Rationalizable Outcome Distributions: A Markov Characterization*

Olivier Gossner† Rafael Veiel‡

Preliminary draft

First version: October 20, 2022

This version: December 12, 2022

Abstract

We study (interim correlated) rationalizability in a game with incomplete information. We characterize the recursive set of possible rationalizable hierarchies through a finite automaton, and provide a revelation principle that characterizes the distributions over these hierarchies that arise from any common prior. We show that a simple and finitely parametrized class of information structures, Stationary Common Automaton Markov Priors (SCAMP), is sufficient to generate every outcome distribution induced by general common prior information structures. Using this result, we characterize the set of rationalizable distributions as a convex polyhedron.

*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 with incomplete information, we are interested in the two following questions. First, what are the possible distributions of outcomes when the information received by players varies, and second, how can we, given such a distribution, construct an information structure that generates it? This question is of theoretical interest *per se* and has implications on information design and robustness.

Firstly, from the theoretical point of view, Aumann (1974, 1987) argues that to study outcomes of a game, it is important to encompass information that players may possess, while being as agnostic as possible on the nature of this information. The corresponding solution concept is correlated equilibrium and its extensions to incomplete information (Forges, 1986) and to Bayes Correlated Equilibrium (Bergemann and Morris, 2016). Each correlated equilibrium distribution is then associated with an information structure that implements it in a canonical way, where each player is informed of their action.

Secondly, 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.

Thirdly, an outcome of a game of complete information is robust if it survives with high enough probability in every neighboring game with incomplete information (Kajii and Morris, 1997). Understanding all outcomes of games with incomplete information thus provides a key to, in turn, characterizing robust outcomes.

Which outcomes are induced by an information structure in a given game depends on the solution concept considered. Correlated equilibria arise from considering the Nash equilibria of games with incomplete information. In this paper we rely on (Interim Correlated) Rationalizability (ICR) (Dekel et al., 2007), which is an extension of correlated rationalizability (Bernheim, 1984; Pearce, 1984) from complete information to incomplete information.

A correlated equilibrium is given by an information structure together with a Bayesian Nash equilibrium of the corresponding game with incomplete information. When considering information design, correlated equilibria thus provide weak implementation of the target outcome distribution:

there exists a Nash equilibrium of the game with incomplete information that yields that outcome distribution. The concept implicitly assumes that the designer has the power not only to disseminate information but also to enforce coordination on a Nash equilibrium between players.

Rationalizability does not assume coordination on an equilibrium but only common knowledge 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. Implementation in rationalizable strategies is thus both weaker than implementation in dominant strategies, and stronger than Nash implementation.

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 so obtained admits a simple description using finitely many parameters.

To obtain our characterization, we establish a revelation principle for ICR. By such a result, we mean a canonical space of signals and a family of distributions on this space such that every implementable distribution is obtained through a distribution in this family. We seek the smallest possible canonical space. It is already known (see the e-mail game from Rubinstein, 1989) that information structures with finite support are not a large enough class to generate all possible outcome distributions in finite games. It was also shown by Dekel et al. (2005) that canonical information structures on the universal type space Mertens and Zamir (1985) are precisely identified to rationalizable action sets in all possible games, while Gossner and Mertens (2001) established a similar result for zero-sum two-player games. However, universal type spaces are too large and complex to permit an operational study of rationalizability in a given game, which is an important challenge we overcome. An operational and sufficient set of signals must thus be larger than finite, and at most countably infinite.

In correlated equilibria, when applying the revelation principle, a signal for a player is identified with an action recommendation for this player. It is without loss of generality to focus on distributions over signals such that following the recommendations is incentive compatible for each player. If we generalized this principle to ICR, the set of signals would correspond to the set of rationalizable action sets for a player, and upon receiving a certain signal, the set of rationalizable actions would precisely be that signal. Such a generalization cannot work, however, as the collection of rationalizable sets in a finite game is finite, and, as noted above, finite information structures are not sufficient to generate all rationalizable distributions.

To understand our revelation principle, we need to distinguish the two notions of an ICR set versus an ICR hierarchy. A type of a player in a Harsanyi type space (Harsanyi, 1967) is given by a belief on other players' types and the payoff-relevant state of nature. To such a type corresponds a set of undominated actions in the game, which constitute the 1st level of that type's ICR hierarchy. Once a player with a certain type assumes that none of her opponents' types choose a dominated action, that player can eliminate further all dominated actions. We thus obtain the 2nd level of the type's ICR hierarchy. This elimination process defines a (non-increasing) ICR hierarchy for each possible type, and the ICR set of the type is the limit set, also given by the intersection, of the hierarchy. While the family of ICR sets in a game is finite, the family of ICR hierarchies is countable.

We show the following revelation principle result: a type of a player in a Harsanyi type space is associated with both an ICR hierarchy of that player and a belief over other players' ICR hierarchies. This ICR hierarchy and belief viewed itself as a Harsanyi type, is associated with an ICR hierarchy that coincides precisely with itself. Hence, the countable set of ICR hierarchies is sufficient to establish a revelation principle. Furthermore, given any distribution on ICR hierarchies, we show that this distribution arises from some common prior Harsanyi type space if and only if it satisfies a family of obedience constraints that are expressed directly on the distribution.

Thus, the set of distributions on ICR hierarchies is entirely characterized by these obedience constraints. However convenient this representation is, it is still countably dimensional and entails countably many obedience constraints. On the other hand, if one is interested in the outcomes of a game, characterized by ICR sets, ICR hierarchies contain redundant information, as multiple ICR hierarchies converge to the same ICR set. Consequently, one is interested in studying an appropriate subclass of distributions on ICR hierarchies which would be low dimensional and at the same time yield every possible distribution on ICR sets.

We show that an appropriate class consists of what we call Stationary Canonical Automaton Markov Priors (or SCAMP for short). Starting with any game 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. We say that a probability distribution on $K \times \Omega^n$ is a SCAMP if it 1/ satisfies the obedience constraints (it is thus a Canonical Prior), 2/ is Markovian and 3/ is such that the limit distribution of the Markov chain is the same as the distribution on terminal states obtained after a finite number of iterations (Stationarity). The first property ensures that the distribution of states of nature and hierarchies arises from a common prior, and provides a common prior that implements it, which is the distribution on the automaton paths itself. The second property implies that the class of information structures and processes considered is finitely dimensional and parametrized, and simple to generate. Finally, the third property ensures that the set of SCAMP can is characterized by the Obedience Constraints on finite histories. Therefore, SCAMP provides a finitely dimensional class of processes characterized by finitely many equations.

Our Theorem 5.1 shows that for any finite game, there exists a finite automaton such that all distributions on ICR sets are induced by SCAMP on this automaton. Thus, SCAMP provides a finitely dimensional parametrization of distributions over ICR hierarchies that induce all possible distributions on ICR sets. Because of the revelation principle above, every SCAMP can also be viewed as an information structure that induces the relevant rationalizable distribution. Hence, SCAMP provides a sufficient class of information structures for the design of information under ICR. Since they are finitely generated, they provide a simple class of information structures against which all information design and robustness under ICR can be benchmarked.

State-of-the-art results on information design under rationalizability Morris et al. (2020) and on robustness to incomplete information Oyama and Takahashi (2020) provide a good understanding of, but are limited to, binary action supermodular games. The method used in these papers relies on information structures that can be reinterpreted as Markovian under the assumption that a state of the automaton corresponds to a profile of action sets to the players.[1] For binary supermodular games, our automaton has precisely this property. For more general games, we show that the appropriate automaton possesses in general several states with the same action sets for all players. Our contribution thus sheds light on the reasons that explain the

[1]Other papers that use Markovian information structures include the original email game Rubinstein (1989) and global games Carlsson and Van Damme (1993); Morris and Shin (2003).

current literature limitations and offers a construction that overcomes these.

We use SCAMP to obtain a characterization of the set of rationalizable distributions in a game. We show that this set of probability distributions is given by a finite family of linear inequalities, from which it follows that it is a convex polyhedron. The structure of rationalizable distributions is thus simple and similar to that of correlated equilibrium distributions. This linear structure is also a virtue when considering applications.

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 show how rationalizable hierarchies are generated through a finite automaton in Section 4.1. In Section 5.1 we characterize rationalizable distributions.

2 Example: A Technology coordination game

We illustrate the concepts and results of the paper in a game of technology coordination. Two players, 1 and 2, 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, a dominates b (irrespectively of player 2's choices) iff

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

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

For player 2, b dominates a iff

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

and there are no beliefs such that a dominates b.

For $n \geq 1$, let us denote $\mathrm{R}_i^n = \mathrm{R}_i^n(t_i)$ the set of actions which survive n rounds of deletion of strategies given i's beliefs. We just have established:

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

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.

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, are the sequences $\mathrm{R}_i^0 = ab$, $\mathrm{R}_i^1, \dots, \mathrm{R}_i^n, \dots$ that appear with positive probability in some common prior model.

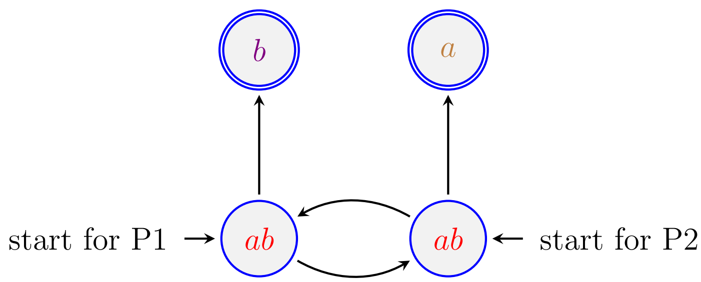

Figure 1: Automaton for one player in the technology example. There are 4 states and each state contains an action set. The initial state is on the left (player 1) or on the right (player 2). Double circled states are terminal ones.

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. Since every pair of types (t_1, t_2) in a 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 on paths when P varies?

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

Our revelation principle (Theorem 4.2) shows that P arises as a distribution on K and type hierarchies from some information structure if and only if it arises from itself viewed as an information structure. Furthermore, this is the case if and only if, for every $n \geq 0$, player i's action set ω_i^n associated to ω is precisely R_i^n. From the above, this is characterized for player 1 by the system of equations:

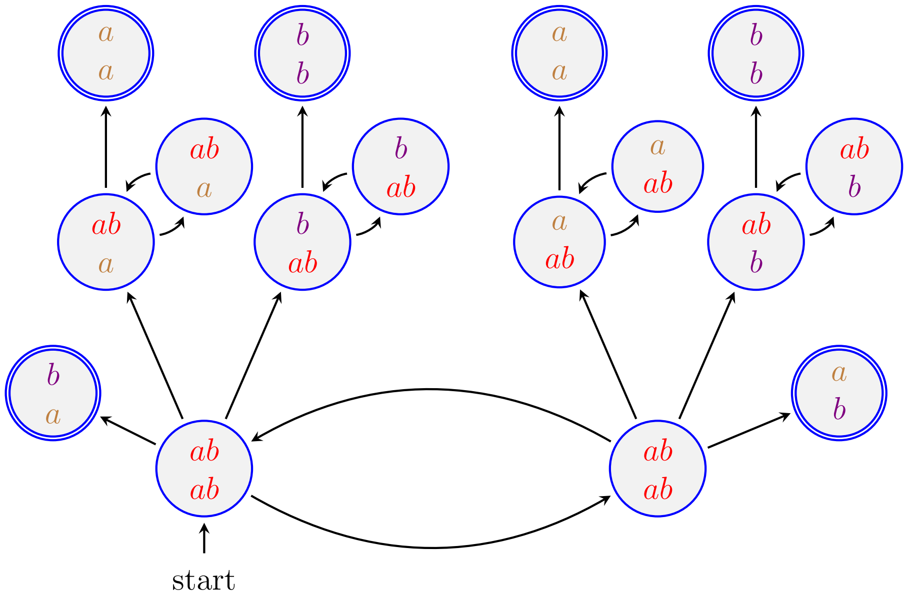

Figure 2: Automaton for both players in the technology example. Each state contains an action set for player 1 (top) and for player 2 (bottom). Arrows indicate possible transitions.

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

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 the probability 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 as the sequence of i-th coordinate of all states of the automaton, and where, the n-th state contains precisely the n-th ICR sets for both players.

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 Stationary 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 Stationary.

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, which 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 of figure 3. Furthermore, it is possible to show that this transformation doesn't violate the obedience constraints whenever 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. Conditional on exiting during a cycle, the probability of reaching terminal nodes is independent of the number of cycles. This implies the following *stationarity* property: the probability on terminal nodes of the Markov chain is given by the conditional probability on these nodes after 3 stages of the process.

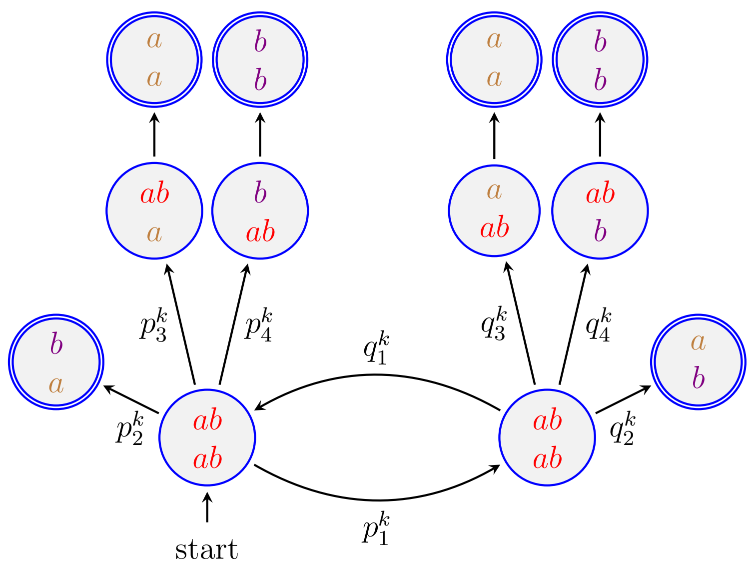

Figure 3: SCAMP generating point distributions. When a single arrow leaves a state, this arrow has probability 1. Transitions may depend on the state of nature $k \in \{G, B\}$.

Stationarity 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 P^3 on $k, \omega^1, \omega^2, \omega^3$ that satisfy OCs. This yields a characterization of the set of rationalizable distributions as a (not necessarily closed) convex polyhedron. For the technology adoption game, we illustrate the payoffs generated by point distributions in Figure 4.

The set of point rationalizable distributions, hence their payoffs, is a subset of correlated equilibria and of their payoffs. The reason is that, rationalizability is a more permissive concept than Nash equilibrium, and therefore rationalizable implementation is more stringent than correlated implementation.

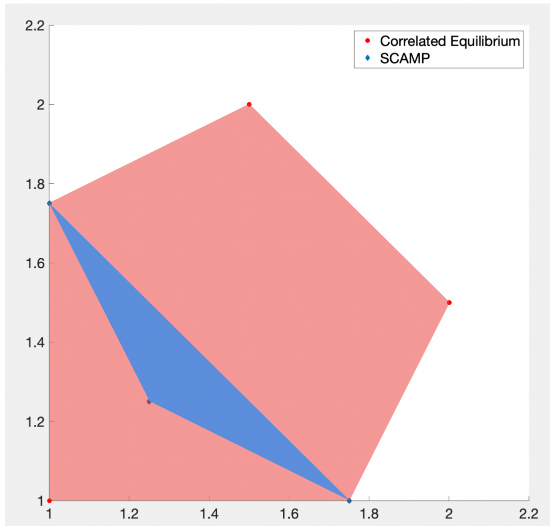

Figure 4: Payoffs generated by correlated equilibria (in red) and by point rationalizable distributions in the technology choice game.

3 Model

We fix a finite set I 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.[2] An *common prior*, denoted P, is given by a family of measurable type spaces T_i, 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 (Dekel et al., 2007) is defined as fol-

[2] 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 $(X_i)_{i \in I}$, we let $X = \prod_i X_i$ and $X_{-i} = \prod_{j \neq i} X_j$, for $i \in I$. 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 $\mathrm{marg}_{x_1,\ldots,x_n}(p)$.

lows.[3] 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 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:

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

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 $i \in I$ 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 $K \times B_{-i}$. $\mathrm{R}_i^{m+1}(t_i)$ 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

$$\mathrm{R}_i^\infty(t_i) = \bigcap_m \mathrm{R}_i^m(t_i). \tag{3.4}$$

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

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

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.

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

4.1 ICR hierarchies and Strategic Automaton

We define the set of ICR hierarchies as the set of all possible hierarchies arising from common prior models.

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, $(\mathrm{R}^m(t))_m \in S$.*

An automaton is a triple $(\Omega, \beta, \leq)$ given by a finite set of states Ω together with an action map $\beta_i : \Omega \to B_i$ for every player i and a binary successor relation $\leq$ on states. A *cycle* is a collection of states $c = \{\omega_1, \dots, \omega_m\}$ so that $\omega_1 \leq \dots \leq \omega^m \leq \omega^1$. We say that $\leq$ is a tree of cycles if the following two properties are satisfied

(i) Every $\omega \in \Omega$ is element of at most one cycle.

(ii) Every $\omega \in \Omega$ has at most one predecessor that is not in the same cycle as ω.

Definition 4.2 (Strategic Automaton)**.** *An automaton $(\Omega, \beta, \leq)$ is a strategic automaton if $\leq$ is a tree of cycles and*

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

For every $m \in \mathbb{N}$, let $S^m = \{(s^0, \dots, s^m) : s \in S\}$. We now sketch a construction of a strategic automaton.

Construction 4.1 (Construction of a Strategic Automaton)**.** *We will construct an automaton from the set S. Let $\omega^0 = S$ and $\beta_i(\omega^0) = A_i$. For any $m \in \mathbb{N}$ and a truncated sequence $(s^0, \dots, s^m) \in S^m$, define the set of m-order tails and action labels for any player i,*

$$\begin{aligned} \omega^m(s^0, \dots, s^m) &= \{(s^m, s^{m+1}, \dots) \in B^{\mathbb{N}} : (s^0, \dots, s^m, s^{m+1}, \dots) \in S\} \\ \beta_i^m(\omega^m(s^0, \dots, s^m)) &= s_i^m. \end{aligned} \tag{4.1}$$

Define the collection of m-order tails $\Omega^m = \{\omega^n(s^0, \dots, s^n) : s \in S, n \leq m\}$. Define the successor relation for any $\omega, \omega' \in \Omega^m$, $\omega \leq^m \omega'$ if and only if

$$\{(b_1, b_2, \dots) : (b_0, b_1, b_2, \dots) \in \omega'\} \subseteq \omega. \tag{4.2}$$

We obtain a triple $(\Omega, \beta, \leq)$, where $\Omega = \bigcup_{m \in \mathbb{N}} \Omega^m$, for every $m \in \mathbb{N}$, if $\omega, \omega' \in \Omega^m$ then $\beta_i(\omega) = \beta_i^m(\omega)$ and $\omega \leq \omega' \iff \omega \leq^m \omega'$. It remains to show that this triple is a strategic automaton.

The result below states that $(\Omega, \beta, \leq)$ is a strategic automaton, and in particular that the set Ω is finite.

Theorem 4.1. $(\Omega, \beta, \leq)$ *is a strategic automaton.*

We give a sketch of the argument for Theorem 4.1 and provide the proof in the Appendix. For any collection of action set profiles $B' \subseteq B$ define its *maximal* subset $\max B'$ as the collection of $b \in B'$ so that there does not exist $\hat{b} \in B' \setminus \{b\}$ satisfying

$$b_i \subseteq \hat{b}_i, \ \forall \ i. \tag{4.3}$$

For the proof of Theorem 4.1 we proceed by partitioning S into first-order maximal sequences, second-order maximal sequences, etc. First-order maximal sequences are sequences that only make maximal transitions. Second-order maximal sequences are sequences that branch out of first-order maximal sequences at some round m and make maximal transitions henceforth. We define higher-order maximal sequences analogously. We first argue that the entries of m-order maximal sequences are best replies to beliefs on $m-1$-order maximal sequences. Moreover, the number of m-order maximal sequences that branch out of a given $m-1$-order maximal sequence at a given round n is bounded, where this bound depends only on the size of the action sets.

Starting with first-order maximal sequences, we compute for every n the set of second-order maximal sequences that branch out of first-order maximal sequences at round n. Since the set of tails of second-order maximal sequences that branch out at a given n is bounded, the set will begin to cycle over n. Through an inductive argument, we then show that tails of all higher-order maximal sequences that branch out of lower-order sequences at any round n will cycle over n with some periodicity z. Since every sequence is maximal of some order, we conclude that transitions after a sequence $s^0, \ldots, s^m$ depend on this sequence modulo cycles of size z. Hence the transition correspondence, which for every m maps a truncated sequence $(s^0, \ldots, s^m) \in S^m$ to the collection of action set profiles s^{m+1} so that $(s^0, \ldots, s^m, s^{m+1}) \in S^{m+1}$, depends on a finite summary statistic of the collection of all truncated sequences. We obtain a representation of the Automaton constructed above, where the strategic state of a truncated hierarchy depends on all the action set profiles it has visited modulo cycles of length z.

4.2 Revelation principle

We now 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.

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. The following result characterizes the set of such distributions. Note that every distribution $P \in \Delta(K \times B^{\mathbb{N}})$ can be viewed itself a common prior in which $B_i^{\mathbb{N}}$ is the set of types for player i, $(k, (b_i^n)_{i,n})$ is drawn according to P, and each player i is informed of her corresponding sequence of action sets $(b_i^n)_n$.

Theorem 4.2 (Revelation principle)**.** *For $P \in \Delta(K \times B^{\mathbb{N}})$ the three conditions are equivalent:*

1. $P \in \mathcal{P}$
2. $P = P_{\mathrm{R}}$*, 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*:*

$$s_i^m = \mathrm{br}_i(P_S(k, s_{-i}^{m-1} | s_i)) \text{ for a.s. in } s_i = (s_i^m)_m \tag{4.4}$$

Hence every distribution on K and ICR hierarchies can be implemented through itself viewed as a common prior. In turn, these distributions is entirely characterized through the Obedience Constraints.

5 Rationalizable Distributions

In this section we characterize the set of distributions on $K \times B$ that arise from common prior models.

Given a Strategic Automaton $(\Omega, \beta, \preceq)$, a path is a sequence $\boldsymbol{\omega} = (\omega^0, \omega^1, \dots)$ satisfying $\omega^m \preceq \omega^{m+1}$ for all $m \in \mathbb{N}$. For every player i a path gives rise 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 $\boldsymbol{\omega} = (\omega^0, \omega^1, \dots)$ in its support is a path. A process on the automaton P is *Markov* if for every $m \in \mathbb{N}$,

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

It is *canonical* if

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

A process on the automaton defines a common prior, where each player i is privately informed of the sequence s_i. We will thus refer to the sequence s_i as a type. For a process P, we say that a type s_i of player i satisfies the Obedience Constraint at level $m \in \mathbb{N}$ when

$$s_i^m = \mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^{m-1}}(P(\cdot, \cdot | s_i))), \quad P\, a.s. \tag{5.3}$$

Note that P is canonical if and only if Obedience Constraints are satisfied for every type of every player and every round. For canonical Markov processes, k and 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}.

The entries of every path $\boldsymbol{\omega} = (\omega^0, \omega^1, \dots)$ converge to a cycle $c_\infty(\boldsymbol{\omega}) \subseteq \Omega$, given by the set of states along the path are visited infinitely often. For process P let $C_{\infty,P} := \{c_\infty(\boldsymbol{\omega}) : \mathrm{marg}_{\Omega^{\mathbb{N}}}(P)(\boldsymbol{\omega}) > 0\}$ denote the collection of cycles which are limits of paths in the support of P.

For every process P, player i, round $m \in \mathbb{N}$ and sequence s_i, define the associated *information set*, $I_{i,P}^m(s_i)$, as the set of states of nature and automaton states that are in the support of type s_i's beliefs at round m,

$$I_{i,P}^m(s_i) = \{(k, \omega) \in K \times \Omega : P(k, \omega^m = \omega | s_i) > 0\}. \tag{5.4}$$

The collection of s_i's information sets is denoted $\mathcal{I}_{i,P}(s_i) := \{I_{i,P}^m(s_i) : m \in \mathbb{N}\}$.

A process P is *stationary* after round m if for every $k \in K$, every cycle $c \in C_{\infty,P}$ and $\underline{\omega}$ satisfying $\omega^0 \leq \underline{\omega}$,

$$\frac{\sum_{\boldsymbol{\omega}:\omega^m \in c} P(\boldsymbol{\omega} | k, \omega^1 = \underline{\omega})}{\sum_{\tilde{c} \in C_{\infty,P}} \sum_{\boldsymbol{\omega}:\omega^m \in \tilde{c}} P(\boldsymbol{\omega} | k, \omega^1 = \underline{\omega})} = \sum_{\boldsymbol{\omega} \in c_\infty^{-1}(c)} P(\boldsymbol{\omega} | k, \omega^1 = \underline{\omega}), \tag{5.5}$$

and for every player i,

$$\mathcal{I}_{i,P}(s_i) = \mathcal{I}_{i,P}(\hat{s}_i) \implies P(k, \omega^1 | s_i) = P(k, \omega^1 | \hat{s}_i), \quad P\, a.s. \tag{5.6}$$

We say that P is stationary if it is stationary after some finite round m. Stationarity thus imposes two requirements. First, it requires that, conditional on the first transition and k, the limit distribution of P be given by the distribution on terminal nodes after m rounds. Second, it requires the finite collection of information sets of a player's action set sequence to be a sufficient statistic for her belief about the first transition and k. Since action labels are constant on cycles, the first condition allows us to recover the outcome distribution of a process by considering its distribution on a finite truncation of all paths. Note that for a Markov process that is stationary after round m, every information set is reached by some sequence within the first m rounds. We will show that for stationary Markov processes, checking a finite number of obedience constraints is enough to establish if the process is also canonical.

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

We now show that SCAMP is a finite dimensional class of processes. For any information set $I \in \mathcal{I}_{i,P}(s_i)$ of player i's type s_i, define its maximal elements

$$\overline{I} = \{(k, \omega) \in I : \forall \omega' \text{ s.t. } \omega \leq \omega', (k, \omega') \notin I\} \tag{5.7}$$

Conditional obedience constraints for a player i with type s_i at $m \leq \mathbb{N}$ take the form

$$s_i^m = \mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^m}(P(\cdot, \cdot | s_i, \overline{I_{i,P}^m}(s_i)))), \quad P \text{ a.s.} \tag{5.8}$$

Lemma 5.1 (SCAMP is Finite Dimensional). *Every Markov process P that is stationary after round m on a Strategic Automaton that satisfies obedience and conditional constraints for all $n \leq m$ is SCAMP.*

We prove Lemma 5.1 as follows: By definition, a stationary Markov process only has a finite number of first-order beliefs. We need to show that if obedience and conditional obedience hold on the first m coordinates, obedience holds everywhere. To show this, we show that, conditional on k and the first transition ω^1, the longer a sequence stays at a given information set, the more probability it assigns to higher ranked automaton states. Since action labels of automaton states are monotonic with respect to this ranking, we conclude that if obedience holds on the first m rounds, then "sub-obedience" must hold everywhere: Every action label contains all best-replies to beliefs on lower action labels. We then use conditional obedience to show that this containment is never strict. By conditional obedience, no matter how much mass is assigned to the highest ranked automaton states in the information set, i's best reply remains unchanged. We then conclude that obedience must hold everywhere.

We show that there is a strategic automaton so that all the SCAMP it induces are sufficient to obtain all outcome distributions.

Theorem 5.1 (Sufficiency of SCAMP)**.** *For every finite game there exists a strategic automaton so that the set of SCAMP induce all outcome distributions.*

Depending on the game, the strategic automaton with this property may be larger than the automaton we constructed in Section 4.1. This can happen for three reasons:

(i) If the probability of leaving a state whose label contains more than one action for a player is greater than zero under a SCAMP, then that state will have probability zero in the corresponding limit distribution. Some outcome distributions may allow for non-singleton outcomes to be realized as well as strict subsets of them. In this case we may need to duplicate the states with non-singleton labels to allow for some paths to converge at that a state and others to leave the state.

(ii) As we have seen in Example 2, cycles with multiple automaton states can arise. For outcome distributions whose support on ICR-hierarchies is restricted, additional cycles may be required to achieve this outcome. Indeed, suppose we added a third state k = Bad state 2 to the payoffs in Example 2. At Bad state 2, payoffs are given by the matrix below

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

Bad state 2

With this additional state, both players' first order rationalizable action sets can be one of a, b or $\{a, b\}$. So in Figure 2, both automaton states in the cycle now have the same transitions. Then ICR-hierarchies can be described with a strategic automaton that has a cycle of a single state as the one shown in Figure 5 below:

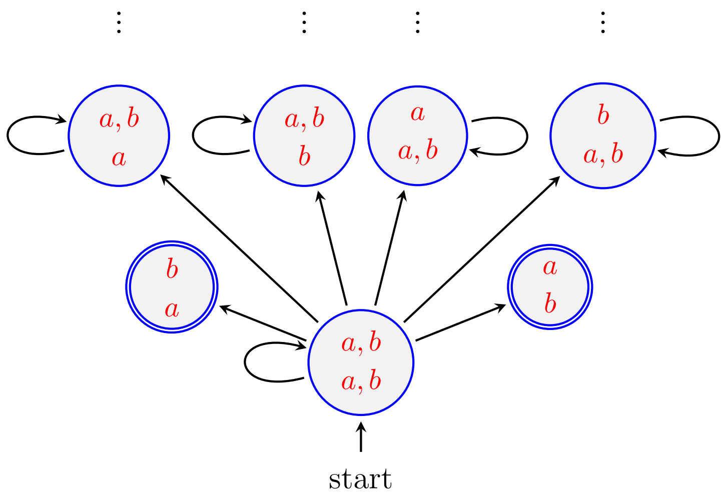

Figure 5: Cycle generating ICR-hierarchies when the Bad state 2 is added.

If we want to generate the outcomes when the probability of Bad state 2 is zero with SCAMP, we need the automaton with a larger cycle as the transitions with under this restriction are not Markovian with the cycle in Figure 5.

(iii) To obtain Stationarity, we want the automaton to be balanced after every initial state: For every initial state ω^0, the collection of paths that start at ω^0 and reach a terminal state pass through the same number of automaton states and cycles.

Conditions (i) - (iii) are extensions of the automaton we constructed in Section 4.1 and can be obtained by duplicating states. We provide a general construction for a strategic automaton that is large enough for SCAMP to implement every outcome in Appendix A.2. In particular, we provide a bound on the number of duplicates that are necessary.

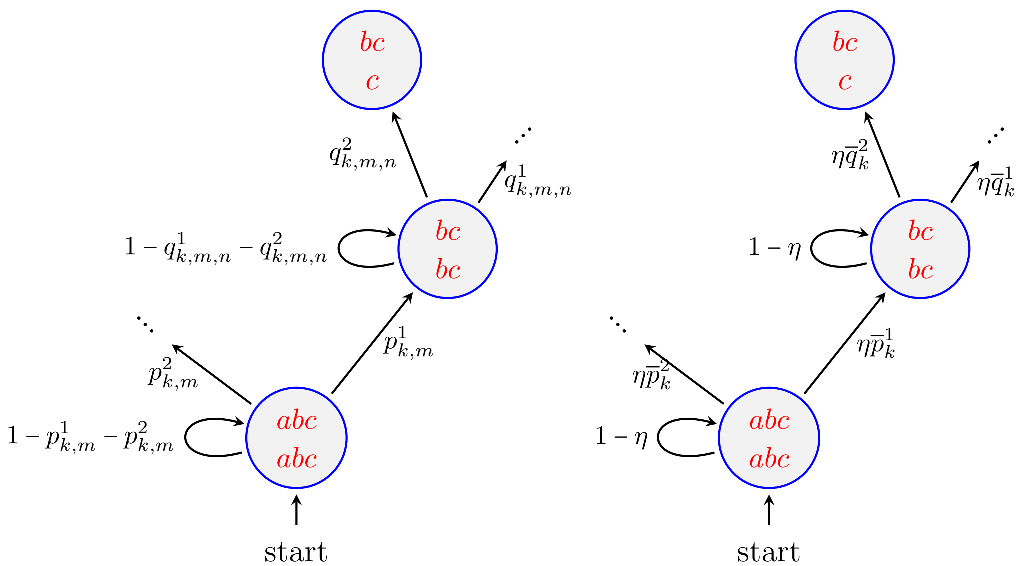

Figure 6: Averaging process to SCAMP.

For every $k \in K$, consider a process on the paths of the joint Automaton in the left panel of Figure 6.

1. We first compute the average exit probability for any pair of successive cycles.

$$\begin{aligned}\overline{p}^1_k &= \frac{\sum_m p^1_{k,m} P(m)}{\sum_m P(m)(p^1_{k,m} + p^2_{k,m})},\\ \overline{q}^2_k &= \frac{\sum_n \sum_m p^1_{k,m} P(m|n) q^2_{k,m,n} P(n)}{\sum_m p^1_{k,m} P(m) \sum_n (q^1_{k,m,n} + q^2_{k,m,n}) P(n|m)},\end{aligned} \tag{5.9}$$

where $P(m) = \prod_{l \le m}(1 - p^1_{k,l} - p^2_{k,l})$ and $P(n|m) = \prod_{l \le n}(1 - q^1_{k,m,l} - q^2_{k,m,l})$. This induces a Stationary Markov Prior $\overline{P}$ when we pick a constant cycling probability $\eta > 0$. Note that the limit distribution of the Stationary Markov Prior coincides with that of $(p^1_{k,m}, p^2_{k,m}, (q^1_{k,m,n}, q^2_{k,m,n})_n)_m$,

$$\overline{P}^\infty(bc, c) = \overline{p}^1_k \overline{q}^2_k = \sum_n \sum_m p^1_{k,m} P(m|n) q^2_{k,m,n} P(n) = P^\infty(bc, c). \tag{5.10}$$

The stationary distribution $\overline{P}^\infty$ of this Markov prior, that is, the limit distribution on the terminal states depends only on the exit probability $\overline{p}^l_k, \overline{q}^l_k$.

2. We then write down obedience constraints for the induced Stationary Markov prior and see that η enters monotonically in the obedience constraints. Consider a sequence for i player whose action is the bottom entry in each automaton state, $(abc, \dots, bc, \dots, bc, c, \dots)$. If player i considers two sequences possible: One where player $-i$ switched to c before her and one where $-i$ did not. An obedience constraint of player i at the round where she eliminates b on a game with payoffs $u_i(k, a_i, a_{-i})$ takes the form

$$\begin{aligned} 0 < \eta \overline{q}_k^2 p(k) &\min_{a_{-i} \in \{b,c\}} (u_i(k, c, a_{-i}) - u_i(k, b, a_{-i})) \\ &+ \overline{q}_k^1 p(k) \min_{a_{-i} \in \{c\}} (u_i(k, c, a_{-i}) - u_i(k, b, a_{-i})). \end{aligned} \tag{5.11}$$

If P satisfies obedience and $\eta = 1$, then by the linearity of our transformation, the above inequality holds. The cycling probability enters monotonically in every obedience constraint: The lower the cycling probability the higher the probability that higher-ranked states are reached at any given round. So a lower cycling probability encourages eliminating more actions: If an agent is eliminating an action at a certain round, she will do so in response to other agents eliminating actions. However, she doesn't know if she is on a path where other agents eliminated their actions before her or after. Paths, where she is the first to eliminate an action, will cycle longer through states where no agent changes actions than paths where she is not. A lower cycling probability means her beliefs assign less weight to her being the first player to eliminate the action and thus causing the path to leave the cycle. The fact that a lower cycling probability encourages eliminating more actions then follows from the monotonicity of the best-reply operator: The set of a player's best replies decreases if her beliefs assign more mass to smaller action sets. Hence the inequality (5.11) holds for any choice of η. An important condition for this to work is that the agent's information set, i.e. the sequences they consider possible does not change much. To ensure this, we need the right automaton structure. As we have seen, this procedure would not work if the left automaton had a cycle with two different transitions but the right automaton remains unchanged: The transitions at even rounds are different from the transitions at odd rounds. In Appendix A.2 we provide a general construction.

5.1 Characterization of Rationalizable Outcomes

Fix a strategic automaton verifying the statement in Theorem 5.1 and let $m^* = |\Omega|$. A distribution $p \in \Delta(K \times \Omega^{m^*})$ satisfies obedience constraints if for $0 < m \leq m^*$ and all $s_i = (\beta_i(\omega^0), \dots, \beta_i(\omega^{m^*}))$ in the support of p,

$$s_i^m = \mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^{m-1}}(p(\cdot, \cdot | s_i))). \tag{5.12}$$

Written out, expression (5.12) is a system of linear inequalities: For every action $a_i' \in s_i^{m-1} \smallsetminus s_i^m$ there is $a_i \in s_i^m$ so that

$$0 < \sum_{(k, s_{-i}, s_i)} p(k, s^{m-1} | s_i) \min_{\sigma_k \in s_{-i}^{m-1}} (u_i(k, a_i, \sigma_k) - u_i(k, a_i', \sigma_k)), \tag{5.13}$$

moreover, for every $a_i \in s_i^m$ and every $a_i' \in A_i$,

$$0 \leq \sum_{(k, s_{-i})} p(k, s_{-i}^{m-1} | s_i) \max_{\sigma_k \in s_{-i}^{m-1}} (u_i(k, a_i, \sigma_k) - u_i(k, a_i', \sigma_k)). \tag{5.14}$$

Let $\mathcal{O}^{m^*} \subseteq \Delta(K \times \Omega^{m^*})$ denote the set of distributions on $K \times \Omega^{m^*}$ that satisfy obedience constraints. Let the set of terminal states of $p \in \Delta(K \times \Omega^{m^*})$ be given by

$$\overline{X}(p) = \{(k, \omega^{m^*}) \in K \times \overline{\Omega} : p(k, (\omega^0, \dots, \omega^{m^*})) > 0\}, \tag{5.15}$$

Letting $p^{m^*}(k, \omega) = \sum_{x \in \Omega^{m^*} : x^{m^*} = \omega} p(k, x)$, define the limit probability of p

$$\overline{p}(k, b) = \frac{\sum_{\omega : \beta(\omega) = b} p^*(k, \omega)}{p^{m^*}(\overline{X}(p))}. \tag{5.16}$$

The limit probability $\overline{p}$ satisfies limit-obedience if for every b in its support and every player i, $b_i = \mathrm{br}_i(\overline{p}(\cdot, \cdot | b_i))$. Let $\mathcal{O}^\infty \subseteq \Delta(K \times B)$ denote the set of probabilities satisfying limit-obedience.

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

$$\mathcal{O} = \{\overline{p} : p \in \mathcal{O}^{m^*}\} \cap \mathcal{O}^\infty. \tag{5.17}$$

The relative closure of this set is a convex polyhedron:

Lemma 5.2 (Linearity of $\mathcal{O}$)**.** *The relative closure of the set $\mathcal{O}$ is a convex polyhedron.*

Indeed, clearly $\mathcal{O}^{m^*}$ and $\mathcal{O}^{\infty}$ are polyhedra. Then the simple coordinate projection of the polyhedron $\mathcal{O}^{m^*}$ of distributions onto the coordinates consisting of the set of terminal paths and K is a lower dimensional polyhedron. The only issue is that the scale is wrong. So consider the cone generated by this lower dimensional polyhedron. For each point in the lower dimensional polyhedron there is a unique positive number that scales it into a probability. So the intersection with the cone and the simplex gives us the space we are looking for. The cone and the simplex are polyhedra. So their intersection is too. We now show that $\mathcal{O}$ coincides with the set of all outcome distributions.

Theorem 5.2 (Necessity and Sufficiency of $\mathcal{O}$ for SCAMP)**.**

(i) Every SCAMP P induces a distribution $p_P \in \mathcal{O}$ through its marginal on $K \times \Omega^{m^*}$ so that the limit probability $\overline{p}_P$ coincides with the outcome distribution of P.

(ii) For every $p \in \mathcal{O}$ there exists SCAMP P_p so that the limit probability $\overline{p}$ coincides with the outcome distribution of P_p.

Part (i) of Theorem 5.2 is an immediate consequence of the fact that P is SCAMP. Part (ii) is constructive and follows the same steps as the proof of Theorem 5.1. We thus obtain our characterization of outcome distributions:

Corollary 5.1 (Linearity of Outcomes)**.** *For every finite game, the set of outcome distributions is equal to $\mathcal{O}$. Its relative closure is a convex polyhedron.*

References

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 (2005): "Topologies on types," *Harvard Institute of Economic Research Discussion Paper*.

——— (2007): "Interim correlated rationalizability," *Theoretical Economics*.

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

Gossner, O. and J.-F. Mertens (2001): "The Value of information in Zero-Sum Games," mimeo.

Harsanyi, J. C. (1967): "Games with incomplete information played by "Bayesian" players, Part I. The basic model," *Management science*, 14, 159–182.

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.

Mertens, J.-F. and S. Zamir (1985): "Formulation of Bayesian analysis for games with incomplete information," *International journal of game theory*, 14, 1–29.

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 (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

A.1 Revelation Principle

Theorem 4.2 *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*:*

$$s_i^m = \mathrm{br}_i(P_S(k, s_{-i}^{m-1}|s_i)) \text{ a.s. in } s_i = (s_i^m)_m \tag{A.1}$$

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 $((\mathrm{R}_i^m)_m : T_i \to S_i)_i$ satisfies: for every player i, $t_i \in T_i$ and $m \in \mathbb{N}$,

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

Then note that for every $b_i \in B_i$, $\mathrm{br}_i(b_i)^{-1} \subseteq \Delta(K \times B_{-i})$ is convex. Write $\tilde{P} \coloneqq P \circ (\mathrm{id} \times \mathrm{R})^{-1}$ with conditional probabilities $(\tilde{P}_i : S_i \to \Delta(K \times S_{-i})_i$ 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 Strategic Automata

A *monotone stochastic transformation* for player i is a map $\rho_i : K \times B_{-i} \to \Delta(B_{-i})$ 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* $\rho_i: K \times B_{-i} \to \Delta(B_{-i})$ *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* $b_{-i} \in B_{-i}$,

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

Proof. Consider any conjecture $\sigma_i: K \times B_{-i} \to \Delta(A_{-i})$ 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 $\sigma_i \circ \rho_i$, given for every $a_{-i} \in A_{-i}, k \in K, b'_{-i} \in B_{-i}$ by

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

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). $\square$

Let $B^{(0)}$ and $B^{(1)}$ be two copies of the set B of action set profiles. Consider a distribution $p \in \Delta(K \times (B^{(0)} \times B^{(1)}))$ on states of nature and transitions on action set profiles so that for every i and every transition (s_i^0, s_i^1) in the support of p,

$$\mathrm{br}_i(\mathrm{marg}_{K \times b_{-i}^{(0)}} p(\cdot, \cdot | s_i^0, s_i^1)) = s_i^1. \tag{A.9}$$

Let $\mathbb{P}$ denote the collection of initial transition probabilities $p \in \Delta(K \times B \times B)$ satisfying this best response condition.

Definition A.1 (p-Best-Reply Hierarchies)**.** *We define the set* S_p *of* p*-Best-Reply Hierarchies recursively as follows:*

$$\begin{aligned} S_p^m &:= \left\{ (s^0, s^1, \ldots, s^m) \in B^{m+1} : \begin{array}{c} \exists\ p^m \in \Delta(K \times (S_p^m))\ s.t. \\ \mathrm{marg}_{k,(s^0,s^1)}(p^m) = p, \\ \forall\ i,\ \forall\ 0 < \ell \leq m, \\ s_i^\ell = \mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^{\ell-1}}(p(\cdot, \cdot | s_i^m)) \end{array} \right\} \\ S_p &:= \{(s^0, s^1, \ldots) : \forall\ m \geq 0,\ (s^0, \ldots, s^m) \in S_p^m\}. \end{aligned} \tag{A.10}$$

Let

$$\mathbb{S}_p := \bigcup_{m=1}^{\infty} S_p^m, \tag{A.11}$$

denote the collection of all truncated sequences in S_p. Define the transition correspondence on truncated sequences $(s^0, \ldots, s^m) \in \mathbb{S}_p$

$$\kappa_p(s^0, \ldots, s^m) := \{b \in B : (s^0, \ldots, s^m, b) \in S_p^{m+1}\}. \tag{A.12}$$

By the revelation principle (Lemma 4.2), for any given $p \in \mathbb{P}$, the sequences S_p coincide with the set of ICR-hierarchies that can arise from common priors $P \in \mathcal{P}$ satisfying

$$P \circ (\mathrm{id} \times (\mathrm{R}^0 \times \mathrm{R}^1))^{-1} = p. \tag{A.13}$$

For any collection of action set profiles $B' \subseteq B$ define its *maximal* subset $\max B'$ as the collection of $b \in B'$ so that there does not exist $\hat{b} \in B' \setminus \{b\}$ satisfying

$$b_i \subseteq \hat{b}_i, \ \forall \ i. \tag{A.14}$$

Let the maximal transition correspondence be defined as

$$\bar{\kappa}_p(s^0, \ldots, s^m) := \max \kappa_p(s^0, \ldots, s^m), \ \forall \ (s^0, \ldots, s^m) \in \mathbb{S}_p. \tag{A.15}$$

For every truncated sequence $(s^1, \ldots, s^m) \in \mathbb{S}_p$, the set of *sequentially maximal tails* is defined as the following collection of sequences,

$$\bar{S}_p(s^0, \ldots, s^m) := \left\{ \hat{s} \in S_p : \begin{array}{c} \forall \ \ell \leq m, \ \hat{s}^\ell = s^\ell, \\ \forall \ \ell > m, \ \hat{s}^\ell \in \bar{\kappa}(\hat{s}^0, \ldots, \hat{s}^{\ell-1}) \end{array} \right\}. \tag{A.16}$$

Claim A.2. *For all $m \in \mathbb{N}$ and $p \in \mathbb{P}$, $|\overline{S}_p(s^0, \ldots, s^m)| \leq 2^{|B|}$.*

Proof. For any $s \in \overline{S}_p(s^0, \ldots, s^m)$ let

$$B^m(s) := \{s^n : n \geq m\}. \tag{A.17}$$

Then for any distinct $s, \hat{s} \in \overline{S}_p(s^0, \ldots, s^m)$ there must be a pair $(b, \hat{b}) \in B^m(s) \times B^m(\hat{s})$ so that neither $b \subseteq \hat{b}$ nor $\hat{b} \subseteq b$. Otherwise, the sequence which first reaches a round with an action set profile that is strictly contained in that of the other sequence would not make a maximal transition and thus not be element of $\overline{S}_p(s^0, \ldots, s^m)$. But then we have that there can only be as many sequences as there are families $B' \subseteq s^m$ where all elements inside B' are ordered according to set-inclusion but contain elements which are not comparable with some element in every other such family. The number of such families is clearly bounded by the total number of subsets of B. □

For every subset $C \subseteq S_p \times B^{\mathbb{N}}$, define the best-response operator on sequence pairs,

$$\mathbf{B}_p(C) := \left\{ (s, \bar{s}) \in S_p \times B^{\mathbb{N}} : \begin{array}{c} \exists\ p^\infty \in \Delta(K \times C) \text{ s.t.} \\ \operatorname{marg}_{k,(s^0,s^1)}(p^\infty) = p, \\ \forall\ i,\ \forall\ m > 0, \\ s_i^m = \operatorname{br}_i(\operatorname{marg}_{k,s_{-i}^{m-1}}(p^\infty(\cdot,\cdot|s_i))) \\ \bar{s}_i^m = \operatorname{br}_i(\operatorname{marg}_{k,\bar{s}_{-i}^{m-1}}(p^\infty(\cdot,\cdot|s_i,\bar{s}_i))) \end{array} \right\}. \tag{A.18}$$

Let $\mathcal{S}_{1,p} := \overline{S}_p(s^0)$ denote the collection of sequentially maximal tails. These are first-order maximal sequences.

$$\mathcal{R}_{1,p}^1 := \left\{ (s, \bar{s}) \in \mathcal{S}_{1,p} \times B^{\mathbb{N}} : \bar{s} \in \bigcup_{\tilde{s}^1 \in \kappa_p(s^0)} \overline{S}_p(s^0, \tilde{s}^1) \right\} \tag{A.19}$$

Given $\mathcal{R}_{1,p}^{m-1}$, define

$$\mathcal{R}_{1,p}^m := \left\{ (s, \bar{s}) \in \mathbf{B}_p(\mathcal{R}_{1,p}^{m-1}) : \begin{array}{c} s \in \mathcal{S}_{1,p}, \\ \bar{s} \in \bigcup_{\tilde{s}^m \in \kappa_p(s^0,\dots,s^{m-1})} \overline{S}_p(s^0, \dots, s^{m-1}, \tilde{s}^m) \end{array} \right\} \tag{A.20}$$

Given $(\mathcal{R}_{l-1,p}^m)_{m \in \mathbb{N}}$, we define the set of l-order maximal sequences

$$\mathcal{S}_{l,p} := \bigcup_{m \in \mathbb{N}} \{ (s^0, \dots, s^{m-1}, \bar{s}^0, \bar{s}^1, \dots) : (s, \bar{s}) \in \mathcal{R}_{l-1,p}^m \} \tag{A.21}$$

sequence $(\mathcal{R}_{l,p}^m)_{m \in \mathbb{N}}$ as follows: For $m < l$, let $\mathcal{R}_{l,p}^m := \mathcal{R}_{l-1,p}^m$. For $m \geq l$,

$$\mathcal{R}_{l,p}^m := \left\{ (s, \bar{s}) \in \mathbf{B}_p(\mathcal{R}_{l,p}^{m-1}) : \begin{array}{c} s \in \mathcal{S}_l, \\ \bar{s} \in \bigcup_{\tilde{s}^m \in \kappa_p(s^0,\dots,s^{m-1})} \overline{S}_p(s^0, \dots, s^{m-1}, \tilde{s}^m) \end{array} \right\}. \tag{A.22}$$

Note that construction $\mathcal{R}_{l,p}^m$ is "Markovian" in the sense that the tails $\bar{s}$ that are being added only depend on beliefs with support on tails that were added at $\mathcal{R}_{l,p}^{m-1}$. We now show that all sequences in S_p are obtain this way:

Claim A.3. $S_p = \bigcup_{l \in \mathbb{N}} \mathcal{S}_{l,p}$.

Proof. We will show that for all $m \leq l$,

$$\begin{aligned} \{ (s^0, \dots, s^m) : s \in \mathcal{S}_{l,p} \} &= \{ (s^0, \dots, s^m) : s \in S_p \}, \\ \mathcal{S}_{l,p} &= \{ s \in \mathcal{S}_{l,p} : \exists\ \bar{s} \text{ s.t. } (s, \bar{s}) \in \mathbf{B}_p(\mathcal{R}_{l,p}^m) \}. \end{aligned} \tag{A.23}$$

For $l = 1$, the first condition holds by definition fo $\mathcal{S}_{1,p}$. Moreover, by the monotonicity of br_i, (Claim A.1) we conclude that $\mathcal{S}_{1,p}$ is best-reply closed and so the second condition also holds. Suppose now that both conditions hold for $l-1$. Then for $m \leq l$,

$$\{(s^0, \ldots, s^{m-1}) : \exists\ \bar{s} \text{ s.t. } (s, \bar{s}) \in \mathcal{R}^m_{l-1,p}\} = \{(s^0, \ldots, s^{m-1}) : s \in S_p\}. \tag{A.24}$$

Then

$$\left\{(s^0, \ldots, s^{m-1}, \tilde{s}^m) : \begin{array}{c} s \in \mathcal{S}_{l-1,p}, \\ \tilde{s}^m \in \kappa_p(s^0, \ldots, s^{m-1}) \end{array}\right\} = \{(s^0, \ldots, s^m) : s \in S_p\}. \tag{A.25}$$

Since the second condition in (A.23) holds for $l-1$ we have that $\mathcal{R}^m_{l-1,p} \subseteq \mathcal{R}^m_{l,p}$. Then it must be that

$$\{(s^0, \ldots, s^m) : \mathcal{S}_{l,p}\} = \{(s^0, \ldots, s^m) : s \in S_p\}, \tag{A.26}$$

and so the result follows. □

Claim A.4. *There is finite z so that for all $p \in \mathbb{P}$ and for all $m, l \in \mathbb{N}$,*

$$\mathcal{R}^m_{l,p} = \mathcal{R}^{m+z}_{l,p}.$$

Proof. Fix $p \in \mathbb{P}$. Define the set of sequences in $\mathcal{S}_{l,p}$ that branch out of $\mathcal{S}_{l-1,p}$ at round $n \in \mathbb{N}$,

$$\hat{\mathcal{S}}_{l,n,p} := \mathcal{S}_{l-1} \cup \{s \in \mathcal{S}_{l,p} : \forall\ \hat{s} \in \mathcal{S}_{l-1,p},\ s^n \neq \hat{s}^n\}. \tag{A.27}$$

Note that for every $l \in \mathbb{N}$, every $n \in \mathbb{N}$ and for every sequence $s \in \hat{\mathcal{S}}_{l,n,p}$, there is a distribution $q \in \Delta(K \times \hat{\mathcal{S}}_{l,n-1,p} \cup \mathcal{S}_{l-1,p})$ so that for all $m \in \mathbb{N}$ and every player i,

$$s^m_i = \mathrm{br}_i(\mathrm{marg}_{k, s^{m-1}_{-i}}(q(\cdot, \cdot | s_i))). \tag{A.28}$$

We proceed by induction on l. Let $l = 2$. From Claim A.2 we obtain that the cardinality of the set $\hat{\mathcal{S}}_{2,n,p}$ is bounded by $2^{|B|}|\mathcal{S}_2| < \infty$ for all $n \in \mathbb{N}$. Moreover, we established in the proof of Claim A.3 that $\mathcal{S}_{1,p}$ is best-reply closed by appealing to the monotonicity property of br_i (Claim A.1), i.e. for all $m \in \mathbb{N}$, $\mathcal{S}_{1,p} \subseteq \{s \in S_p : \exists \bar{s} \text{ s.t. } (s, \bar{s}) \in \mathbf{B}_p(\mathcal{R}^m_{1,p})\}$. Hence $\mathbf{B}_p$ will start to cycle after some number $z_{1,p} \in \mathbb{N}$ of iterations on $\mathcal{R}^1_{1,p}$: $\mathcal{R}^m_{1,p} = \mathcal{R}^{m+z_{1,p}}_{1,p}$, for all $m \in \mathbb{N}$. For any l, note again that the set $\hat{\mathcal{S}}_{l,n,p}$ is bounded by $2^{|B|}|\mathcal{S}_{l,p}| < \infty$. Since $\mathcal{S}_{l,p}$ is a finite, best-reply closed set (see proof of Claim A.3), we deduce that $\mathbf{B}_p$ will again start to cycle after some number $z_{l,p}$ of iterations on $\mathcal{R}^1_{l,p}$: $\mathcal{R}^m_{l,p} = \mathcal{R}^{m+z_l}_{l,p}$.

Finally, note that there is $\ell_p < |B|$ so that for all $l \geq \ell_p$ and all $m \in \mathbb{N}$, $\mathcal{R}^m_{l,p} = \mathcal{R}^m_{l+1,p}$ - otherwise there would be a sequence which makes infinitely many eliminations. Note that this bound is uniform over all $p \in \mathbb{P}$. Since $|\hat{\mathcal{S}}_{n,l,p}| \leq 2^{|B|}$ also holds for all choices of $p \in \mathbb{P}$, we deduce that $\sup_{p\in\mathbb{P}} z_{\ell_p,p} < \infty$. But then there is a finite z (e.g. the least common multiple of all the finite collection $\{z_{\ell_p,p} : p \in \mathbb{P}\}$) so that $\forall\ p \in \mathbb{P},\ \forall\ m, l \in \mathbb{N},\ \mathcal{R}^m_{l,p} = \mathcal{R}^{m+z}_{l,p}$. $\square$

Claim A.5 (Automaton for S_p)**.** *There is an automaton $(\Omega, \beta, \preceq)$ so that for every $p \in \mathbb{P}$ and every sequence $s \in S_p$, there is a path $(\omega^0, \omega^1, \dots)$ so that*

$$s = (\beta(\omega^0), \beta(\omega^1), \dots).$$

Proof. For every $p \in \mathbb{P}$, every truncated sequence $(s^0, \dots, s^m) \in \mathbb{S}_p$ and profile $b \in B$ define

$$m(b|s^0, \dots, s^m) := |\{\ell \leq m : b = s^\ell\}|. \tag{A.29}$$

Define the summary statistic for z satisfying the statement in Claim A.4

$$\xi_z(s^0, \dots, s^m) := ((s^0, \dots, s^{\min\{z,m\}}), \hat{\xi}_z(s^0, \dots, s^m)), \tag{A.30}$$

where $\hat{\xi}_z(s^0, \dots, s^m) := \{(s^\ell, m(s^\ell|s^0, \dots, s^m) \mod z) : \ell \leq m\}$. Then for every $p \in \mathbb{P}$ and truncated sequences $h, h' \in \mathbb{S}_p$

$$\xi_z(h) = \xi_z(h') \implies \kappa_p(h) = \kappa_p(h'). \tag{A.31}$$

By Claim A.3 we obtain an automaton $(\Omega, \beta, \preceq)$, where

$$\begin{aligned} \Omega &:= \bigcup_{p\in\mathbb{P}} \{\xi_z(h) : h \in \mathbb{S}_p, p \in \mathbb{P}\} \\ \beta(\xi_z(s^0, \dots, s^m)) &:= s^m,\ \forall\ s \in S, m \in \mathbb{N}, \\ \omega \preceq \omega' &\iff \exists s \in S \text{ s.t.} \\ &\qquad \omega = \xi_z(s^0, \dots, s^m) \text{ and} \\ &\qquad \omega' = \xi_z(s^0, \dots, s^{m+1}). \end{aligned} \tag{A.32}$$

$\square$

For every $P \in \mathcal{P}$ let its support in S be denoted $S_P := \{s \in S : \sum_{k\in K} P(k,s) > 0\}$. Furthermore, let $p = \mathrm{marg}_{k,s^0,s^1}(P)$. Since $S_P \subseteq S_p$, we obtain the following corollary.

Corollary A.1. *There is an automaton* $(\Omega, \beta, \leq)$ *so that for every* $P \in \mathcal{P}$ *and every sequence* $s \in S_P$*, there is a path* $(\omega^0, \omega^1, \dots)$ *so that*

$$s = (\beta(\omega^0), \beta(\omega^1), \dots).$$

The construction has shown the following three properties:

(i) If a cycle was left, it is never revisited.

(ii) If s and s' reach a state ω at rounds m and n respectively, then the set of states visited by s before round m equals the set of states visited by s' before round n.

(iii) All cycles have constant size z and for every cycle c and $\omega \in c$ there is a unique state $\omega' \in c$ so that $\omega \leq \omega'$.

We will use properties (i)-(iii) to establish the sufficiency of information design via SCAMP and obtain a linear characterization of all achievable outcome distributions.

A.3 Sufficiency of SCAMP

For $P \in \mathcal{P}$, we denote by ν_P its outcome distribution. Say that $P, P' \in \mathcal{P}$ are *outcome-equivalent* if $\nu_P = \nu_{P'}$.

We say that an outcome distribution is *stronger* than another if conditional on a state of nature the latter first order stochastically dominates the former with respect to set-inclusion for all players:

Definition A.2 (Stronger Outcome Distribution)**.** *For* $\nu, \nu' \in \Delta(K \times B)$ *we say that* ν *is* stronger *than* ν' *if there exists monotone stochastic transformation* $\rho: K \times B \to \Delta(B)$ *so that*

$$\nu(k, b) = \nu' \circ \rho(k, b) := \sum_{b'} \rho(b | k, b') \nu'(k, b'). \tag{A.33}$$

Definition A.3 (Strongest Outcome Distribution)**.** *An outcome distribution* $\nu \in \mathcal{V}$ *is* strongest *if every outcome distribution* $\nu' \in \mathcal{V}$*, which is stronger than* ν *also satisfies* $\nu = \nu'$*.*

Let $\mathcal{P}^* \subseteq \mathcal{P}$ denote the collection of priors in $\mathcal{P}$ with a strongest outcome distribution.

A.3.1 Finiteness of SCAMP

Let P be Stationary and Markov. For every player i, type s_i and information set $I \in \mathcal{I}_{i,P}(s_i)$, let $M(I, s_i) \in \mathbb{N} \cup \{\infty\}$ denote the number of rounds m so that $I = I_P^m(s_i)$,

$$M(I, s_i) := |\{m \in \mathbb{N} : I = I_P^m(s_i)\}|, \tag{A.34}$$

where the cardinality of a countable set is set to ∞. Let $\underline{m}(s_i) := \min\{m \in \mathbb{N} : I = I_P^m(s_i)\}$ denote the first round where s_i has information set I.

Define the *transitive closure* of $\preceq$, denoted $\overline{\preceq}$, as follows: $\omega \overline{\preceq} \omega'$ if and only if there is a list of states $\omega_0, \ldots, \omega_n$ so that

$$\omega = \omega_0 \preceq \cdots \preceq \omega_n = \omega'. \tag{A.35}$$

Note that cycles correspond to the equivalence classes of $\overline{\preceq}$.

Say that a process P satisfies *sub-obedience* if for every $m \in \mathbb{N}$ and every player i,

$$\mathrm{br}_i(\mathrm{marg}_{k, s_{-i}^{m-1}}(P(\cdot, \cdot | s_i))) \subseteq s_i^m, \ P \text{ a.s.} \tag{A.36}$$

Claim A.6. *If P is Stationary after round m^* and Markov and satisfies obedience for all players' types on rounds $m \le m^*$ then it satisfies sub-obedience for all $n \in \mathbb{N}$.*

Proof. For any type s_i, $I \in \mathcal{I}_{i,P}(s_i)$ and $x \in \mathbb{N}$ let $s_i(I, x)$ be the type satisfying

$$s_i(I,x)^m := \begin{cases} s_i^m, \text{ if } m \le \underline{m}(s_i) + M(I, s_i) \\ s_i^{\underline{m}(s_i) + M(I, s_i)} \text{ if } m \in \{\underline{m}(s_i) + M(I, s_i), \ldots, \underline{m}(s_i) + M(I, s_i) + x\} \\ s_i^{m-x}, \text{ otherwise.} \end{cases}$$

Since P is Markov we have that there is an infinite collection $N(s_i, I) \subseteq \mathbb{N}$ of numbers x so that $s_i(I, x)$ has positive probability under P.

Suppose $M(I, s_i) > 1$. Then there is a cycle c so that $I \cap (K \times c)$ is non-empty and consists of states which are minimal in I with respect to the partial pre-order $\preceq$ (otherwise I would be an information set for at most one round). Since every state in a cycle has a unique successor in the cycle, to each such cycle we can associate a cycling probability $\eta_{c,k,P} \in [0, 1]$ under P,

$$\eta_{c,k,P} := \prod_{\omega \in c} P(c | k, \omega). \tag{A.37}$$

For any $\omega' \in c$ and path $\boldsymbol{\omega}$ define $\overline{m}(\boldsymbol{\omega}, c) := \min\{m \in \mathbb{N} : \omega^m \in c\}$, where $\min\{\varnothing\} := \infty$. Moreover, define $\mathcal{H}_P^m(c, \omega') := \{\boldsymbol{\omega} : \overline{m}(\boldsymbol{\omega}, c) = m, \ \omega^m = \omega'\}$.

Then for any list of distinct successors $\omega_0 \leq \cdots \leq \omega_{n-1} \leq \omega_n$, where $\omega_0, \ldots, \omega_{n-1} \in c$ and $m, \underline{m} \in \mathbb{N}$ where $m \geq \underline{m}$, $\omega_0 \in \{\hat{\omega} : \exists\ k \in K,\ (k, \hat{\omega}) \in I^{\underline{m}}_{i,P}(s_i)\}$ and $\beta_i(\omega_n) = s_i^m, \ldots, \beta_i(\omega_0) = s_i^{m-n}$,

$$\begin{aligned}
P(\{\boldsymbol{\omega} : \omega^{m-1} = \omega_{n-1}, \omega^m = \omega_n\} | k, s_i, \mathcal{H}^{\underline{m}}_P(c, \omega_0)) &= \frac{\eta^{\ell_{m,\underline{m}}}_{c,k,P} \prod_{h=1}^{n-1} P(\omega_{h+1}|k, \omega_h)}{\sum_{\boldsymbol{\omega} \in \mathcal{H}^{\underline{m}}_P(c,\omega_0) : \beta(\boldsymbol{\omega}) = s_i} P(\boldsymbol{\omega}|k)} \\
&= \frac{\eta^{\ell_{m,\underline{m}}}_{c,k,P} \prod_{h=1}^{n-1} P(\omega_{h+1}|k, \omega_h)}{\sum_{l \leq \ell_{m,\underline{m}}} \eta^{l}_{c,k,P} \prod_{h=1}^{n-1} P(\omega_{h+1}|k, \omega_h)},
\end{aligned}$$

where $\ell_{m,\underline{m}} \coloneqq \frac{m - \underline{m} - n}{|c|}$ is the looping number. Since c is a minimal cycle in I, we may pick $\underline{m} = \underline{m}(s_i)$ and $m = \underline{m}(s_i) + M(I, s_i) + x$. Deduce that for any $x, x' \in N(s_i, I)$ so that $x \leq x'$,

$$\begin{aligned}
&\operatorname{marg}_{\omega^{\underline{m}(s_i)+M(I,s_i)+x}} P(\{\omega' : \omega \leq \omega'\} | k, s_i(I, x), \mathcal{H}^{\underline{m}(s_i)}_P(c, \omega_0)) \\
&\qquad \leq \operatorname{marg}_{\omega^{\underline{m}(s_i)+M(I,s_i)+x'}} P(\{\omega' : \omega \leq \omega'\} | k, s_i(I, x'), \mathcal{H}^{\underline{m}(s_i)}_P(c, \omega_0)).
\end{aligned} \tag{A.38}$$

Since β is monotone with respect to $\leq$, as x increases, beliefs conditional on k and conditional on paths entering cycle c at ω_0 become weakly stronger. Since P is stationary, $P(k, \omega^1 | s_i(I, x)) = P(k, \omega^1 | s_i(I, x'))$ and so

$$\operatorname{marg}_{\omega^{\underline{m}(s_i)+M(I,s_i)+x'}} P(\cdot, \cdot | s_i(I, x'), \mathcal{H}^{\underline{m}(s_i)}_P(c, \omega_0)) \circ (\mathrm{id} \times \beta_{-i})^{-1} \tag{A.39}$$

is weakly stronger than

$$\operatorname{marg}_{\omega^{\underline{m}(s_i)+M(I,s_i)+x}} P(\cdot, \cdot | s_i(I, x), \mathcal{H}^{\underline{m}(s_i)}_P(c, \omega_0)) \circ (\mathrm{id} \times \beta_{-i})^{-1}. \tag{A.40}$$

Stationarity implies that $P(\mathcal{H}^{\underline{m}(s_i)}_P(c, \omega_0) | k, \omega^1, s_i(I, x))$ and $P(k, \omega^1 | s_i(I, x))$ are both independent of x and so beliefs on I conditional on $s_i(I, x)$ become stronger as x. Finally, note that m^* must be large enough so that every information set is reached by some type before round m^*. Hence the result follows from the monotonicity of br_i established in Claim A.1. □

Claim A.7. *Let P be a Stationary after round m^* and Markov. If P satisfies sub-obedience constraints at all rounds and conditional-obedience for all $m \leq m^*$ then it is SCAMP.*

Proof. Fix $m \in \mathbb{N}$ and s_i. We will show that obedience holds at round m. By sub-obedience, we have that

$$\mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i))) \subseteq s_i^{m+1}, P \text{ a.s.} \tag{A.41}$$

By stationarity, we have that for every $n \leq m^*$, $\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i))$ is a weaker distribution than $\mathrm{marg}_{k,s_{-i}^n}(P(\cdot,\cdot|\hat{s}_i,\overline{I_{i,P}^m(s_i)}))$. To see this, note the result holds for terminal information sets. Define

$$X_k := \{(k,\boldsymbol{\omega}) : \exists\ \omega' \precsim \omega^n \text{ s.t. } (k,\omega') \in \overline{I_{i,P}^m(s_i)}\}. \tag{A.42}$$

Then stationarity implies that $P(\cdot|X_k, C_\infty) = P|_{X_k} \circ (\mathrm{id} \times c_\infty)^{-1}$, where $P|_{X_k}(\boldsymbol{\omega}) := P(\boldsymbol{\omega}|X_k)$ and $C_\infty := \cup_{\boldsymbol{\omega}:P(\boldsymbol{\omega})>0} c_\infty(\boldsymbol{\omega})$. Recall that partial pre-order $\preceq$ has the property that every automaton state has a unique, strict predecessor. Then we have that

$$P(\omega|\overline{I_{i,P}^m(s_i)}) = \sum_{c:\omega \precsim c} P(c|X_k, C_\infty) \tag{A.43}$$

But then by conditional obedience,

$$s_i^{m+1} = \mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i,\overline{I_{i,P}^m(s_i)}))). \tag{A.44}$$

But then by the monotonicity property of br_i established in Claim A.1 we conclude that

$$\mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i,\overline{I_{i,P}^m(s_i)}))) \subseteq \mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i))). \tag{A.45}$$

But then we have obedience, i.e. $\mathrm{br}_i(\mathrm{marg}_{k,s_{-i}^m}(P(\cdot,\cdot|s_i))) = s_i^{m+1}$. But then P is canonical on all cycles, which thus implies that its distribution on limit cycles $P \circ (\mathrm{id} \times c_\infty)^{-1}$ induces the outcome distribution on P, i.e.

$$\nu_P = P \circ (\mathrm{id} \times \beta \circ c_\infty)^{-1}. \tag{A.46}$$

□

We thus obtain the finiteness result for SCAMP:

Lemma 5.1 *Every Stationary Markov prior P on a Strategic Automaton that that satisfies obedience constraints and limit-obedience for all $n \leq |\Omega|$ is SCAMP.*

Proof. Immediate consequence of Claim A.6 and Claim A.7. □

A.3.2 SCAMP Automaton

Definition A.4 (Automaton Representation)**.** *For every $P \in \mathcal{P}$ and Strategic Automaton $(\Omega, \beta, \preceq)$, define an automaton representation of sequences $s \in S_P$ as paths via a state map*[4]*: $\tau \colon s \mapsto (\omega^0, \omega^1, \dots)$, where $s^m = \beta(\omega^m)$ for all m.*

For any $m \in \mathbb{N}$, let $\tau^m(s) := \mathrm{proj}_{\omega^m}(\tau(s))$. Let P_τ denote the induced distribution on K and automaton paths. For a given choice of representation for each $P \in \mathcal{P}$, let $\mathcal{P}_\tau \subseteq \Delta(K \times \Omega^{\mathbb{N}})$ denote the set of induced automaton representations. Depending on the automaton, representations need not be unique.

Definition A.5. *Say that a Strategic Automaton $(\hat{\Omega}, \hat{\beta}, \hat{\preceq})$ is a* sub-automaton *of $(\Omega, \beta, \preceq)$ if 1) $\hat{\Omega} \subseteq \Omega$, 2) for all $\omega \in \hat{\Omega}$, $\hat{\beta}(\omega) = \beta(\omega)$ and 3) for all $\omega, \omega' \in \hat{\Omega}$, $\omega \hat{\preceq} \omega' \iff \omega \preceq \omega'$.*

Fix a prior $P \in \mathcal{P}$ with a representation $\hat{\tau}$ on a Strategic Automaton $(\hat{\Omega}, \hat{\beta}, \hat{\preceq})$. We now construct an extension $(\Omega, \beta, \preceq)$ on which P will admit an outcome equivalent SCAMP.

We now define a new map τ on truncated sequences $(s^0, \dots, s^m)$.

For any $\omega' \in \hat{\Omega}$, let $m^1_{\omega'}$ be the shortest truncation so that ω' is reached

$$m^1_{\omega'} := \min\{m \in \mathbb{N} : \exists\ s \in S_P \text{ s.t. } \tau^m(s) = \omega'\}. \tag{A.47}$$

Consider the action set transitions at any round m

$$\kappa^m_{\omega'} := \bigcup_{s \in S_P \text{ s.t. } \tau(s^0, \dots, s^m) = \omega'} \kappa(s^0, \dots, s^m) \tag{A.48}$$

Define

$$\begin{aligned} N(\omega', 1) &:= \{n \geq m^1_{\omega'} : \kappa^n_{\omega'} \subseteq \kappa^{m^1_{\omega'}}_{\omega'}\} \\ m^2_{\omega'} &:= \min\{n \geq m^1_{\omega'} : n \notin N(\omega', 1)\} \\ N(\omega', 2) &:= \{n \geq m^2_{\omega'} : \kappa^n_{\omega'} \subseteq \kappa^{m^2_{\omega'}}_{\omega'}\} \\ m^3_{\omega'} &:= \min\{n \geq m^2_{\omega'} : n \notin N(\omega', 2)\} \\ N(\omega', 3) &:= \{n \geq m^3_{\omega'} : \kappa^n_{\omega'} \subseteq \kappa^{m^3_{\omega'}}_{\omega'}\} \\ &\ \ \vdots \end{aligned} \tag{A.49}$$

[4] An example is, in particular, the map ξ_z as defined in (A.30), where z satisfies the cyclicality condition of Corollary A.1.

Note that the sequence $(N(\omega', n), m^n_{\omega'})_n$ converges after a finite number N of rounds.

We thus define τ as follows

$$\tau^m(s) := \left(\hat{\tau}^m(s), \sum_{n \leq N} n\mathbf{1}_{N(\hat{\tau}^m(s),n)}(m)\right) \tag{A.50}$$

We define $\Omega := \{\tau^m(s) : s \in S_P, m \in \mathbb{N}\}$ and $\beta(\tau^m(s)) = s^m$. Note that the additional copies of each $\omega' \in \hat{\Omega}$ will form a cycle, where $(\omega', 1)$ is the entry point and for each $n \leq N$, the successors of (ω', n) are given by the successors of truncated sequences $(s^0, \ldots, s^{m^n_{\omega'}})$.

The key property of this extension is that every sequence reaching a state $\omega \in \Omega$ will have an information set that is contained in the information set of the sequence that first reaches ω. Hence, if we average out over all transition probabilities at state ω to obtain a Markov transition, obedience constraints of this Markov prior will be a sum over all obedience constraints of P.

We extend the automaton $(\Omega, \beta, \preceq)$ constructed above to have the following property for expositional convenience:

(iv) there is $\omega \in \Omega$ so that for every path $\boldsymbol{\omega}$,

$$\omega^1 = \omega \iff \{c : \exists\ m \text{ s.t. } \omega^m \in c\} = \{c_\infty(\boldsymbol{\omega})\}. \tag{A.51}$$

Condition (iv) states that the first transition of a path determines if the path converges at the first cycle it reaches or not. Any strategic automaton can easily be extended by duplicating all strategic states below the lowest-ranked cycle. This allows for an automaton representation of any prior that separates all paths which converge at their first cycle from all other paths. We consider a second extension, also for expositional convenience.

For any process $P \in \Delta(K \times \Omega^{\mathbb{N}})$, define the induced binary relation $\preceq_P$, where $\omega \preceq_P \omega'$ if and only if there is $m \in \mathbb{N}$ and path $\boldsymbol{\omega}$ so that $\text{marg}_{\Omega^{\mathbb{N}}}(P)(\boldsymbol{\omega}) > 0$, $\omega^m = \omega$, and $\omega^{m+1} = \omega'$. Recall that $\mathcal{C}$ denotes the collection of cycles of the automaton under $\preceq$. Let $\mathcal{C}_P$ denote the collection of cycles (i.e. equivalence classes) under $\preceq_P$. Let the induced partial order on those cycles be denoted $\preceq_{\mathcal{C}_P}$. Finally, let $\overline{\preceq}_P$ and $\overline{\preceq}_{\mathcal{C}_P}$ denote the corresponding transitive closures.

(v) Every P admits a representation so that there exist $z, M \in \mathbb{N}$ so that for all paths $\boldsymbol{\omega}$ satisfying $\text{marg}_{\Omega^{\mathbb{N}}}(P)(\boldsymbol{\omega}) > 0$,

$$\begin{aligned} \left|\{n \in \mathbb{N} : \omega^n \notin \cup_{c \in \mathcal{C}_P} c\}\right| &= M \\ |c| &= z,\ \forall\ c \in \mathcal{C}_P. \end{aligned} \tag{A.52}$$

Definition A.6 (SCAMP Automaton). *Let the SCAMP automaton* $(\Omega, \beta, \leq)$ *be the automaton constructed above from the strategic automaton obtained in Claim A.5, suitably extended to satisfy* (*iv*) *and* (*v*).

From now on we will be using only the SCAMP automaton. Let P be Markov. To each such cycle c we can associate a cycling probability $\eta_{c,k} \in [0,1]$ under P,

$$\eta_{c,k} := \prod_{\omega \in c} P(c|k, \omega). \tag{A.53}$$

For any two ordered cycles $c, c' \in \mathcal{C}_P$ of P, let $X_P(c, c')$ denote the number of states, which are not cycles under P, that are ranked between a state in c and a state in c',

$$X_P(c,c') := \{\omega : \exists\ \omega' \in c \text{ s.t. } \omega' \precsim \omega\} \cap \{\omega : \exists\ \omega' \in c' \text{ s.t. } \omega \precsim \omega'\}. \tag{A.54}$$

Let $(c_m(s))_{m \leq m(s)}$ denote the cycles as constructed above for any $s \in S_P$ where $c_{m(s)}(s)$ is the last cycle visited by s. Define the set of sequences which transition from c to $\hat{c}$,

$$\mathcal{H}_P(c,\hat{c}) := \left\{ s \in S_P : \begin{array}{c} \exists\ m \leq n \text{ s.t.} \\ \{\tau^m(s), \ldots, \tau^n(s)\} = X_P(c,\hat{c}) \end{array} \right\}. \tag{A.55}$$

Define the average exit probability of P

$$\overline{\mu}_P(\hat{c}|k,c) := \frac{P(\mathcal{H}_P(c,\hat{c})|k)}{\sum_{c<\tilde{c}} P(\mathcal{H}_P(c,\tilde{c})|k)}. \tag{A.56}$$

Note that a cycling probability $\eta = (\eta_{k,c})_{k,c} \in [0,1]^{K \times \mathcal{C}_P}$, an exit probability $\mu \colon K \times \mathcal{C} \to \Delta(\mathcal{C})$ with $\mu(c|k,c') > 0 \implies c' \leq_{\mathcal{C}_P} c$, and a marginal probability on K together induce a Markov prior $P_{\mu,\eta}$: For every $k \in K$ and $s \in S_P$,

$$P_{\mu,\eta}(s|k) := \prod_{m<m(s)} (\eta_{k,c_m(s)})^{\ell_m(s)} (1 - \eta_{k,c_m(s)}) \mu(c_{m+1}(s)|k, c_m(s)), \tag{A.57}$$

where ℓ_m is the looping number, i.e.

$$\ell_m(s) := \frac{1}{z} \left| \left\{ n \in \mathbb{N} : \forall\ l < z,\ \tau^{n+l}(s) \in c_m(s) \right\} \right|, \tag{A.58}$$

and $z \in \mathbb{N}$ is the size of a cycle.

Definition A.7 (Simple Cycling Probability)**.** *A cycling probability* $\eta = (\eta_{k,c})_{k,c} \in [0,1]^{K\times\mathcal{C}}$ *is* simple *if for all* $k \in K$ *and cycles* $c, \hat{c} \in \mathcal{C}$*,*

$$\eta_{k,c}, \eta_{k,\hat{c}} > 0 \implies \eta_{k,c} = \eta_{k,\hat{c}}.$$

Simple Cycling probabilities imply Stationary Markov priors:

Claim A.8. *For any simple cycling probability,* η *the induced Markov prior* $P_{\overline{\mu}_P,\eta}$ *is stationary.*

Proof. Consider a type s that visits cycles $c_1, \dots c_{n_m(s)}$, where $c_1 \preceq_{\mathcal{C}_P} \cdots \preceq_{\mathcal{C}_P} c_{n_m(s)-1} \preceq_{\mathcal{C}_P} c_\infty(\tau(s)) = c_{n_m(s)}$ and m is the first round where $\tau^m(s) \in c_{n_m(s)}$. By property (v) of the automaton representation, if s loops for a total of ℓ_m times before round m, then all sequences $\hat{s}$ which reach a terminal node at round m for the first time must also loop through ℓ_m cycles before round m. Define

$$\overline{S}^m := \{s \in S_{P_{\overline{\mu}_P,\eta}} : m = \min\{n : \tau^n(s) \in c_\infty \circ \tau(s)\}\}, \tag{A.59}$$

By property (iv), if $n_m(s) > 1$, then conditional on $\tau^1(s) = \omega^1$ and conditional the event $\overline{S}^m$, we can write[5] for any simple transition probability η:

$$\begin{aligned} P_{\overline{\mu},\eta}(s^0, \dots, s^m | k, \omega^1, \overline{S}^m) &= \frac{\prod_{1<h\le n_m(s)} \eta^{\ell_m}_{k,c_{h-1}} \overline{\mu}(c_h|k, c_{h-1})}{\sum_{\hat{s}\in\overline{S}^m:\tau^1(\hat{s})=\omega^1} \prod_{1<h\le n_m(\hat{s})} \eta^{\ell_m}_{k,c_{h-1}} \overline{\mu}(c_h|k, c_{h-1})} \\ &= \frac{\prod_{1<h\le n_m(s)} \overline{\mu}(c_h|k, c_{h-1})}{\sum_{\hat{s}\in\overline{S}^m:\tau^1(\hat{s})=\omega^1} \prod_{1<h\le n_m(\hat{s})} \overline{\mu}(c_h|k, c_{h-1})}. \end{aligned}$$

Let m be the first round so that for every $c \in C_{\infty,P}$, there exists $s \in S_P$ so that $\tau^m(s) \in c_\infty(\tau(s))$. From the definition of $\overline{\mu}$ (see expression (A.56)) and the fact that every path of a representation eventually reaches a cycle in $C_{\infty,P}$ we then obtain that

$$\begin{aligned} \sum_{\hat{s}\in\overline{S}^m:\tau^1(\hat{s})=\omega^1} \prod_{1<h\le n_m(\hat{s})} \overline{\mu}(c_h|k, c_{h-1}) &= \sum_{s\in S_{P_{\overline{\mu},\eta}}:\exists\ n \text{ s.t. } \tau^n(s)\in c_\infty(\tau(s))} P(s|k, \omega^1) \\ &= 1. \end{aligned}$$

Hence for every $k \in K$, every cycle $c \in C_{\infty,P}$,

$$\begin{aligned} \frac{P_{\overline{\mu},\eta}(\tau^m(s) \in c|k, \omega^1)}{\sum_{\hat{s}\in\overline{S}^m:\tau^1(\hat{s})=\omega^1} P_{\overline{\mu},\eta}(\hat{s}|k, \omega^1)} &= \prod_{1<h\le n_m(s)} \overline{\mu}(c_h|k, c_{h-1}) \\ &= \sum_{\boldsymbol{\omega}\in c_\infty^{-1}(c)} P_{\overline{\mu},\eta}(\boldsymbol{\omega}|k, \omega^1) \end{aligned}$$

[5] Note that we don't need to condition $\mu(c_h|k, c_{h-1})$ on ω^1 since every automaton state has a unique predecessor that is not in the same cycle.

It remains to show that for every $s, \hat{s} \in S_{P_{\overline{\mu},\eta}}$ and player i,

$$\mathcal{I}_{i,P}(s_i) = \mathcal{I}_{i,P}(\hat{s}_i) \implies P(k,\omega^1|s_i) = P(k,\omega^1|\hat{s}_i), \quad P\,a.s. \tag{A.60}$$

It will be enough to show that $P(k,\omega^1|s_i)$ does not depend on the cycling probability. Let $c_h(\tau(s))$ be the h-th cycle passed by $\tau(s)$. We consider two cases:

(i) Suppose $P_{\overline{\mu}_P,\eta}(\{(k,\hat{s}) : k \in K, c_1(\tau(\hat{s})) = c_\infty(\tau(\hat{s}))\}|s_i) = 0$. Then by properties (iv) and (v) and the fact that the cycling probability is simple, the argument above implies that i's beliefs on (k,ω^1) conditional on s_i do not depend on the cycling probability. Hence $P(k,\omega^1|s_i) = P(k,\omega^1|\hat{s}_i)$.

(ii) Consider now the case where $P_{\overline{\mu}_P,\eta}(\{(k,\hat{s}) : k \in K, c_1(\tau(\hat{s})) = c_\infty(\tau(\hat{s}))\}|s_i) > 0$. Then it must be that for every h, $\beta_i(c_h(\tau(\hat{s}))) = s_i^m$ for all $m \geq \overline{m}(s_i) := \min\{n \in \mathbb{N} : \exists \tilde{s} \text{ s.t. } P(\tilde{s}|s_i) > 0, \tau^n(\tilde{s}) \in c_1(\tau(\tilde{s})) = c_\infty(\tau(\tilde{s}))\}$. But $\overline{m}(s_i) = \overline{m}(\hat{s}_i)$ and we must have that $s_i = \hat{s}_i$. So again, $P(k,\omega^1|s_i) = P(k,\omega^1|\hat{s}_i)$.

□

A.3.3 Sufficiency of SCAMP

Claim A.9. *Let $P \in \mathcal{P}$ and let $\overline{\mu}_P$ be given by (A.56). Then there exists a simple cycling probability η so that the induced Markov prior $P_{\overline{\mu},\eta}$ is an outcome equivalent SCAMP.*

Proof. We will first show that we can find a simple cycling probability so that sub-obedience holds everywhere. We then argue that conditional obedience also holds. Claim A.8 implies $\overline{\mu}_P$ induces a Stationary Markov process. But then by Claim A.7 we deduce that $\overline{P}$ is SCAMP.

We start by showing sub-obedience. For any ω and k, define

$$\begin{aligned}
\underline{G}_i^{a_i,a_i'}(k,\omega) &:= \min_{\sigma_k \in \beta_{-i}(\omega)} (u_i(k,a_i,\sigma_k) - u_i(k,a_i',\sigma_k)), \\
\overline{G}_i^{a_i,a_i'}(k,\omega) &:= \max_{\sigma_k \in \beta_{-i}(\omega)} (u_i(k,a_i,\sigma_k) - u_i(k,a_i',\sigma_k)).
\end{aligned} \tag{A.61}$$

For $\omega \precsim \omega'$ we have that $\beta_j(\omega') \subseteq \beta_j(\omega)$ for any player j. Hence for every player i, $k \in K$ and $a_i, a_i' \in A_i$,

$$\begin{aligned}
\underline{G}_i^{a_i,a_i'}(k,\omega) &\leq \underline{G}_i^{a_i,a_i'}(k,\omega'), \\
\overline{G}_i^{a_i,a_i'}(k,\omega) &\geq \overline{G}_i^{a_i,a_i'}(k,\omega').
\end{aligned} \tag{A.62}$$

Let $\overline{P} \coloneqq P_{\overline{\mu}_P,\eta}$ denote the Markov prior with a simple cycling probability $\eta = (\eta_{k,c})_{k,c}$. We now show that $\overline{P}$ is canonical for any simple cycling probability that preserves the support of limit cycles, $\{(k,c) : \overline{P} \circ (\mathrm{id} \times c_\infty)^{-1}(k,c) > 0\} = \{(k,c) : P \circ (\mathrm{id} \times c_\infty)^{-1}(k,c) > 0\}$. Let $s \in S_{\overline{P}}$ and fix any round $m \in \mathbb{N}$. We can write obedience constraints at information set $I^m_{i,\overline{P}}(s_i)$ for $a_i \in s_i^{m+1}$, $a'_i \in s_i^m \setminus s_i^{m+1}$

$$0 < \sum_{(k,\omega) \in I^m_{i,\overline{P}}(s_i)} \eta^{\ell(\omega)} \prod_{1 < h \leq n(\omega)} \overline{\mu}_P(c_{\omega,h}, |k, c_{\omega,h-1}) P(\omega^1|k) P_K(k) \underline{G}_i^{a_i,a'_i}(k,\omega), \tag{A.63}$$

moreover, for every $a_i \in s_i^m$ and every $a'_i \in A_i$,

$$0 \leq \sum_{(k,\omega) \in I^m_{i,\overline{P}}(s_i)} \eta^{\ell(\omega)} \prod_{1 < h \leq n(\omega)} \overline{\mu}_P(c_{\omega,h}, |k, c_{\omega,h-1}) P(\omega^1|k) P_K(k) \overline{G}_i^{a_i,a'_i}(k,\omega), \tag{A.64}$$

where a path reaching ω must pass through cycles $c_{\omega,1}, \dots, c_{\omega,n(\omega)}$ and the looping number $\ell(\omega)$ is defined as the number of loops required to reach ω at round m. Note that if $\omega \overline{\precsim} \hat{\omega}$,

$$\ell(\omega) \geq \ell(\hat{\omega}). \tag{A.65}$$

Indeed, if a state is ranked lower at the round m, it must be that it cycled weakly more often.

For any information set under P, $I_i \in \mathcal{I}_{i,P}(s_i)$ and cycle c, define

$$\begin{aligned} \mathcal{H}_P(c|I_i) &\coloneqq \left\{ s \in \cup_{\hat{c}} \mathcal{H}_P(c,\hat{c}) : \exists\ n \in \mathbb{N} \text{ s.t. } I^n_{i,P}(s_i) = I_i \right\} \\ \mathcal{H}_P(I_i) &\coloneqq \left\{ s \in S_P : \exists\ n \in \mathbb{N} \text{ s.t. } I^n_{i,P}(s_i) = I_i \right\} \end{aligned} \tag{A.66}$$

Let $\tau^m(s) = \omega$. Then by construction of the SCAMP automaton (see Definition A.6), we have that for all $\hat{s} \in S_P$ so that $\tau^n(\hat{s}) = \omega$ for some $n \in \mathbb{N}$,

$$I^n_{i,P}(\hat{s}_i) \subseteq I^m_{i,\overline{P}}(s_i). \tag{A.67}$$

Summing the right hand side of the sub-obedience constraints (A.63) of player i under P at information set I_i over all sequences in $\mathcal{H}_P(I_i)$,

$$\sum_{(k,\omega) \in I_i} \rho(\omega, I_i, k) \alpha^{a_i,a'_i}_{k,\omega}, \tag{A.68}$$

where $\alpha_{k,\omega}^{a_i,a_i'} := P_K(k)\underline{G}_i^{a_i,a_i'}(k,\omega)$ and

$$\rho(\omega, I_i, k) := \prod_{1<h\leq n(\omega)} P(\mathcal{H}_P(c_{\omega,h}|I_i) \mid k, \mathcal{H}_P(c_{\omega,h-1}|I_i)). \tag{A.69}$$

Substituting the definition of $\overline{\mu}_P$ into the right hand side of (A.63),

$$0 < \sum_{I_i\subseteq I_{i,\overline{P}}^m(s_i)} \sum_{(k,\omega)\in I_i} \eta^{\ell(\omega)}\rho(\omega, I_i, k)P(\mathcal{H}_P(I_i)|k)\alpha_{k,\omega}^{a_i,a_i'}. \tag{A.70}$$

Hence the inequality (A.63) holds for any interior choice of η. The same argument implies that inequality (A.64) holds for $\eta = 1$. We now show that this means conditional obedience holds. Conditional obedience takes the form: For $a_i \in s_i^{m+1}$, $a_i' \in s_i^m \setminus s_i^{m+1}$

$$0 < \sum_{(k,\omega)\in\overline{I_{i,\overline{P}}^m(s_i)}} \prod_{1<h\leq n(\omega)} \overline{\mu}_P(c_{\omega,h},|k, c_{\omega,h-1})P(\omega^1|k)P_K(k)\underline{G}_i^{a_i,a_i'}(k,\omega), \tag{A.71}$$

moreover, for every $a_i \in s_i^m$ and every $a_i' \in A_i$,

$$0 \leq \sum_{(k,\omega)\in\overline{I_{i,\overline{P}}^m(s_i)}} \prod_{1<h\leq n(\omega)} \overline{\mu}_P(c_{\omega,h},|k, c_{\omega,h-1})P(\omega^1|k)P_K(k)\overline{G}_i^{a_i,a_i'}(k,\omega). \tag{A.72}$$

Indeed, by conditioning on reaching maximal states, implies that all sequences must cycle for the same number of times on the SCAMP automaton (by properties (iv) and (v) of the SCAMP Automaton). We proceed by backwards induction:

(i) First suppose that $s_i^m = s_i^{m+l}$ for all $l > 0$ (i.e. s_i converges before or at round m). Then there is a terminal information set I_i^* so that the beliefs of s_i are weakly stronger and $\eta = 1$ because the cycle is terminal. But then conditional obedience also holds since $\overline{I_{i,\overline{P}}^m(s_i)}$ contains weakly lower ranked states than I_i^* but weakly higher ranked than $I_{i,\overline{P}}^m(s_i)$. Since we established sub-obedience on $I_{i,\overline{P}}^m(s_i)$, we have (A.71) and since we have obedience on I_i^*, we must also have (A.72).

(ii) Given our argument in (i) we now proceed by backwards induction: Suppose that for every $(k,\omega) \in \overline{I_{i,\overline{P}}^m(s_i)}$, every successor ω' satisfying both $\omega \precsim \omega'$ and $\omega' \notin \overline{I_{i,\overline{P}}^m(s_i)}$, we have shown obedience (i.e. conditions

(A.63) and (A.64) at every information set I_i' that contains ω'. We conclude that conditional obedience holds: $\overline{I^m_{i,\overline{P}}(s_i)}$ contains weakly lower ranked states than I_i' but weakly higher ranked than $I^m_{i,\overline{P}}(s_i)$. Since we established sub-obedience on $I^m_{i,\overline{P}}(s_i)$, we have (A.71) and since we have obedience on I_i' per inductive hypothesis, we must also have (A.72).

By Claim A.8, for any choice simple cycling probability we have that $\overline{P}$ is a Stationary Markov process. We have shown that it satisfies sub-obedience and conditional obedience. But then by Claim A.7 we deduce that $\overline{P}$ is SCAMP. □

Theorem 5.1 *For every finite game there exists a strategic automaton so that the set of SCAMP induce all outcome distributions.*

Proof. An immediate consequence of Claim A.9 and the definition of the SCAMP automaton. □

A.4 Linear Characterization of Rationalizable Outcomes

Define the *Generating Set* X as the set of paths on the automaton of length $m^* = |\Omega|$,

$$X := \left\{(k, \omega^0, \dots, \omega^{m^*}) : \omega^0 \leq \dots \leq \omega^{m^*}\right\}. \tag{A.73}$$

For any $p \in \Delta(X)$, let its support be X_p. For every player i, define the equivalence classes $\mathcal{X}_{i,p}$,

$$\mathcal{X}_{i,p} := \left\{\left\{(k,e) \in X_p : \beta_i(e^0) = \beta_i(\hat{e}^0), \dots, \beta_i(e^{m^*}) = \beta_i(\hat{e}^{m^*})\right\} : (\hat{k}, \hat{e}) \in X_p\right\}. \tag{A.74}$$

A distribution $p \in \Delta(X)$ satisfies *obedience* if for every $(k,e) \in X_p$, every i, every m, $X_{i,p} \in \mathcal{X}_{i,p}$ so that $(k,e) \in X_{i,p}$ and $m \leq m_p^*$,

$$\mathrm{br}_i(p|_{X_{i,p}} \circ (\mathrm{id} \times \beta_{-i} \circ \mathrm{proj}_{m-1})^{-1}) = \beta_i(e^m), \tag{A.75}$$

where $p|_{X_{i,p}}(k,e') := p(k,e')/p(X_{i,p})$, for all $k \in K$ and $e' \in X_{i,p}$. The collection of distributions $p \in \Delta(X_p)$ satisfying expression (A.75) for every player i: For every i, $X_{i,P} \in \mathcal{X}_{i,p}$, $e \in X_{i,p}$, $m \leq m^*$ and every $a_i' \in \beta_i(e^{m-1}) \setminus \beta_i(e^m)$ there is

$a_i \in \beta_i(e^m)$ so that

$$\begin{aligned}0 &< \sum_{(k,e)\in X_{i,p}} p(k, e^{m-1}) \min_{\sigma_k \in \Delta(\beta_{-i}(e^{m-1}))} \sum_{a_{-i}} \sigma_k(a_{-i})(u_i(k, a_i, a_{-i}) - u_i(k, a_i', a_{-i})) \\ &= \sum_{(k,e)\in X_{i,p}} p(k, e^{m-1}) \min_{\sigma_k \in \beta_{-i}(e^{m-1})} (u_i(k, a_i, \sigma_k) - u_i(k, a_i', \sigma_k)),\end{aligned}$$

moreover, for every $a_i \in \beta_i(e^m)$ and every $a_i' \in A_i$

$$\begin{aligned}0 &\leq \sum_{(k,e)\in X_{i,p}} p(k, e^{m-1}) \max_{\sigma_k \in \Delta(\beta_{-i}(e^{m-1}))} \sum_{a_{-i}} \sigma_k(a_{-i})(u_i(k, a_i, a_{-i}) - u_i(k, a_i', a_{-i})) \\ &= \sum_{(k,e)\in X_{i,p}} p(k, e^{m-1}) \max_{\sigma_k \in \beta_{-i}(e^{m-1})} (u_i(k, a_i, \sigma_k) - u_i(k, a_i', \sigma_k))\end{aligned}$$

Let $\mathcal{O}^{m^*} \subseteq \Delta(K \times \Omega^{m^*})$ denote the set of distributions on $K \times \Omega^{m^*}$ that satisfy obedience constraints. Let the set of terminal states of $p \in \Delta(K \times \Omega^{m^*})$ be given by

$$\overline{X}(p) = \{(k, \omega^{m^*}) \in K \times \overline{\Omega} : p(k, (\omega^0, \dots, \omega^{m^*})) > 0\}, \tag{A.76}$$

Letting $p^{m^*}(k, \omega) = \sum_{x \in \Omega^{m^*} : x^{m^*} = \omega} p(k, x)$, define the limit probability of p

$$\overline{p}(k, b) = \frac{\sum_{\omega:\beta(\omega)=b} p^{m^*}(k, \omega)}{p^{m^*}(\overline{X}(p))}. \tag{A.77}$$

The limit probability $\overline{p}$ satisfies limit-obedience if for every b in its support and every player i, $b_i = \mathrm{br}_i(\overline{p}(\cdot, \cdot | b_i))$. Let $\mathcal{O}^\infty \subseteq \Delta(K \times B)$ denote the set of probabilities satisfying limit-obedience.

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

$$\mathcal{O} = \{\overline{p} : p \in \mathcal{O}^{m^*}\} \cap \mathcal{O}^\infty. \tag{A.78}$$

The relative closure of this set is a convex polyhedron:

Lemma 5.2. *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 $\overline{p}^*$,

$$\overline{p}^*(k, b) = \sum_{\omega:\beta(\omega)=b} p^*(k, \omega)$$

Then we have that the relative closure of $\{\overline{p}^* : p \in \mathcal{O}^{m^*}\}$ is a convex polyhedron. Finally, note that the collection $\{p^{m^*}(\overline{X}(p)) : p \in \mathcal{O}^{m^*}\} \subseteq [0,1]$ is an interval $[\underline{x}, \overline{x}]$. Let $\overline{X}$ denote the set of states of nature K and sequences in Ω^{m^*} that reach a terminal state. The we have that

$$\{\overline{p} : p \in \mathcal{O}^{m^*}\} = \left\{\frac{1}{x}\overline{p}^* : p \in \mathcal{O}^{m^*}, x \in [\underline{x}, \overline{x}]\right\} \cap \Delta(\overline{X}). \tag{A.79}$$

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

We now show that $\mathcal{O}$ coincides with the set of all outcome distributions. For any $(k, e) = (k, (\omega^0, \ldots, \omega^{m^*})) \in X$, let the ordered list of cycles visited by e be denoted by

$$c_1(e), \ldots, c_{m(e)}(e). \tag{A.80}$$

For every distribution $p \in \Delta(X)$, we can obtain an average exit probability of p

$$\overline{\mu}_p(c, \hat{c}|k) := p(\mathcal{H}_p^{m^*}(c, \hat{c})|k), \tag{A.81}$$

where

$$\mathcal{H}_p^{m^*}(c, \hat{c}) := \left\{(s^0, \ldots, s^{m^*}) : \begin{array}{c} \exists\ m \leq n \leq m^* \text{ s.t.} \\ \{\tau^m(s), \ldots, \tau^n(s)\} = X_P(c, \hat{c}) \end{array}\right\}. \tag{A.82}$$

For any choice of simple cycling probability η, the triple $(\overline{\mu}_p, \eta, \mathrm{marg}_K(p))$ induces a distribution on X, denoted p_P, where for every $k \in K$, $(k, e) = (k, (\omega^0, \ldots, \omega^{m_P^*})) \in X_P$

$$p_{\overline{\mu}_p, \eta}(e|k) := \prod_{m < m(e)} (\eta_{k, c_m(e)})^{\ell_m(e)} (1 - \eta_{k, c_m(e)}) \mu(c_{m+1}(e)|k, c_m(e)), \tag{A.83}$$

where (with slight abuse of notation) the looping number $\ell_m(e)$ is defined analogously to (A.58)

$$\ell_m(e) := \frac{1}{z} \left| \left\{ n \leq m^* : \forall\ l < z,\ e^{n+l} \in c_m(e) \right\} \right|. \tag{A.84}$$

Corollary A.2. *If $p \in \Delta(X)$ satisfied obedience then there is simple cycling probability η so that the Markov prior induced by $(\overline{\mu}_p, \eta, marg_K(p))$ is SCAMP whose outcome distribution is equal to $\overline{p}$.*

Proof. This follow readily from the construction and the arguments in Claim A.9 applied to □

We thus obtain our characterization of strongest rationalizable outcomes:
Theorem 5.2

(i) Every SCAMP P induces a distribution $p_P \in \mathcal{O}$ through its marginal on $K \times \Omega^{m^*}$ so that the limit probability $\overline{p}_P$ coincides with the outcome distribution of P.

(ii) For every $p \in \mathcal{O}$ there exists SCAMP P_p so that the limit probability $\overline{p}$ coincides with the outcome distribution of P_p.

Proof. (i) Follows from the definition of SCAMP. (ii) Is an immediate consequence of Claim A.2. □