Information Design for Rationalizability*

Olivier Gossner† Rafael Veiel‡

First version: October 20, 2022

This version: November 14, 2024

Abstract

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.

*We are grateful to Stephen Morris, Daisuke Oyama, Satoru Takahashi, Jonathan Weinstein, Muhamet Yildiz for valuable comments and helpful discussions. We also thank seminar participants at SAET, Grenoble, Marseille, NYU-AD, Harvard/MIT, UCLA USC, CalTech, Columbia, UBC-SFU, Northwestern, Dauphine, LSE, IHP and the One World Mathematical Game Theory Seminar. This work was financially supported by the French National Research Agency (ANR), Investissements d'Avenir (ANR-11-IDEX0003/LabEx Ecodec/ANR-11-LABX-0047) (Gossner).

†CNRS - CREST - École Polytechnique and London School of Economics

‡Massachusetts Institute of Technology

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. (2007). 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., 2020; 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. (2007). 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, 2016), 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 5.2, 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 (1993) and Morris and Shin (2003). 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 proof strategy consists of three main steps, where each step introduces new conceptual tools. In the first one, we study the possible laws followed by 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, as well as a sufficient set of information structures to generate those, 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 limit distributions 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 also build a dictionary that allows to reinterpret SCAMP through generalizations of correlated equilibria, the email game, and global games.

We now detail these three main steps: a revelation principle for ICR hierarchies, SCAMP, and the dictionary for information structures.

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 (the state of nature and) players' ICR hierarchies. Conversely, any probability distribution on (nature and) profiles of ICR hierarchies 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, in this new information structure, 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 are closed forms that characterize the distributions on ICR hierarchies arising from any potential common priors.

Now that we have pinned 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 $K \times \Omega^{\mathbb{N}}$ induces, through the mappings from Ω to action sets, a distribution on (K and) sequences of action sets for all players.

We say that a probability distribution on $K \times \Omega^{\mathbb{N}}$ 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 (1993) 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 (1989).

The rest of the paper is organized as follows. In Section 2 we illustrate our concepts and results in a game of technology adoption. We present the model in Section 3 and introduce Strategic Automata as our main tool to represent ICR hierarchies in Section 4. Our main results are presented in Section 5. Section 7 discusses the complexity and interpretation of SCAMP and relates SCAMP to existing work. Finally, in Section 6 we apply our methods and results to the design of information in common value priority systems.

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

	a	b
a	$2,1$	$0,0$
b	$0,0$	$1,2$

G

	a	b
a	$0,1$	$0,0$
b	$2,2$	$1,0$

B

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, as 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 $\mathrm{R}_i^n = \mathrm{R}_i^n(t_i)$ the set of actions that survive n rounds of deletion of dominated strategies given i's beliefs. We just have established:

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

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:

$$\mathrm{R}_1^{n+1} = \begin{cases} a & \text{if } 3p_1(\mathrm{R}_2^n{=}a, G) - p_1(G) > 2p_1(B) - p_1(\mathrm{R}_2^n{=}b, B) \\ b & \text{if } p_1(\mathrm{R}_2^n{=}a, B) + p_1(B) > 2p_1(G) - 3p_1(\mathrm{R}_2^n{=}b, G) \\ ab & \text{otherwise} \end{cases}$$

and for player 2:

$$\mathrm{R}_2^{n+1} = \begin{cases} a & \text{if } p_2(\mathrm{R}_2^n{=}b, B) + p_2(B) > 2p_2(G) - 3p_2(\mathrm{R}_2^n{=}a, G) \\ b & \text{if } 3p_2(\mathrm{R}_2^n{=}b, G) - p_2(G) > 2p_2(B) - p_2(\mathrm{R}_2^n{=}a, B) \\ ab & \text{otherwise} \end{cases}$$

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 $\mathrm{R}_1^n = ab$, for n odd we may have $\mathrm{R}_1^{n+1} = ab$ or $\mathrm{R}_1^{n+1} = a$ but not $\mathrm{R}_1^{n+1} = b$ and for n even we may have $\mathrm{R}_1^{n+1} = ab$ or $\mathrm{R}_1^{n+1} = b$ but not $\mathrm{R}_1^{n+1} = a$. A symmetric property holds for player 2.

2.1 Strategic Automaton

The possible ICR hierarchies for each player are summarized on the automaton of Figure 1. The state labeled with"start for Pi", is the initial (or 0-th) level of for player i, $\mathrm{R}_i^0 = ab$. 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 $\mathrm{R}_i^0 = ab$, $\mathrm{R}_i^1, \ldots, \mathrm{R}_i^n, \ldots$ that appear with positive probability in some common prior model.

Figure 2 allows to visualize the possible joint ICR hierarchies for both players as the set of infinite sequences starting at the initial state and that follow arrows, possibly reaching a terminal state.

Let us call Ω the set of 16 states of figure 2. Each state $\omega \in \Omega$ is labeled with an action set ω_i for each player i. Since every pair of types (t_1, t_2) in a

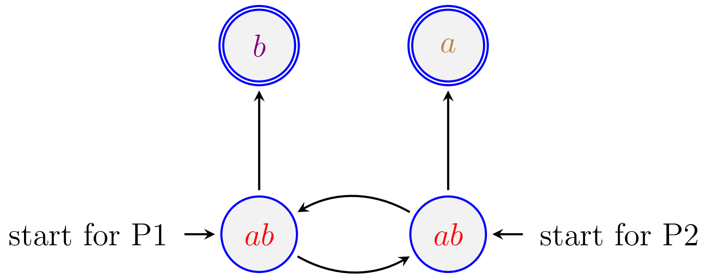

Figure 1: Automaton for one player in the technology example. There are 4 states and each state contains an action set. The initial state is on the left (player 1) or on the right (player 2). Double circled states are terminal ones.

type space can be mapped to a path in the automaton, it follows that every prior P 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 P varies?

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

Our revelation principle (Theorem 4.1) shows that a distribution on $K \times \Omega^{\mathbb{N}}$ arises from $(\mathrm{R}_i^n)_n$ applied to some information structure P if and only if, in the information structure where i is informed of $(\omega_i^1, \ldots, \omega_i^n, \ldots)$, i's ICR hierarchy is precisely $(\omega_i^1, \ldots, \omega_i^n, \ldots)$. From the derivation of ICR hierarchies from types above, this is the case when for every n, and when p_i denotes $P(k, \omega_{-i}^1, \ldots, \omega_{-i}^n, \ldots | \omega_i^1, \ldots, \omega_i^n, \ldots)$:

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

and for player 2:

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

These equations, which we call *Obedience Constraints*, are expressed directly on conditional beliefs p_i, and thus depend only on the probability

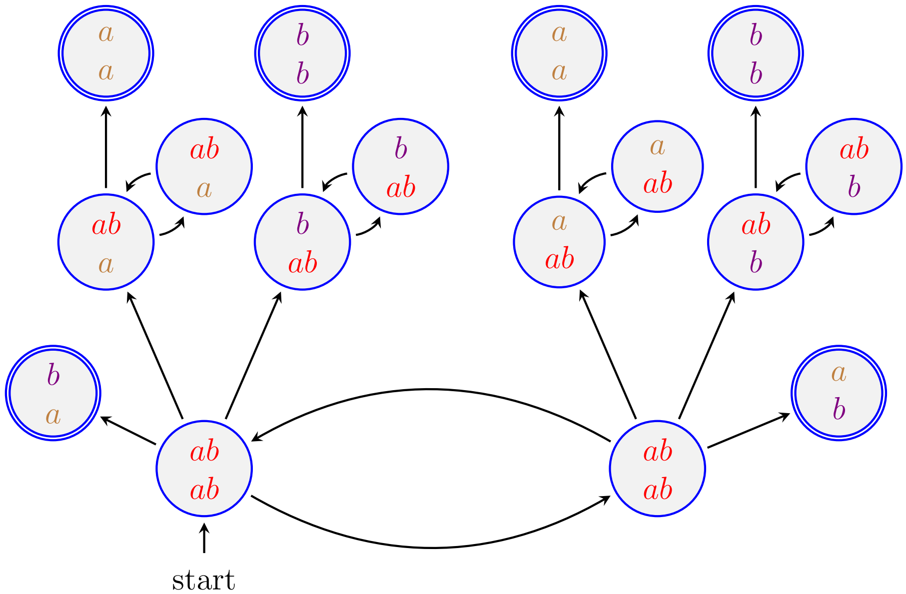

Figure 2: Automaton for both players in the technology example. Each state contains an action set for player 1 (top) and for player 2 (bottom). Arrows indicate possible transitions.

distribution P on $K \times \Omega^{\mathbb{N}}$.

The revelation principle thus fully characterizes the possible distributions of $(k, (\mathrm{R}_1^n)_n, (\mathrm{R}_2^n)_n)$ that may arise in any common prior model. It also characterizes information structures that yield these distributions, as *canonical* information structures 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.

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. Remember that for a type t_i of player i in a type space, the set of rationalizable distributions is obtained as $\mathrm{R}_i^\infty = \mathrm{R}_i^\infty(t_i) = \cap_n \mathrm{R}_i^n(t_i)$. We say that a distribution μ on $K \times (\{ab, a, b\})^2$ is *rationalizable* if there exits a common prior P such that the induced distribution of $(k, \mathrm{R}_1^\infty, \mathrm{R}_2^\infty)$ is μ.

Our SCAMP revelation principle, Theorem 5.1 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 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). By applying the same transformation whenever possible, we obtain a Markov chain of the form in Figure 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 during a cycle, the probability of reaching terminal nodes 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.

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

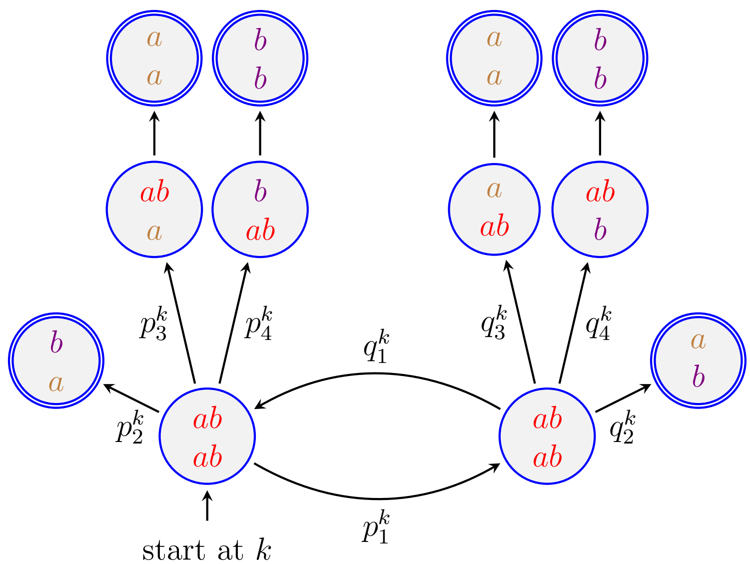

Figure 3: SCAMP generating point distributions. When a single arrow leaves a state, this arrow has probability 1. Transitions may depend on the state of nature $k \in \{G, B\}$.

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 4.

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

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 cycles).

In our technological coordination game, each player receives a recommendation to play either a or b corresponding to the terminal state of the

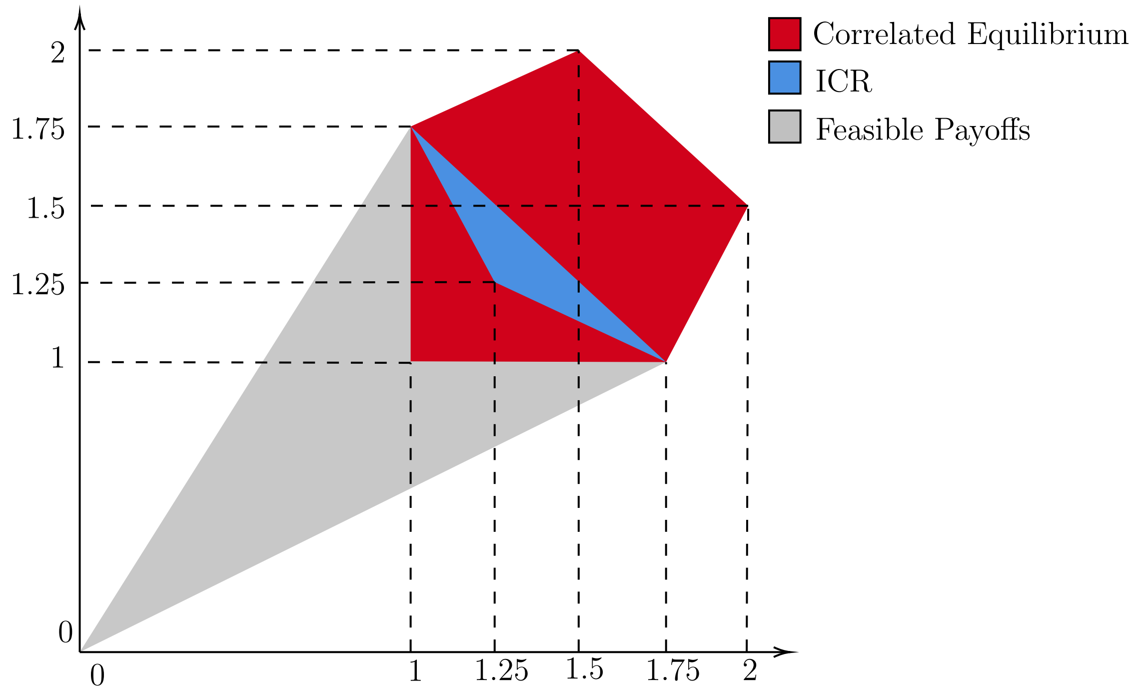

Figure 4: Point rationalizable payoffs (in blue), correlated equilibrium payoffs (in red) and feasible payoffs (in grey) in the coordination game.

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 $(ab, \ldots, ab, a, a, \ldots)$ 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 rounds. Players' private signals are thus slightly out of sync. This is very similar to Rubinstein (1989)'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, 1993), 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 Model

General notations For any mapping $f\colon X \to Y$ and any subset $E \subseteq X$ we write $f(E) \coloneqq \{f(x) : x \in E\}$. 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.

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

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

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 $u_i\colon K \times A \to \mathbb{R}$, 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 over $K \times T$ admitting a conditional probability $P(\cdot|\cdot)\colon T_i \to \Delta(K \times T_{-i})$ for every player i such that for every $k \in K$, 1) $t_i \mapsto P(k, X_{-i}|t_i)$ is measurable for every measurable set $X_{-i} \subseteq T_{-i}$ and 2) $P(k, X_i \times X_{-i}) = \int_{X_i} P(k, X_{-i}|t_i) dP(t_i)$, for every $k \in K$ and measurable sets $X_i \subseteq T_i$, $X_{-i} \subseteq T_{-i}$.

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 Interim Correlated Rationalizability (Dekel et al., 2007) is the outcome of the process of elimination of dominated strategies, or equivalently never best-responses, in the agent normal form of the game with incomplete information. It is defined as follows.[1] Let B_i denote the collection of non-empty subsets of A_i, and define a *conjecture for player i* as a map $\sigma_i : K \times B_{-i} \to \Delta(A_{-i})$ such that the support

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

of $\sigma_i(k, a_{-i})$ is included in b_{-i} for every $k \in K$ and $b_{-i} \in B_{-i}$. A belief $p \in \Delta(K \times B_{-i})$ and a conjecture σ_i induce a probability distribution $\langle \sigma, p \rangle \in \Delta(K \times A_{-i})$ given by:

$$\langle \sigma_i, p \rangle(k, a_{-i}) = \sum_{b_{-i} \in B_{-i}} p(k, b_{-i}) \ \sigma_i(k, b_{-i})(a_{-i}). \tag{3.1}$$

Player i's best-reply map $\mathrm{br}_i \colon \Delta(K \times B_{-i}) \to B_i$ is defined by:

$$\mathrm{br}_i(p) = \bigcup_{\sigma} \left\{ \arg\max_{a_i \in A_i} \mathbb{E}_{\langle \sigma, p \rangle} u_i(\cdot, a_i, \cdot) \right\}, \tag{3.2}$$

where $\mathbb{E}_{\langle \sigma, p \rangle}$ is the expectation with respect to the distribution $\langle \sigma, p \rangle$. The ICR *hierarchy* $(\mathrm{R}_i^m(t_i))_{m \geq 0}$ of a type $t_i \in T_i$ is defined iteratively:

i) For every $i \in N$ and $t_i \in T_i$, $\mathrm{R}_i^0(t_i) = A_i$,

ii) For $m \geq 0$, $P(\cdot|t_i) \in \Delta(K \times T_{-i})$ and the (measurable) map R_{-i}^m from T_{-i} to B_{-i} induce a belief $P(\cdot|t_i) \circ (\mathrm{id} \times \mathrm{R}_{-i}^m)^{-1}$ on $K \times B_{-i}$, and $\mathrm{R}_i^{m+1}(t_i)$ is the set of best-responses to this belief:

$$\mathrm{R}_i^{m+1}(t_i) = \mathrm{br}_i(P(\cdot|t_i) \circ (\mathrm{id} \times \mathrm{R}_{-i}^m)^{-1}). \tag{3.3}$$

The set of rationalizable actions associated to t_i is

$$\mathrm{R}_i^{\infty}(t_i) = \bigcap_{m \in \mathbb{N}} \mathrm{R}_i^m(t_i). \tag{3.4}$$

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

$$\mu_P = P \circ (\mathrm{id} \times \mathrm{R}^{\infty})^{-1}. \tag{3.5}$$

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

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 in common priors model.

4.1 Revelation principle

We define the set of ICR hierarchies as the set of all possible hierarchies that arise as we vary the common prior.

Definition 4.1 (ICR Hierarchies). *The set of* ICR *hierarchies is the minimal subset $S \subseteq B^{\mathbb{N}}$ so that for every common prior model P and every t in the support of P, $(\mathrm{R}^m(t))_m \in S$.*

We first characterize the distributions on K and ICR hierarchies arising from common priors as well as information structures that implement those hierarchies through a revelation principle on S.

Every common prior P induces, through ICR and the identity on K, a distribution P_{R} on $K \times B^{\mathbb{N}}$. Let $\mathcal{P}$ denote the set of such distributions when we vary over all common priors on all possible type spaces. The following result characterizes $\mathcal{P}$. Note that every distribution $P \in \Delta(K \times B^{\mathbb{N}})$ can be viewed itself a common prior in which the set of types for player i is $B_i^{\mathbb{N}}$, $(k, (b_i^n)_{i,n})$ is drawn according to P, and each player i is informed of her corresponding type $(b_i^n)_n \in B_i^{\mathbb{N}}$.

Theorem 4.1 (Revelation Principle). *For $P \in \Delta(K \times B^{\mathbb{N}})$ the three conditions are equivalent:*

1. $P \in \mathcal{P}$;
2. (a) $\mathrm{R}(s) = s$, P *a.s.* ;

   (b) $P = P_{\mathrm{R}}$, *i.e., P is the image of itself viewed as a common prior ;*
3. $P(s^0{=}A) = 1$ *and P satisfies the family of* obedience constraints*:*

$$s_i^m = \mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^{m-1}} P(\cdot,\cdot|s_i)) \text{ for a.e. } s_i = (s_i^m)_m. \tag{4.1}$$

The proof of Theorem 4.1 can be found in Appendix A.1. The Revelation Principle implies that every distribution on K and ICR hierarchies can be implemented through itself viewed as a common prior. In turn, these distributions are entirely characterized by the family of obedience constraints (OC's), one for each type and for each level m.

4.2 Strategic Automaton

In this section, we show that the set S of possible ICR hierarchies coincides with the set of paths of a finite automaton. This finite structure captures the recursivity of the best-response operator of the game.

An *automaton* is a triple $(\Omega, \beta, \preceq)$ given by a finite set of states Ω together with an action map $\beta_i \colon \Omega \to B_i$ for every player i and a binary successor relation $\preceq$ on states. A *cycle* is an ordered collection of states $c = \{\omega^1, \ldots, \omega^n\} \subseteq \Omega$ so that $\omega^h \preceq \omega^{h+1}$ for all $h < n$ and $\omega^1 = \omega^n$. A *path* is a sequence $(\omega^0, \omega^1, \ldots)$ 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 corresponds to the set of ICR hierarchies.

Definition 4.2 (Strategic Automaton)**.** *A strategic automaton is an automaton $(\Omega, \beta, \preceq)$ such that:*

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

To prove existence of a strategic automaton, we construct one directly. We provide a construction of an automaton whose paths give rise to S and then prove that it is finite and thus a strategic automaton.

The initial state is given by $\omega^0 = S$ and $\beta_i(\omega^0) = A_i$ for every i. For every $m \in \mathbb{N}$, let $S^* = \{(s^0, \ldots, s^m) : s \in S, m \in \mathbb{N}\}$ be the set of truncated sequences of S. The set of tails associated to a truncated sequence $(s^0, \ldots, s^m) \in S^*$ is:

$$\tau(s^0, \ldots, s^m) = \{(s^m, s^{m+1}, \ldots) \in B^{\mathbb{N}} : (s^0, \ldots, s^m, s^{m+1}, \ldots) \in S\}, \quad (4.2)$$

and the set of states of the automaton is the collection of tail sets:

$$\Omega = \{\tau(s^*) : s^* \in S^*\}.$$

For every $\omega \in \Omega$, action labels for player i are given by $\beta_i(\omega) = s_i^m$ for $(s^m, s^{m+1}, \ldots) \in \omega$, where it follows from (4.2) that the definition is independent of the choice of $(s^m, s^{m+1}, \ldots) \in \omega$.

We finally define the successor relation for $\omega, \omega' \in \Omega$ by $\omega \preceq \omega'$ if and only if there exists $s^0 \in B$ so that

$$\{(s^0, s^1, s^2, \ldots) : (s^1, s^2, \ldots) \in \omega'\} \subseteq \omega. \quad (4.3)$$

We rely on the following result from Gossner and Veiel (2024).

Theorem 4.2. *$(\Omega, \beta, \preceq)$ is a strategic automaton.*

5 Rationalizable Distributions

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

5.1 SCAMP

A path on a strategic automaton $(\Omega, \beta, \preceq)$, 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 $m \in \mathbb{N}$,

$$P(\omega^{m+1}|k, \omega^0 \dots \omega^m) = P(\omega^{m+1}|k, \omega^m), \quad P\, a.s. \tag{5.1}$$

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

$$\mathrm{R}(s) = s, \quad P\, a.s. \tag{5.2}$$

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

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 5.1 (Sufficiency of SCAMP)**.** *For every finite game, there exists a strategic automaton on which SCAMP induces all outcome distributions.*

To prove Theorem 5.1, we start with a canonical prior represented as a process on a strategic automaton, and we modify it into a SCAMP on a new automaton. The new automaton depends on the strategic automaton we start with, but is independent of the canonical prior. The construction involves averaging probabilities of transitions in order to make the new process Markov, and modifying the automaton in such a way that the new one has at most one cycle. We illustrate the construction in a simple example below, and then present the general construction.

Example Figure 5 below describes the transition probabilities of a canonical prior P (top) on a hypothetical strategic automaton. In this example, we assume that for every action of every player, there is a state where this action is dominant. Figure 6 describes the transition probabilities of a Markov chain on the same automaton. For every $m \in \mathbb{N}$, $p^1_{k,m}$ is the probability making the left-most transition conditional on cycling for m rounds and conditional on state $k \in K$,

$$p^1_{k,m} = P(s^{m+1} = (ab, a) : k, s^n = (ab, ab), \forall\ n \leq m). \tag{5.3}$$

The transition probabilities $p^2_{k,m}, \ldots, p^4_{k,m}$ are defined similarly. For every $m \in \mathbb{N}$ and state $k \in K$, the probability of cycling m times conditional on k be given by $P(m|k) = \prod_{n \leq m}(1 - \sum_{l=1}^{4} p^l_{k,m})$.

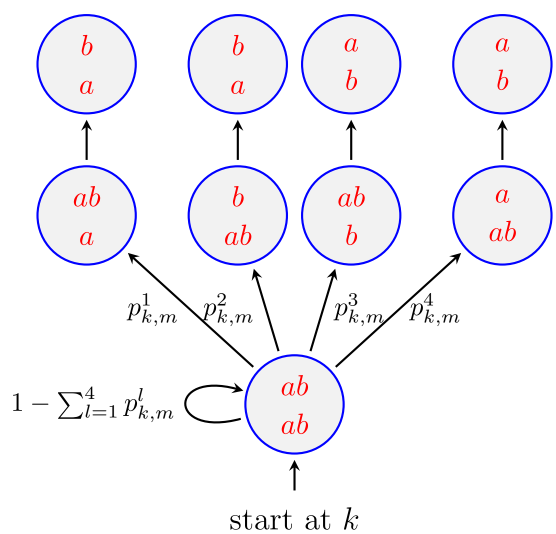

Figure 5: Canonical Prior as Process on SCAMP automaton.

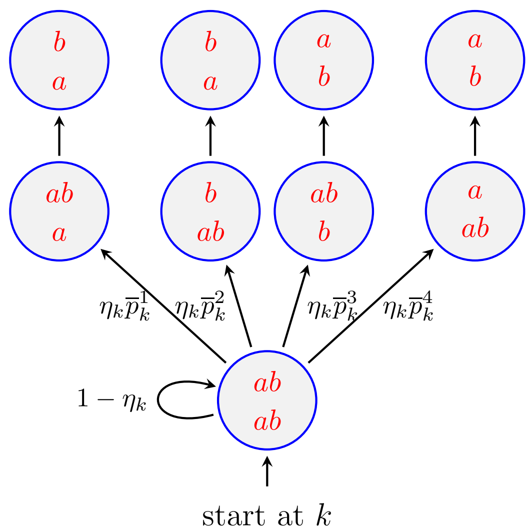

Figure 6: Averaging process to SCAMP.

We now construct a SCAMP process, represented in Figure 6, which induces the same outcome distribution as the original process P. The Markov transitions are parametrized by a family η_k, $\bar{p}_k^1, \bar{p}_k^2, \bar{p}_k^3, \bar{p}_k^4$ and depend on k. We need $\overline{p}_k = (\overline{p}_k^1, \ldots, \overline{p}_k^4)_k$ and $(\eta_k)_k$ to satisfy the following three conditions[2]:

i) Outcome Equivalence:

$$\begin{aligned} \overline{p}_k^1 + \overline{p}_k^2 &= (p_{k,1}^1 + p_{k,1}^2) + \sum_{m=2}^{\infty} (p_{k,m}^1 + p_{k,m}^2) P(m-1|k), \\ \overline{p}_k^3 + \overline{p}_k^4 &= (p_{k,1}^3 + p_{k,1}^4) + \sum_{m=2}^{\infty} (p_{k,m}^3 + p_{k,m}^4) P(m-1|k). \end{aligned} \tag{5.4}$$

ii) Obedience of type where player 1 transitions to a at the first round:

$$\sum_{k \in K} P(k) \eta_k \overline{p}_k^1 (u_1(k, a, \alpha_2) - u_1(k, b, \alpha_2)) > 0, \ \forall \ \alpha_2 \in \{a, b\}. \tag{5.5}$$

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

iii) Obedience of types transitioning to a at round $m+1$ for player 1:

$$\begin{aligned}\sum_{k\in K} P(k)\big(&\eta_k\ (1-\eta_k)^{m-1}\ \overline{p}_k^1\ (u_1(k,a,b)-u_1(k,b,b)) \\ &+\eta_k\ (1-\eta_k)^{m}\ \overline{p}_k^2(u_1(k,a,\alpha_2)-u_1(k,b,\alpha_2))\big)>0, \\ &\forall\ \alpha_2\in\{a,b\},\ \forall\ m\geq 1.\end{aligned} \tag{5.6}$$

Since P satisfies obedience, we know that

$$\sum_{k\in K} P(k)p_k^1(u_1(k,a,\alpha_2)-u_1(k,b,\alpha_2))>0,\ \forall\ \alpha_2\in\{a,b\}. \tag{5.7}$$

Consider setting $\eta_k=1$ whenever the utility difference of condition (ii) is positive for all actions α_{-i} - when a (respectively b for the other player) is a dominant action - and setting $\eta_k=\zeta\in(0,1)$ when the utility difference is negative for some action α_{-i}. By convexity of the OCs, condition (iii) is satisfied when choosing $\overline{p}_k^l=p_k^l+\sum_{m=2}^{\infty}p_{k,m}^l P(m-1|k)$ at each $l=1,\dots,4$ and any such choice of $(\eta_k)_k$. Note that this choice of $\overline{p}$ also satisfies condition (i), which is independent of $(\eta_k)_k$. Given our choice of $\overline{p}$, there must be ζ small enough so that

$$\zeta\overline{p}_k^1\leq p_k^1,\ \forall\ k \text{ s.t. } \min_{\alpha_2\in\{a,b\}}(u_1(k,a,\alpha_2)-u_1(k,b,\alpha_2))<0. \tag{5.8}$$

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.

5.2.1 General Construction of Markov Process

An important challenge in the general construction is that not all priors and automata are as well behaved as the one in the example above. The paths of a prior P may loop through multiple cycles, they may pass through several states before and after each cycle. In this section we provide an overview of the general construction that takes care of these situations. Much of the work will involve bringing priors into a form that resembles the automaton in the example of the previous subsection.

Construction of Extended Automaton We start out with the strategic automaton $(\Omega,\beta,\preceq)$ defined in Section 4.2, which we then extend in order to prove Theorem 5.1. In preparation for our construction, we will restrict attention to a subclass of well-behaved priors in $\mathcal{P}$, which we call product priors that are closed from below. To define this property, we consider the following partial order on sequences $s, \hat{s} \in S$,

$$s \subseteq^*_S \hat{s} \iff s_i^m \subseteq \hat{s}_i^m, \ \forall\, i, m \text{ and } B(s) = B(\hat{s}), \tag{5.9}$$

where $B(s) \coloneqq \{s^m : m \in \mathbb{N}\}$. Define the set of lower bounds of $s, s' \in S$,

$$\underline{S}(s, s') \coloneqq \{\overline{s} \in S : \overline{s} \subseteq^*_S s, \overline{s} \subseteq^*_S s'\}. \tag{5.10}$$

Say that P is *closed from below* if for all $s, s' \in S$ in the support of P so that $\underline{S}(s, s') \neq \varnothing$, we also have that

$$P(\underline{S}(s, s')) > 0. \tag{5.11}$$

For any player i, we write $\hat{s} \subseteq^*_{S_i} s$ when for every $m \in \mathbb{N}$, $\hat{s}_i^m \subseteq s_i^m$, and $B(s) = B(\hat{s})$. A prior $P \in \mathcal{P}$ is a *product prior* if for every i and for every pair of sequences $s, \hat{s} \in S_P$ satisfying $\hat{s} \subseteq^*_{S_i} s$,

$$P(s_i, \hat{s}_{-i}) > 0, \text{ and } P(\hat{s}_i, s_{-i}) > 0. \tag{5.12}$$

We first show that every prior admits an outcome equivalent product prior that is closed from below.

Lemma 5.1. *For every $P \in \mathcal{P}$ there exists an outcome equivalent product prior $P^* \in \mathcal{P}$ that is closed from below.*

We prove Lemma 5.1 by exploiting a key monotonicity property of br, which we state in Claim A.1 in the appendix.

We show that under this class of priors we can construct the Markov process from the $\subseteq^*_S$-minimal sequences: For any $P \in \mathcal{P}$, let $\underline{S}_P$ denote the set of $\subseteq^*_S$-minimal sequences in the support of P,

$$\underline{S}_P \coloneqq \min_{\subseteq^*_S}\{s \in S : P(S) > 0\}. \tag{5.13}$$

For any $s \in S$ in the support of P and player i let $\underline{S}_{i,P}(s)$ denote the $\subseteq^*_S$-minimal sequences in the support of s_i's beliefs:

$$\underline{S}_{i,P}(s) \coloneqq \min_{\subseteq^*_S}\{\tilde{s} \in S : P(\tilde{s}|s_i) > 0\}. \tag{5.14}$$

We say that P has a *seed* if for every $\underline{s} \in \underline{S}_P$, $\underline{S}_{i,P}(\underline{s}) \subseteq \underline{S}_P$.

Lemma 5.2. *Every product prior that is closed from below has a seed.*

We now proceed to the construction of the strategic automaton $(\overline{\Omega}, \overline{\beta}, \overline{\preceq})$ for which Theorem 5.1 holds. In $(\overline{\Omega}, \overline{\beta}, \overline{\preceq})$, each state consists of a pair $\overline{\omega} = (\omega, (\iota_i^M)_i)$, where $\omega \in \Omega$ and $\iota_i^M \subseteq \Omega^M$ describes player i's information set chains of a given length $M \in \mathbb{N}$ that we specify in Lemma 5.3 below. The mapping $\overline{\beta}$ and successor relation $\overline{\preceq}$ are inherited from β and $\preceq$ in a natural way.

Let $P \in \mathcal{P}$ be a product prior that is closed from below. Let $s \in \underline{S}_P$ and fix $m \in \mathbb{N}$. A (s,m)-*chain* is a tuple of sequences

$$(s(0), \ldots, s(m)) \in S^m \tag{5.15}$$

so that $s(0) = s$ and for all $l \leq m$ there is a player i_l so that $s(l) \in \underline{S}_{i_l,P}(s(l-1))$. Let $\mathcal{C}_P(s,m)$ denote the collection of (s,m)-chains. Lemma 5.3 below states that for product priors that are closed from below, all chains must have a repeating entry after some length M. This is an immediate consequence of the finiteness of the set of $\subseteq_S^*$-minimal sequences for product priors that are closed from below.

Lemma 5.3. *There is $M \in \mathbb{N}$ so that for every product prior $P \in \mathcal{P}$ that is closed from below, every $s \in \underline{S}_P$ and every chain $(s(0), \ldots, s(m)) \in \mathcal{C}_P(s,m)$, there are $\underline{m}_s, \overline{m}_s \leq M$ so that*

$$s(\underline{m}_s) = s(\overline{m}_s). \tag{5.16}$$

We now construct the strategic automaton for Theorem 5.1. As mentioned earlier, we do so by extending the states in Ω. Fix $s \in \underline{S}_P$. For every i define the set of (s,M)-chains that start with player i,

$$\mathcal{C}_{i,P}(s) \coloneqq \{(s(0), \ldots, s(M)) \in \mathcal{C}_P(s,M) : s(0)_i = s(1)_i\}, \tag{5.17}$$

where $M \in \mathbb{N}$ verifies the statement in Lemma 5.3. For $m \in \mathbb{N}$, define the m-th extended information set of s,

$$\mathcal{I}_{i,P}^m(s) \coloneqq \{(\tau^m(s(0)), \ldots, \tau^m(s(M))) : (s(0), \ldots, s(M)) \in \mathcal{C}_{i,P}(s)\}, \tag{5.18}$$

where we use the short-hand $\tau^m(\tilde{s}) \coloneqq \tau(\tilde{s}^0, \ldots, \tilde{s}^m)$, for all $\tilde{s} \in S$. The m-th extended information set of player i thus keeps track of the state in the m-th coordinate of every sequence along a chain starting with player i. We will bootstrap the paths of our Markov chain from $\subseteq^*_S$-minimal sequences in $\underline{S}_P$. For the resulting Markov chain to satisfy OC, we need to keep track of all chains and thus of all extended information sets of sequences in $\underline{S}_P$. We let $\mathcal{I}^m_P(s) \coloneqq (\mathcal{I}^m_{i,P}(s))_i$ and define the extended automaton state

$$\bar{\omega}^m_P(s) = (\tau^m(s), \mathcal{I}^m(s)). \tag{5.19}$$

The collection of extended automaton states across all priors $P \in \mathcal{P}$ is finite and is denoted $\overline{\Omega}$. We let the successor relation $\overline{\preceq}$ be defined by $(\omega, \iota)\overline{\preceq}(\hat{\omega}, \hat{\iota})$ when 1) $\omega \preceq \hat{\omega}$, 2) for all i, all $(J^0, \dots, J^M) \in \iota_i$ there is $(\hat{J}^0, \dots, \hat{J}^M) \in \hat{\iota}_i$ so that for every $l \leq M$, $J^l \preceq \hat{J}^l$, and 3) for all i, all $(\hat{J}^0, \dots, \hat{J}^M) \in \hat{\iota}_i$ there is $(J^0, \dots, J^M) \in \iota_i$ so that for every $l \leq M$, $J^l \preceq \hat{J}^l$. Finally, we define the label on extended automaton states $\overline{\beta}(\omega, \iota) \coloneqq \beta(\omega)$.

Construction of Paths We now proceed to specifying the paths in the support of our Markov process. For every $s \in \underline{S}_P$ the *m-order branch* $\bar{t}^m_{s,P}$ is the ordered tuple of distinct elements in $\{\bar{\omega}^n_P(s) : n \leq m\}$ according to $\overline{\preceq}$. For every $m \in \mathbb{N}$ let $\mathcal{T}^m_P$ denote the collection of m-order branches. For every $t \in \mathcal{T}^m_P$ and any player i let

$$t_i = (\overline{\beta}_i(t^0), \dots, \overline{\beta}_i(t^m)). \tag{5.20}$$

In our construction of the Markov process, we will average over the probabilities of paths that pass through the same branch. For this procedure to preserve OCs, branches need to be of a certain length, m^*, which is bounded by the parameters of the game. Lemma 5.4 below states that there is a length m^* so that for every $\subseteq^*_S$-minimal sequence $s \in \underline{S}_P$, every player i and any round $n \leq m^*$, the branches of the minimal sequences in the support of s_i's beliefs, $\underline{S}_{i,P}(s_i)$, have the same action labels as the branch of s at round n.

Lemma 5.4. *There is $m^* \in \mathbb{N}$ so that for all product priors $P \in \mathcal{P}$ that are closed from below, for every $s \in \underline{S}_P$, every player i and every $\hat{s} \in \underline{S}_{i,P}(s_i)$,*

$$(\bar{t}^{m^*}_{s,P})_i = (\bar{t}^{m^*}_{\hat{s},P})_i. \tag{5.21}$$

We conclude from Lemma 5.4 that the support of the obedience constraints for $\subseteq^*_S$-minimal types at every transition is preserved when represented as branches. We now bootstrap all paths in the support of our Markov process by adding at most one cycle to each branch. Let m^* satisfy the statement in Lemma 5.4. For every branch $t \in \mathcal{T}^{m^*}_P$ let $C(t)$ and $\bar{C}(t)$ denote the first and terminal cycle in $t = (t^0, \ldots, t^{m^*})$, respectively. For every $s \in \underline{S}_P$ and any choice of $l \in \mathbb{N}$, define the path

$$v(s,l) := ((\bar{t}^{m^*}_{s,P})^0, \ldots, \underbrace{C(\bar{t}^{m^*}_{s,P}), \ldots, C(\bar{t}^{m^*}_{s,P})}_{l-1 \text{ times}}, \ldots, \bar{C}(\bar{t}^{m^*}_{s,P}), \ldots). \tag{5.22}$$

Define the set of paths of the Markov process

$$\overline{\mathbf{\Omega}} := \{v(s,l) : s \in \underline{S}_P, l \in \mathbb{N}\}. \tag{5.23}$$

For every path $v = (v^m)_{m\in\mathbb{N}} \in \overline{\mathbf{\Omega}}$, define $v_i := (\beta_i(v^m))_{m\in\mathbb{N}}$. For every player i define

$$\overline{\mathbf{\Omega}}_i(v_i) := \{\hat{v} \in \overline{\mathbf{\Omega}} : \hat{v}_i = v_i\}. \tag{5.24}$$

Construction of Markov Process In order to define the Markov process, we must now assign probabilities to elements in K and the paths in $\overline{\mathbf{\Omega}}$. For every $s \in \underline{S}_P$, let its upper-contour set be denoted by

$$\overline{S}_P(s) := \{\tilde{s} \in S_P : s \subseteq^*_S \tilde{s}\}. \tag{5.25}$$

So we set

$$\overline{P}_\zeta(k, v(\underline{s}, l)) := \zeta^l(1-\zeta)\ P(\{k\} \times \overline{S}_P(\underline{s})), \tag{5.26}$$

for any cycling probability $\zeta \in [0,1]$, every $\underline{s} \in \underline{S}_P$ and $l \in \mathbb{N}$. For every $\underline{s} \in \underline{S}_P$ and type v_i define the minimal number of loops required to bring a $\subseteq^*_S$-minimal sequence into the information set of an arbitrary Markov-type v_i,

$$\ell(\underline{s}, v_i) := \min\{l \in \mathbb{N} : v(\underline{s}, l) \in \overline{\mathbf{\Omega}}_i(v_i)\}. \tag{5.27}$$

Lemma 5.5 below states that the $\subseteq^*_S$-minimal sequences and the minimal sequences in the support of their beliefs are always aligned. The result is an immediate consequence of Lemma 5.4.

Lemma 5.5. *For every $\underline{s} \in \underline{S}_P$ and any $\tilde{s} \in \underline{S}_{i,P}(\underline{s})$,*

$$\ell(\underline{s}, v_i) = \ell(\tilde{s}, v_i). \tag{5.28}$$

Finally, we show that there exists ζ small enough so that the Markov process $\overline{P}_\zeta$ is outcome equivalent to P, which concludes the argument.

5.3 Characterization of Rationalizable Outcomes

We now show that the set of outcomes is characterized by a convex polyhedron. This result follows from the sufficiency of SCAMP established in Theorem 5.1 and exploits the "simple" property in SCAMP. Let m^* be the length of branches used in the construction of the Markov process above. A distribution $p \in \Delta(K \times \overline{\Omega}^{m^*})$ satisfies OCs on $\overline{\Omega}^{m^*}$ if for every $0 < m \leq m^*$ and all $s_i = (\overline{\omega}_i^0, \ldots, \overline{\omega}_i^{m^*})$,

$$s_i^m = \mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^{m-1}}(p(\cdot, \cdot | s_i))). \tag{5.29}$$

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

$$\bar{p}(k, b) = \sum_{(\overline{\omega}^0, \ldots, \overline{\omega}^{m^*}) \in X^* : \overline{\beta}(\overline{\omega}^{m^*}) = b} p((k, (\overline{\omega}^0, \ldots, \overline{\omega}^{m^*})) | X^*). \tag{5.30}$$

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

$$b_i = \mathrm{br}_i(\bar{p}(\cdot, \cdot | b_i)). \tag{5.31}$$

Let $\mathcal{O}^\infty \subseteq \Delta(K \times B)$ denote the set of probabilities $\bar{p}$ satisfying $b_i = \mathrm{br}_i(\bar{p}(\cdot, \cdot | b_i))$ 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.

$$\mathcal{O} = \{\bar{p} : p \in \mathcal{O}^{m^*}\} \cap \mathcal{O}^\infty. \tag{5.32}$$

The relative closure of this set is a convex polyhedron:

Lemma 5.6 (Linearity of $\mathcal{O}$)**.** *The relative closure of the set $\mathcal{O}$ is a convex polyhedron.*

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

$$\mathcal{O}^* = \{\nu_P : P \in \mathcal{P}\}. \tag{5.33}$$

Theorem 5.2 (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 5.1 can be used to map every distribution in $\mathcal{O}^{m^*}$ into a SCAMP. Since every branch has at most one cycle (simple), its outcome distribution coincides with $\bar{p}$. Conversely, every SCAMP induces a distribution in $\mathcal{O}^{m^*}$ through its marginal probability on the first m^* rounds.

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 (1989). They can also be represented as information structures that are widely used in economic applications: Asynchronous information[3] provision (as in Abreu and Brunnermeier, 2003), and a common state with additive noise (as in Carlsson and Van Damme, 1993). 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 $\tau_i \in \{0, \ldots, n_i\}$, where $n_i \leq |\overline{\Omega}|$. 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

$$y_i = \theta + \tau_i. \tag{5.34}$$

[3] See Morris (2014) for a more complete overview.

Theorem 5.3 (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 5.3 is an immedaite consequence of Theorem 5.1: 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 $\overline{\beta}_i$.

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 (1993). 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 5.3 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 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.

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.

6.1 Model

Two buyers want to use a service of unknown quality[4] $k \in K = \{1, \dots, \bar{k}\}$, 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:

$$u_i(k, a_i, a_j) = \begin{cases} k - a_i, \text{ if } a_j < a_i, \\ \alpha_k k - a_i, \text{ if } 0 < a_i = a_j, \\ 0, \text{ otherwise,} \end{cases} \tag{6.1}$$

where $\alpha_k \in [0, 1]$ represents an externality of consuming the service simultaneously when the quality is k and the utility of consuming last is set to zero.

[4] We let K be an interval of $\mathbb{N}$ to simplify the analysis.

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

$$V(P) = \sum_{k,b} \min_{a_i \in b_i, a_j \in b_j} (a_i + a_j)\ \nu_P(k, b). \tag{6.2}$$

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 $\overline{a}_k$ the least priority with this property: $\overline{a}_k = \max\{a \in A : a \leq \alpha_k k\}$.

6.2 Complete Information Benchmark

As a benchmark, we compute the outcomes when the quality is publicly observed. Let $\mathrm{R}_i(k) \subseteq A$ 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 6.1. $\mathrm{R}_i(k) = \begin{cases} 0, \ldots, \overline{a}_k & \text{if } \alpha_k > \frac{k-1}{k} \\ \overline{a}_k & \text{otherwise.} \end{cases}$

Proof. Assume first $\alpha_k > \frac{k-1}{k}$ and let $a_j \leq \overline{a}_k$. 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 $a_i \leq \overline{a}_k$ is a best-reply to itself in state k, hence is rationalizable. Moreover, every $a > \overline{a}_k$ is dominated in state k by $\overline{a}_k$.

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 $a_1, \ldots \overline{a}_{k-1}$ by successive outbidding. is a best-reply as long as $\alpha_k k \geq a_j$. $\square$

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 $\overline{a}_k$ 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 $(\overline{a}_k)_k$. 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. $\overline{a}_k - \alpha_k k$, shrinks to zero.

Lemma 6.2. *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

$$\begin{aligned} V(P) &\leq \sum_{s \in S} \left(\mathbb{E}_P(\alpha_k k|s_i) + \mathbb{E}_P(\alpha_k k|s_j)\right) P(s) \\ &\leq 2\mathbb{E}_P(\alpha_k k) = \sum_{k \in K} 2\overline{a}_k P_K(k). \end{aligned} \tag{6.3}$$

□

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 $k < \overline{k}$, then the upper bound is attained by an ANIS information structure where every player i is privately informed of the following signal:

$$z_i = (a_k, \theta_k + \varepsilon_i), \tag{6.4}$$

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 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.

6.4 SCAMP/ANIS

In Figure 7 we show an automaton that attains the upper bound provided above: all terminal nodes have singleton action labels given by $\overline{a}_k$, for every

k and every player. Indeed, each automaton state in Figure 7 contains an action set for each player of the form $\{a, \ldots, \overline{a}_k\}$, where $a \leq \overline{a}_k$. These action sets are identified in the figure by their minimal action. For example the node containing "$\begin{matrix} a_k \\ a_k+1 \end{matrix}$" represents an automaton state labeled with action sets $\{a_k, \ldots, \overline{a}_k\}$ for player 1 and $\{a_k + 1, \ldots, \overline{a}_k\}$ 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.

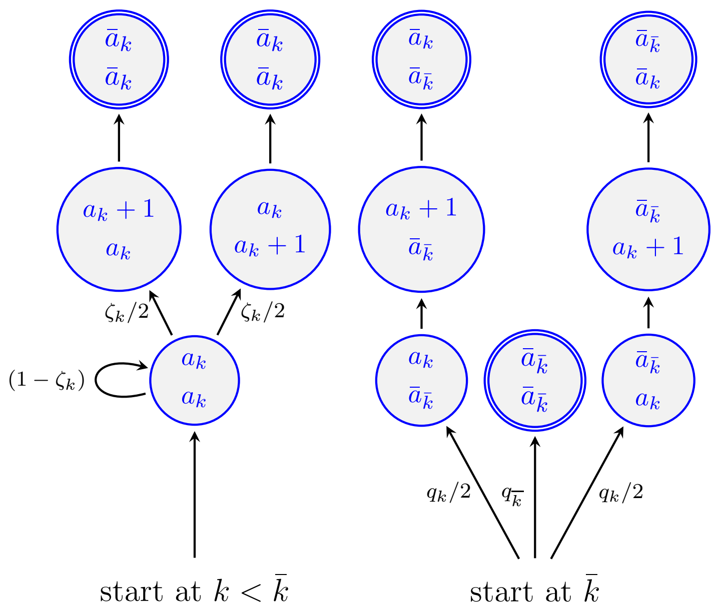

Figure 7: SCAMP Automaton

The SCAMP automaton in Figure 7 has two key parameters for every quality $k \in K$: ζ_k and q_k. The parameter q_k parametrizes the transition probabilities on the right part of the automaton when the state is $\bar{k}$. Positive q_k means that any player observing a sequence of the form $(a_k, a_k + 1, \ldots, \bar{a}_k, \bar{a}_k, \ldots)$ is uncertain if the quality is k or $\bar{k}$. By attaching just the right probability to $\bar{k}$, we can make OC hold for that sequence. The parameter ζ_k appears on the left part of the automaton and parametrizes the probability of exiting the cycle. After observing a sequence

$(a_k, a_k, a_k + 1, \ldots, \bar{a}_k, \bar{a}_k, \ldots)$, a player knows that the state is k, but may not know if their opponent knows this.

Proposition 6.1 below states that, under certain conditions on the parameters $(\alpha_k)_k$, the automaton in Figure 7 can be populated with transition probabilities $(\zeta_k, q_k)_k$ so that OC holds for all sequences. This establishes our (approximate) surplus extraction result:

Proposition 6.1 (Optimal Information Disclosure: ANIS)**.** *Suppose $\alpha_{\bar{k}} \leq \frac{\bar{k}-1}{\bar{k}}$ is close enough to one so that $\zeta_k \in (0,1)$ for all $k < \overline{k}$ and suppose $\alpha_k > \frac{k-1}{k}$ for all $k < \overline{k}$. Then for every P_K the ANIS described in (6.11) attains the upper bound payoff guarantee in Lemma 6.2.*

The relevant conditions for a SCAMP/ANIS to extract (approximately) all of the surplus are twofold and can be explained using Lemma 6.1: first, the externality of simultaneous consumption at the highest quality $\overline{k}$, i.e. $1-\alpha_{\overline{k}}$, 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 $\overline{k}$, the externality, i.e. $1-\alpha_k$ is low. Under complete information, Lemma 6.1 implies that all bids below $\overline{a}_k$ 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 (1989): 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 (1989)) 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

$$s_{k,m,i} = (0, \ldots, \underbrace{a_k, \ldots, a_k}_{m \text{ times}}, a_k + 1, \ldots, \overline{a}_k, \overline{a}_k, \ldots). \tag{6.5}$$

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

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

$$P(k', s_{k,1,i}) = \begin{cases} \frac{\zeta_{k'}}{2} P(k'), \text{ if } k' = k \\ \frac{q_k}{2} P(k'), \text{ if } k' = \overline{k} \\ 0, \text{ otherwise.} \end{cases} \tag{6.6}$$

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

$$P(k', s_{k,1,i}) = \begin{cases} P(\overline{k}), \text{ if } k' = \overline{k} \\ 0, \text{ otherwise.} \end{cases} \tag{6.7}$$

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):

$$\begin{aligned} Q_k \alpha_{\overline{k}} \overline{k} + (1 - Q_k)\alpha_k k - \alpha_k k = Q_k \overline{k} + (1 - Q_k)k - \alpha_k k - 1 \\ Q_k(1 - \alpha_{\overline{k}})\overline{k} + (1 - Q_k)(1 - \alpha_k)k = 1 \\ Q_k = \frac{1 - (1 - \alpha_k)k}{(1 - \alpha_{\overline{k}})\overline{k} - (1 - \alpha_k)k} \end{aligned} \tag{6.8}$$

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

$$Q_k = \frac{q_k/2P(\overline{k})}{\zeta_k/2P(k) + q_k/2P(\overline{k})}. \tag{6.9}$$

That is, $q_k = 2\frac{1-(1-\alpha_k)k}{P(\overline{k})}$ and $\zeta_k = 2\frac{(1-\alpha_{\overline{k}})\overline{k}-1}{P(k)}$. 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 $m > 1$:

$$P(k, s_{k,m,i}, s_{k,n,-i}) = \begin{cases} P(k)(1-\zeta_k)^m \frac{\zeta_k}{2}, \text{ if } n = m \\ P(k)(1-\zeta_k)^{m-1} \frac{\zeta_k}{2}, \text{ if } n = m-1 \\ 0, \text{ otherwise.} \end{cases} \tag{6.10}$$

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 (6.4), so that for each $k \in K$,

$$\begin{aligned} \Pr(\theta_k = n|k) &= \begin{cases} \zeta_k^*(1-\zeta_k^*)^n, \text{ if } k \neq \bar{k} \\ \mathbf{1}_{n=0}, \text{ otherwise.} \end{cases} \\ \Pr(a_{\tilde{k}}|k) &= \begin{cases} (1-q_k^*), \text{ if } \tilde{k} = k \\ q_k^*, \text{ if } \tilde{k} = \bar{k} \\ 0, \text{ otherwise.} \end{cases} \\ \Pr(\varepsilon_{i,k} = 0|k) &= \frac{1-\zeta_k^*}{2-\zeta_k^*} \\ q_k^* = \frac{1-(1-\alpha_k)k}{P(\bar{k})}&, \ \zeta_k^* = \frac{(1-\alpha_{\bar{k}})\bar{k} - 1}{P(k)} \end{aligned} \tag{6.11}$$

7 Discussion

7.1 Structure and Complexity of SCAMP

Consider the strategic automaton used to establish Theorem 5.1. 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,

$$|\{(k, v) \in K \times \overline{\Omega}^{\mathbb{N}} : P(k, v|s_i) > 0\}| \leq L. \tag{7.1}$$

The following result is then an immediate consequence of "simple":

Corollary 7.1 (Bounded OC of SCAMP)**.** *Every SCAMP has bounded obedience constraints.*

In the example of Section 2 we moved from an automaton with multiple cycles (Figure 2) to an automaton with only one cycle (Figure 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.

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 (2015) 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.

7.3 Relation to Literature

SCAMP can be viewed as a generalization of the e-mail game in Rubinstein (1989), 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, 1993; Morris and Shin, 2003), 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.

References

Abreu, D. and M. K. Brunnermeier (2003): "Bubbles and crashes," *Econometrica*, 71, 173–204.

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.

Bergemann, D. and S. Morris (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11, 487–522.

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

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

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

Forges, F. (1986): "An approach to communication equilibria," *Econometrica: Journal of the Econometric Society*, 1375–1385.

——— (1993): "Five legitimate definitions of correlated equilibrium in games with incomplete information," *Theory and decision*, 35, 277–310.

Gossner, O. and R. Veiel (2024): "Strategic Type Spaces," Tech. rep., CREST, Institut Polytechnique de Paris.

Halac, M., E. Lipnowski, and D. Rappoport (2021): "Rank uncertainty in organizations," *American Economic Review*, 111, 757–86.

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

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.

Liu, Q. (2015): "Correlation and common priors in games with incomplete information," *Journal of Economic Theory*, 157, 49–75.

Mathevet, L., J. Perego, and I. Taneva (2020): "On information design in games," *Journal of Political Economy*, 128, 1370–1404.

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*.

Morris, S. and H. S. Shin (2003): "Global games: Theory and applications," in *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1*, Cambridge University Press, 56–114.

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.

Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior Under" Almost Common Knowledge"," *The American Economic Review*, 385–391.

A 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 $f_i \colon X_i \to Y_i$, we let $f \colon X \to Y$ be given by $f(x) = (f_i(x_i))_i$ for $x \in X$ and $f_{-i} \colon X_{-i} \to Y_{-i}$ 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 $x_1, \ldots, x_n$ of a distribution $P \in \Delta(\prod_\ell X_\ell)$ is denoted $\operatorname{marg}_{x_1,\ldots,x_n}(p)$.

We use the following notation. For every player i let $B_i = 2^{A_i} \setminus \{\varnothing\}$ 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.

A.1 Theorem 4.1: Revelation Principle

Theorem 4.1 *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*:*

$$s_i^m = \operatorname{br}_i(\operatorname{marg}_{k, s_{-i}^{m-1}}(P(\cdot, \cdot | s_i))) \quad \textit{for a.e. } s_i = (s_i^m)_m. \tag{A.1}$$

Proof. (1. $\implies$ 3) Let $P \in \Delta(K \times T)$ be a common prior. The induced profile of conditional probabilities $(P_i : T_i \to \Delta(K \times T_{-i}))_i$ is a Harsanyi type space and so the profile of maps $((\mathrm{R}_i^m)_m \colon T_i \to S_i)_i$ satisfies: for every player i, $t_i \in T_i$ and $m \in \mathbb{N}$,

$$\mathrm{R}_i^m(t_i) = \operatorname{br}_i(P_i(t_i) \circ (\mathrm{id} \times \mathrm{R}_{-i}^{m-1})^{-1}). \tag{A.2}$$

Then note that for every $b_i \in B_i$, $\mathrm{br}_i(b_i)^{-1} \subsetneq \Delta(K \times B_{-i})$ is convex. Write $\tilde{P} \coloneqq P \circ (\mathrm{id} \times \mathrm{R})^{-1}$ with conditional probabilities $(\tilde{P}_i \colon S_i \to \Delta(K \times S_{-i})_i$ and so for all $k \in K$, $m \in \mathbb{N}$ and $s \in S$ so that $P(\mathrm{R}_i^{-1}(s_i)) > 0$,

$$\begin{aligned} s_i^m &= \mathrm{br}_i \left(\int_{\{t_i : \mathrm{R}_i^m(t_i) = s_i^m\}} P_i(t_i) \circ (id \times \mathrm{R}_{-i}^{m-1})^{-1} \mathrm{d}P(t_i | \mathrm{R}_i^m = s_i^m) \right) \\ &= \mathrm{br}_i \left(\tilde{P}_i(s_i) \right). \end{aligned} \tag{A.3}$$

So $\tilde{P}$ satisfies (A.1), as required. (3. $\implies$ 2) If $P \in \Delta(K \times B^{\mathbb{N}})$ satisfies (A.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. □

A.2 Theorem 5.1: Sufficiency of SCAMP

A.2.1 Monotonicity Property of br

A *monotone stochastic transformation* for player i is a map $\rho_i \colon K \times B_{-i} \to \Delta(B_{-i})$ so that for every $b \in B$ and $k \in K$,

$$b'_{-i} \subseteq b_{-i}, \ \forall \ b'_{-i} \in \mathrm{supp}(\rho_i(k, b_{-i})). \tag{A.4}$$

Claim A.1 (Monotonicity of br)**.** *For any monotone stochastic transformation $\rho_i \colon K \times B_{-i} \to \Delta(B_{-i})$ and for any $p_i \in \Delta(K \times B_{-i})$,*

$$\mathrm{br}_i(p_i \circ \rho_i) \subseteq \mathrm{br}_i(p_i), \tag{A.5}$$

where for all $k \in K$ and $b_{-i} \in B_{-i}$,

$$p_i \circ \rho_i(k, b_{-i}) \coloneqq \sum_{b'_{-i} \in B_{-i}} \rho_i(b_{-i} | k, b'_{-i}) p_i(k, b'_{-i}). \tag{A.6}$$

Proof. Consider any conjecture $\sigma_i \colon K \times B_{-i} \to \Delta(A_{-i})$ so that $\mathrm{supp}(\sigma(\cdot | k, b_{-i})) \subseteq b_{-i}$ 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

$$\sigma_i \circ \rho_i(a_{-i} | k, b'_{-i}) \coloneqq \sum_{b_{-i}} \sigma_i(a_{-i} | k, b_{-i}) \rho_i(b_{-i} | k, b'_{-i}). \tag{A.7}$$

Since ρ_i is monotone, the conjecture $\sigma_i \circ \rho_i$ also satisfies the support constraint of σ_i. Hence

$$\begin{aligned}
\langle \sigma_i, p_i \circ \rho_i \rangle (k, a_{-i}) &= \sum_{b'_{-i} \in B_{-i}} \left(\sum_{b_{-i} \in B_{-i}} \sigma_i(a_{-i}|k, b_{-i}) \rho_i(b_{-i}|k, b'_{-i}) \right) p_i(k, b'_{-i}) \\
&= \sum_{b'_{-i} \in B_{-i}} \sigma_i \circ \rho_i(a_{-i}|k, b'_{-i}) p_i(k, b'_{-i}) \\
&= \langle \sigma_i \circ \rho_i, p_i \rangle (k, a_{-i}).
\end{aligned} \tag{A.8}$$

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 $\sigma \colon K \times B_{-i} \to \Delta(A_{-i})$, 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,

$$\begin{aligned}
\underline{u}^m_{i,a_i,a'_i}(k, s) &\coloneqq \min_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s^m_{-i}} \sigma(a_{-i}|k, s^m)(u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})) \\
\overline{u}^m_{i,a_i,a'_i}(k, s) &\coloneqq \max_{\sigma \in \Sigma_B} \sum_{a_{-i} \in s^m_{-i}} \sigma(a_{-i}|k, s^m)(u_i(k, a_i, a_{-i}) - u_i(k, a'_i, a_{-i})).
\end{aligned}$$

For every player i and $s \in S_P$ define i's information set

$$S_{i,P}(v_i) \coloneqq \{\tilde{s} \in S_P : \tilde{s}_i = s_i\}. \tag{A.9}$$

The expected payoff increments are denoted

$$\begin{aligned}
\underline{U}^m_{i,P}(s_i, a_i, a'_i) &\coloneqq \sum_{(k,\hat{s}) \in K \times S_{i,P}(s_i)} \underline{u}^m_{i,a_i,a'_i}(k, \hat{s}) P(k, \hat{s}), \\
\overline{U}^m_{i,P}(s_i, a_i, a'_i) &\coloneqq \sum_{(k,\hat{s}) \in K \times S_{i,P}(s_i)} \overline{u}^m_{i,a_i,a'_i}(k, \hat{s}) P(k, \hat{s}).
\end{aligned} \tag{A.10}$$

OCs can be grouped into two sets of constraints:

(i) *Sub-obedience*: $\forall\ a'_i \in s^m_i \setminus s^{m+1}_i, \exists\ a_i \in s^{m+1}_i$ s.t. $0 < \underline{U}^m_i(s_i, a_i, a'_i)$,

(ii) *Super-obedience*: $\forall\ a_i \in s^{m+1}_i, \forall\ a'_i \in A_i$, $0 \leq \overline{U}^m_i(s_i, a_i, a'_i)$.

A.2.2 Closed from Below

For every $P \in \mathcal{P}$, define the sequence support

$$S_P \coloneqq \{s \in S : P(s) > 0\}. \tag{A.11}$$

Lemma A.1. *For every $P \in \mathcal{P}$ there is outcome equivalent prior $P^* \in \mathcal{P}$ that is closed from below.*

Proof. Define the partial order on sets of sequences $V, V' \subseteq S$: $V \ll V'$ if and only if for every $s \in V$ there is $s' \in V'$ so that $s \subseteq_S^* s'$ and for every $s' \in V'$ there is $s \in V$ so that $s \subseteq_S^* s'$. Let

$$\mathcal{B} \coloneqq \{V \subseteq S : \forall\ s, s' \in V,\ \underline{S}(s, s') \neq \varnothing \implies \underline{S}(s, s') \cap V \neq \varnothing\}\,. \tag{A.12}$$

For any distribution $\nu \in \Delta(K \times B)$ define the best response correspondence

$$\mathbb{B}_\nu(V) \coloneqq \bigcup_{p \in \Delta(K \times S) \text{ s.t. } \nu_p = \nu, \text{ and } V \subseteq S_p} \overline{\mathrm{br}}(p), \tag{A.13}$$

where for every $p \in \Delta(K \times S)$,

$$\overline{\mathrm{br}}(p) \coloneqq \left\{ \Big((\mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^m}(p(\cdot, \cdot | s_i))))_i \Big)_{m \in \mathbb{N}} : s \in S_p \right\}. \tag{A.14}$$

Note that for any outcome distribution $\nu \in \Delta(K \times B)$, Claim A.1 implies that the $\mathbb{B}_\nu$ operator is monotonic in $\ll$. Moreover, it can be shown that for every $V \in \mathcal{B}$, $\mathbb{B}_\nu(V) \in \mathcal{B}$. By Knaster-Tarski's fixed point theorem, $\mathbb{B}_\nu$ admits a fixed point in $\mathcal{B}$. We conclude that there is a canonical prior P^* that is closed from below with outcome distribution ν. □

Let $\mathcal{P}^*$ denote the collection of priors in $\mathcal{P}$ which are closed from below. Let $\underline{S}_P$ denote the collection of minimal sequences in S_P,

$$\underline{S}_P \coloneqq \min_{\subseteq_S^*} S_P. \tag{A.15}$$

Lemma A.2. *There exists $M \in \mathbb{N}$ so that for all $P \in \mathcal{P}^*$,*

$$|\underline{S}_P| \leq M. \tag{A.16}$$

Proof. The number of minimal elements $\underline{S} \coloneqq \min_{\subseteq_S^*} S$ is finite. Picking $M \coloneqq |\underline{S}|$ yields the result. □

A.2.3 Closed from Below Product Priors

Lemma A.3. *Every prior $P \in \mathcal{P}^*$ admits an outcome-equivalent product prior $P' \in \mathcal{P}^*$.*

Proof. Fix a pair of sequences $s, \hat{s} \in S_P$ satisfying $\hat{s}_i \subseteq^*_{S_i} s_i$. For any i and $\epsilon \in (0,1)$, define the perturbed prior

$$P_\epsilon(k, \bar{s}) = \begin{cases} (1-\epsilon)P(k,\bar{s}), \text{ if } \bar{s} = \hat{s} \\ \epsilon P(k, \hat{s}), \text{ if } \bar{s} = (s_i, \hat{s}_{-i}) \\ P(k, \bar{s}), \text{ otherwise.} \end{cases} \tag{A.17}$$

We show that there is ϵ small enough so that the perturbed prior is outcome equivalent to P. For any $m \in \mathbb{N}$ and player i, let $\mathrm{R}^m_{i,P_\epsilon}(\cdot)$ represent the m-th round of ICR for player i under prior P_ϵ. We start with $m = 1$ and $j \neq i$. Then we have that $\mathrm{R}^1_{j,P_\epsilon}(\tilde{s}_j) = \tilde{s}^1_j$ for all $\tilde{s} \in S_{P_\epsilon}$. Moreover, $\mathrm{R}^1_{i,P_\epsilon}(\tilde{s}_i) = \tilde{s}^1_i$, for all $\tilde{s}_i \neq s_i$ in the support of P_ϵ. For s_i we have that

$$P_\epsilon(k, s_i) = P(k, s_i) + \epsilon P(k, \hat{s}_i). \tag{A.18}$$

Hence $\mathrm{R}^1_{i,P_\epsilon}(\hat{s}_i) = \hat{s}^1_i \subseteq \mathrm{R}^1_{i,P_\epsilon}(s_i) \subseteq s^1_i$. Deduce that there is ϵ small enough so that sub-obedience holds for any type of any player under the perturbed prior P_ϵ. In particular, $\mathrm{R}^m_{i,P_\epsilon}(\hat{s}_i) \subseteq \mathrm{R}^m_{i,P_\epsilon}(s_i)$, for all $m \in \mathbb{N}$. To establish outcome equivalence, note that for all m, $s \in S_P$ and player i,

$$\lim_{n \uparrow \infty} s^n_i \subseteq \mathrm{R}^m_{i,P_\epsilon}(s_i). \tag{A.19}$$

Let $\hat{v} = (\mathrm{R}^m_{P_\epsilon}(\hat{s}))_m$, $v = (\mathrm{R}^m_{P_\epsilon}(s))_m$ and let $\bar{P}_\epsilon := P_\epsilon \circ (\mathrm{id} \times \prod_m \mathrm{R}^m_{P_\epsilon})^{-1}$ denote the corresponding canonical prior. Then we have that $(v_i, \hat{v}_{-i}) \in S_{\bar{P}_\epsilon}$.

Next, consider the perturbed prior

$$P'_\epsilon(k, \bar{s}) = \begin{cases} (1-\epsilon)P(k,\bar{s}), \text{ if } \bar{s} = s \\ \epsilon P(k, s), \text{ if } \bar{s} = (\hat{s}_i, s_{-i}) \\ P(k, \bar{s}), \text{ otherwise.} \end{cases} \tag{A.20}$$

Let $\hat{v}' = (\mathrm{R}^m_{P'_\epsilon}(\hat{s}))_m$, $v' = (\mathrm{R}^m_{P'_\epsilon}(s))_m$ and let $\bar{P}'_\epsilon := P'_\epsilon \circ (\mathrm{id} \times \prod_m \mathrm{R}^m_{P'_\epsilon})^{-1}$ denote the corresponding canonical prior, where $\mathrm{R}^m_{P'_\epsilon}$ denotes the mth round of delection of R under prior P'_ϵ. Analogous arguments as before show that $(\hat{v}'_i, v_{-i}) \in S_{\bar{P}'_\epsilon}$, where $\bar{P}'_\epsilon$ is outcome equivalent to P. Finally, note that if P is closed from below, then so are both $\bar{P}_\epsilon$ and $\bar{P}'_\epsilon$. One can thus easily mix over all such priors to obtain a product prior. $\square$

Lemma A.4. *Every product prior that is closed from below has a seed.*

Proof. Let $\underline{s} \in \underline{S}_P$ and suppose there is $\hat{s} \in \underline{S}_{i,P}(\underline{s}) \setminus \underline{S}_P$. Then there is $\underline{\hat{s}} \in \underline{S}_P$ so that $\underline{\hat{s}} \subseteq^*_S \hat{s}$. Consider the sequence $\bar{s} = (\hat{s}_i, \underline{\hat{s}}_{-i})$. Note that $\bar{s} \subseteq^*_S \hat{s}$. Since P is a product prior, $\bar{s} \in S_P$ and so $\underline{\hat{s}}_i \neq \hat{s}_i$. But then consider $\tilde{s} = (\underline{\hat{s}}_i, \underline{s}_{-i})$, which again must be in S_P, a contradiction. $\square$

A.2.4 Construction of Automaton and Paths

Lemma A.5. *There is $M \in \mathbb{N}$ so that for every product prior $P \in \mathcal{P}^*$, every $s \in \underline{S}_P$, and every chain $(s(0), \ldots, s(m)) \in \mathcal{C}_P(s, m)$, there are $\underline{m}_s, \overline{m}_s \leq M$ so that*

$$s(\underline{m}_s) = s(\overline{m}_s).$$

Proof. By Lemma A.4, P has a seed and by Lemma A.2 there is M so that for all product priors $P \in \mathcal{P}^*$, $|\underline{S}_P| \leq M$. For every $s \in S_P$ and any chain $(s(0), \ldots, s(m)) \in \mathcal{C}(s, m)$ we thus have that $\underline{m}_s \leq M$, where

$$\underline{m}_s \coloneqq \min\{n \leq m : \exists\ l < n, \text{ s.t. } s(n) = s(l)\}. \tag{A.21}$$

Hence the result. $\square$

Lemma A.6. *There is $m^* \in \mathbb{N}$ so that for all product priors $P \in \mathcal{P}^*$, for every $s \in \underline{S}_P$, every player i and every $\hat{s} \in \underline{S}_{i,P}(s_i)$,*

$$(\bar{t}^{m^*}_{s,P})_i = (\bar{t}^{m^*}_{\hat{s},P})_i. \tag{A.22}$$

Proof. Fix $m \in \mathbb{N}$ and suppose that $\bar{\omega}^m_P(s) = (\omega^m, \iota^m) \neq (\omega^{m+1}, \iota^{m+1}) = \bar{\omega}^{m+1}_P(s)$. Then we must have that $\iota^m_j \neq \iota^{m+1}_j$, for some player j. Suppose without loss of generality that there exists $(J^m(0), \ldots, J^m(M)) \in \iota^m_j \setminus \iota^{m+1}_j$. Then we have that

$$(\tau^m(\hat{s}), J^m(0), \ldots, J^m(M-1)) \in \hat{\iota}^m_i \setminus \hat{\iota}^{m+1}_i. \tag{A.23}$$

Deduce that for all $m \in \mathbb{N}$, $\bar{\omega}^m_P(s) \neq \bar{\omega}^{m+1}_P(s) \iff \bar{\omega}^m_P(\hat{s}) \neq \bar{\omega}^{m+1}_P(\hat{s})$ and so in particular, $(\bar{t}^{m^*}_{s,P})_i = (\bar{t}^{m^*}_{\hat{s},P})_i$ for any m^* greater than the convergence time of $(\varpi^n_P(s))_n$, i.e. greater than

$$\overline{m}(s) \coloneqq \min\{n \in \mathbb{N} : \{\varpi^h_P(s) : h \leq n\} = \{\varpi^h_P(s) : h \in \mathbb{N}\}\}. \tag{A.24}$$

We thus let $m^* \coloneqq \max_{s \in \underline{S}_P} \overline{m}(s)$. $\square$

Lemma A.7. *For every $\underline{s} \in \underline{S}_P$ and any $\tilde{s} \in \underline{S}_{i,P}(\underline{s})$,*

$$\ell(\underline{s}, v_i) = \ell(\tilde{s}, v_i). \tag{A.25}$$

Proof. This is an immediate consequence of the construction of $\overline{\mathbf{\Omega}}$ and Lemma A.6. □

For every $\tilde{v} \in \overline{\mathbf{\Omega}}$ let $\ell^*(\tilde{v})$ and $\underline{s}^*(\tilde{v})$ satisfy $\tilde{v} = v(\underline{s}^*(\tilde{v}), \ell^*(\tilde{v}))$. Define the minimal information set

$$\underline{\overline{\mathbf{\Omega}}}_i(v_i) \coloneqq \{\tilde{v} \in \overline{\mathbf{\Omega}}_i(v_i) : \ell^*(\tilde{v}) \leq \ell^*(\hat{v}),\ \forall\ \hat{v} \in \overline{\mathbf{\Omega}}_i(v_i)\}. \tag{A.26}$$

Let $\underline{\ell}(v_i) \coloneqq \min_{\tilde{v} \in \overline{\mathbf{\Omega}}_i(v_i)} \ell^*(\tilde{v})$.

A.2.5 Obedience

Lemma A.8. *There is a cycling probability ζ so that $\bar{P}_\zeta$ satisfies sub-obedience for every $m \in \mathbb{N}$, every player i and every type v_i.*

Proof. Expected minimal conjectured payoff increments between $a_i, a_i' \in A_i$ take the form:

$$\underline{U}_i^m(v_i, a_i, a_i') = \sum_{(k,v) \in K \times \hat{\mathbf{\Omega}}_i(v_i)} \underline{u}_{i,a_i,a_i'}^m(k, v) \bar{P}_\zeta(k, v) > 0, \tag{A.27}$$

where, by replacing $\hat{\mathbf{\Omega}}_i(v_i)$ with $\underline{\hat{\mathbf{\Omega}}}_i(v_i)$, $\underline{U}_i^m(v_i, a_i, a_i')$ is bounded from below by

$$\sum_{(k,\hat{v}) \in K \times \underline{\hat{\mathbf{\Omega}}}_i(v_i)} \underline{u}_{i,a_i,a_i'}^m(k, \hat{v}) \sum_{(\underline{s},l) \in \underline{S}_P \times \mathbb{N}\ :\ \hat{v} = v(\underline{s},l)} \zeta^l (1-\zeta) P(\{k\} \times \overline{S}_P(\underline{s})) \tag{A.28}$$

By Lemma A.7, we have that

$$\lim_{\zeta \downarrow 0} \frac{\underline{U}_i^m(v_i, a_i, a_i')}{\zeta^{\underline{\ell}(v_i)}} = \sum_{(k,\hat{v}) \in K \times \underline{\hat{\mathbf{\Omega}}}_i(v_i)} \underline{u}_{i,a_i,a_i'}^m(k, \hat{v}) \sum_{\underline{s} \in \underline{S}_P : \hat{v} = v(\underline{s}, \underline{\ell}(v_i))} P(\{k\} \times \overline{S}_P(\underline{s})). \tag{A.29}$$

Since $P \in \mathcal{P}^*$ has a seed, we conclude that

$$\lim_{\zeta \downarrow 0} \frac{\underline{U}_i^m(v_i, a_i, a_i')}{\zeta^{\underline{\ell}(v_i)}} > 0. \tag{A.30}$$

Then there is $\zeta \in (0, 1)$ small enough so that

$$\underline{U}_i^m(v_i, a_i, a_i') > 0, \tag{A.31}$$

which concludes the proof. □

Lemma A.9 (Outcome Equivalence)**.** *$\bar{P}_\zeta$ is outcome equivalent to P.*

Proof. For every $m \in \mathbb{N}$ and every path v in the support of $\bar{P}_\zeta$, let $\mathrm{R}^m_{\bar{P}_\zeta}(v)$ denote the m-th round of ICR under prior $\bar{P}_\zeta$. It follows from the monotonicity of br (See Claim A.1) that for all $m \in \mathbb{N}$,

$$\lim_{n \uparrow \infty} \overline{\beta}(v^n) \subseteq \mathrm{R}^m_{\bar{P}_\zeta}(v). \tag{A.32}$$

By Lemma A.8 we conclude that $\bar{P}_\zeta$ is outcome equivalent to P. ☐

Corollary A.1 (Outcome Equivalent SCAMP)**.** *Every $P \in \mathcal{P}$ admits an outcome-equivalent SCAMP.*

A.2.6 Outcome Distributions

Lemma 5.6. *The relative closure of the set $\mathcal{O}$ is a convex polyhedron.*

Proof. Clearly, the relative closure of $\mathcal{O}^{m^*}$ and that of $\mathcal{O}^\infty$ are convex polyhedra. Consider the un-scaled limit measure $\overline{p}^*$,

$$\overline{p}^*(k,b) \coloneqq \sum_{(\omega^0,\ldots,\omega^{m^*}) \in X^* : \beta(\omega^{m^*}) = b} p^*(k,\omega). \tag{A.33}$$

Then we have that the relative closure of $\{\overline{p}^* : p \in \mathcal{O}^{m^*}\}$ 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 $[\underline{x}, \overline{x}]$. Then we have that

$$\{\overline{p} : p \in \mathcal{O}^{m^*}\} = \left\{ \frac{1}{x}\overline{p}^* : p \in \mathcal{O}^{m^*}, x \in [\underline{x}, \overline{x}] \right\} \cap \Delta(X^*). \tag{A.34}$$

Indeed, each $\overline{p}^*$ has a unique $x \in [\underline{x}, \overline{x}]$ so that $\frac{1}{x}\overline{p}^* \in \Delta(X^*)$, we obtain $\mathcal{O}$ as the intersection of a cone and a simplex, making it a convex polyhedron. ☐

Theorem 5.2. *$\mathcal{O}$ 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 5.1. Then $p^*|_{X^*}$ induces a distribution on the branches of the automaton and thus a transition probability $\mu_{p^*}\colon K \times \Omega \to \Delta(\Omega)$. Using the arguments used to establish Theorem 5.1 we find a cycling probability ζ so that the Markov process induced by $\mu_{p^*}, \zeta, \operatorname{marg}_K(p)$ is SCAMP with outcome distribution p. We thus obtain our characterization of rationalizable outcomes in Theorem 5.2. $\square$