. (1852): “Mémoire sur les nombres premiers.” *Journal de Mathématiques pures et appliquées*, 1, 366–390.
- UI, T. (2001): “Robust equilibria of potential games,” *Econometrica*, 69, 1373–1380.
- VAN ZANDT, T. (2010): “Interim Bayesian Nash Equilibrium on Universal Type Spaces for Supermodular Games,” *Journal of Economic Theory*, 145, 249–263.
- VIVES, X. (2014): “On the possibility of informationally efficient markets,” *Journal of the European Economic Association*, 12, 1200–1239.
- WEINSTEIN, J. AND M. YILDIZ (2007): “A Structure Theorem for Rationalizability with Applications to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.
- WILLARD, S. (2012): *General topology*, Courier Corporation.