

MONOTONICITY AND ROBUST IMPLEMENTATION UNDER FORWARD-INDUCTION REASONING

PIERPAOLO BATTIGALLI

Department of Decision Sciences, Bocconi University and IGIER

EMILIANO CATONINI

New York University Shanghai

In sequential games, the set of paths consistent with rationality and forward-induction reasoning may change nonmonotonically when adding transparent restrictions on players' beliefs. Yet, we prove that—in an incomplete-information environment—predictions become sharper when the restrictions only concern initial beliefs about types. Thus, strong rationalizability for games with payoff uncertainty characterizes the path predictions of forward-induction reasoning across all possible restrictions on players' hierarchies of exogenous beliefs. With this, we can solve an open problem: the implementation of social choice functions through sequential mechanisms under forward-induction reasoning—which considerably expands the realm of implementable functions compared with simultaneous mechanisms (Müller (2016))—is indeed robust in the sense of Bergemann and Morris (2009).

KEYWORDS: Incomplete information, forward induction, strong rationalizability, path monotonicity, robust implementation, sequential mechanisms.

1. INTRODUCTION

WE STUDY THE IMPLICATIONS of belief restrictions for a version of rationalizability in incomplete-information sequential games, which captures forward-induction reasoning. Belief restrictions may be motivated by contextual details of an economic application, such as an objective distribution of the players' payoff types. We prove a monotonicity result for restrictions that concern only exogenous (initial) beliefs, and we apply the result to solve an open problem on robust implementation of social choice functions through sequential mechanisms.¹

The building blocks of our analysis—strong rationalizability and belief restrictions—relate to the notion of **strong belief**. Strong belief in an event E means that E is assigned probability 1 conditional on any observation that does not contradict E . Strong rationalizability is characterized, in epistemic terms, by the assumptions of *rationality and common strong belief in rationality* (Battigalli and Siniscalchi (2002)). In particular, each player is initially certain of the rationality of the coplayers, that is, that they are subjective expected

Pierpaolo Battigalli: pierpaolo.battigalli@unibocconi.it

Emiliano Catonini: emiliano.catonini@nyu.edu

We thank Carlo Andreatta, Alessandro Cherubin, Nicodemo De Vito, Samuele Dotta, Drew Fudenberg, Shuige Liu, Silvia Meneghesso, Mariann Ollár, Andrés Perea, Viola Sigismondi, Marciano Siniscalchi, Nicolas Sourisseau, Joel Watson, Gabriel Ziegler, and three anonymous referees for useful comments. This project is funded by the European Union (ERC, TRAITS-GAMES, 101142844). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Emiliano Catonini gratefully acknowledges the financial support of the National Science Foundation of China, Excellent Young Scientist Program (overseas).

¹The application builds on the work of Bergemann and Morris (2009) and Müller (2016). For an overview of robust implementation, see Bergemann and Morris (2012) and the references therein.

payoff maximizers. Furthermore, if an observed but unexpected move is consistent with the assumption that the copleayers are rational, then such a move is interpreted as intentional and belief in rationality is maintained, leading to inferences about copleayers' private information and future moves. This rationalization of past moves is called "**forward induction**."²

We study belief restrictions that are **transparent** in that there is common belief throughout the game that the restrictions hold. The prior literature has shown that the strategy profiles and paths consistent with strong rationalizability and transparent belief restrictions can change nonmonotonically with respect to the restrictions. Notably, stronger restrictions do not necessarily imply sharper predictions (see, e.g., Battigalli and Friedenberg (2012)).

We prove, however, that stronger restrictions *do imply* sharper path predictions when the restrictions concern only exogenous beliefs, that is, initial beliefs about types. This monotonicity theorem implies that strong rationalizability in the "belief-free" setting without restrictions on beliefs characterizes the path predictions of forward-induction reasoning across all possible restrictions on exogenous beliefs. We apply this result to the theory of robust full implementation.

The power of sequential mechanisms to implement social choice functions (SCFs) was first explored for *complete-information* environments, that is, under the assumption that agents' preferences, or payoff types—even if they are unknown to the planner—are common knowledge among the agents.³ In particular, drawing on work by Abreu and Matsushima (1992), Glazer and Perry (1996) proved that a large class of SCFs can be virtually implemented⁴ by means of perfect-information sequential mechanisms, if players reason by backward induction and play the unique subgame perfect equilibrium. If the domain of preference profiles is finite, this is equivalent to strongly rationalizable virtual implementation, because backward- and forward-induction reasoning yield the same path of play in generic finite games with complete and perfect information (Battigalli (1997), Battigalli and Siniscalchi (2002)).

In environments with *incomplete information*, agents' behavior depends on their beliefs about each other's types. Just like in complete-information environments the planner is not assumed to know agents' commonly known payoff types, in environments with incomplete information the planner may not know what hierarchies of exogenous beliefs the agents can hold and conceive. Formally, the planner may be uncertain about the relevant exogenous type structure, for example, whether agents' beliefs are derived from a common prior on the domain of preference profiles and, if so, what that prior is (see Harsanyi (1967–1968) and Mertens and Zamir (1985)). In compliance with the Wilson's (1987) doctrine, Bergemann and Morris (2009) analyze **robust implementation**, that is, the possibility to implement an SCF independently of the exogenous type structure. They show that robust (virtual) full implementation of SCFs by means of *static* (i.e., simultaneous-move) mechanisms—which amounts to rationalizable implementation—is severely limited when agents' valuations of outcomes exhibit a mild degree of interdependence. Müller (2016)

²In contrast with backward induction, which interprets past unexpected moves as trembles and assumes strategic rationality only for future moves. Forward-induction ideas are also used in the equilibrium refinements literature; see, for example, the survey by Kohlberg (1990). Battigalli and Siniscalchi (2002, 2003) and Battigalli and Catonini (2024) explore relationships between the epistemic game theory approach underlying rationalizability ideas and forward-induction refinements of equilibrium.

³See Moore and Repullo (1988), Chapter 10 of Osborne and Rubinstein (1994) and the references therein.

⁴Virtual implementation of an SCF means that, for each type profile, the outcome predicted by the solution concept can be made arbitrarily close to the outcome prescribed by the SCF. See the cited references.

instead proves that using *sequential* mechanisms and assuming that agents reason by *forward induction*—as captured by strong rationalizability—yields a significant expansion of the implementable SCFs. Yet, it was not known whether implementation in strongly rationalizable strategies is robust to considering contextual restrictions on agents' exogenous interactive beliefs about each other's types.

Our game-theoretic result allows us to solve this open problem. Strongly rationalizable strategies may change nonmonotonically when adding restrictions on exogenous beliefs, but only the induced paths of play matter for the implementation of SCFs. Since we prove that the set of possible paths under stronger restrictions on exogenous beliefs is included in the one obtained with weaker or no restrictions, it follows that *strongly rationalizable implementation is robust* in the aforementioned sense.

The rest of the paper is organized as follows. Section 2 provides a heuristic analysis and additional background. Section 3 contains the game-theoretic preliminaries. Section 4 states and explains the main theorem. Section 5 applies this theorem to the analysis of Bayesian games. Section 6 applies our game-theoretic results to the robust implementation problem. Section 7 discusses extensions. The Appendix collects proofs of claims and lemmas that are omitted from the main body of the paper.

2. HEURISTIC ANALYSIS AND DETAILED BACKGROUND

In this section, we first illustrate intuitions and difficulties behind our main monotonicity result by means of an example (2.1). Next, we set the stage for the robust-implementation application of our result: first, we introduce robust implementation via static mechanisms and its connection with rationalizability (2.2); second, we discuss implementation via sequential mechanisms and illustrate by example the use of forward-induction reasoning to expand the set of implementable SCFs (2.3).

2.1. Strong Rationalizability, Heuristic Analysis of an Example

Strong rationalizability is the iterated elimination, for each payoff type of each player, of the strategies that are not sequential best replies to belief systems, which assign probability 1, as long as possible, to the coplayers' strategies that survive the previous elimination steps.⁵ Considering transparent belief restrictions, we obtain an “umbrella solution concept” called strong **directed** rationalizability, which works in the same way, except that the set of possible belief systems for each type is also restricted exogenously and not just through strategic reasoning. Thus, the elimination procedure is parameterized by a profile $\Delta = (\Delta_i)$ of players' restricted sets of beliefs. For any *fixed* Δ , we obtain a specific solution called “**strong Δ -rationalizability**.” The kind of belief restrictions we analyze in this paper only pertain to the initial beliefs about the payoff types of the coplayers.

In the following example, we illustrate the two elimination procedures and the main hurdle towards proving our general monotonicity result. Since we have not yet introduced all the required formal concepts, the analysis is necessarily heuristic and based on intuition.

⁵In complete-information environments, strong rationalizability used to be called “extensive-form rationalizability” (Pearce (1984), Battigalli (1997)). There exist several versions of strong rationalizability, but Battigalli, Catonini, and Manili (2023) prove that they are equivalent.

EXAMPLE 1: Consider a signaling game between players 1 (sender) and 2 (receiver) where the set of possible payoff types θ_1 of the sender is $\Theta_1 = \{x, y, z\}$, the set of messages/signals is $M = S_1 = \{\ell, r\}$, and the sets of feasible reactions of the receiver are $\mathcal{A}_2(\ell) = \{a, b\}$ after message ℓ and $\mathcal{A}_2(r) = \{c, d, e\}$ after message r . Thus, the set of receiver’s strategies is $S_2 = \mathcal{A}_2(\ell) \times \mathcal{A}_2(r)$, whereas the sender’s strategies and signals coincide. The payoffs are as follows:

Payoffs of 1 and 2:

| after ℓ | a | b |
|----------------|-----|-----|
| $\theta_1 = x$ | 3 1 | 1 0 |
| $\theta_1 = y$ | 1 0 | 1 1 |
| $\theta_1 = z$ | 3 1 | 1 0 |

| after r | c | d | e |
|----------------|-----|-----|-----|
| $\theta_1 = x$ | 0 0 | 0 0 | 0 1 |
| $\theta_1 = y$ | 0 0 | 0 1 | 3 0 |
| $\theta_1 = z$ | 0 1 | 2 0 | 2 0 |

We start with strong rationalizability (i.e., no restrictions on exogenous beliefs).

1. The first step of elimination follows from mere rationality. We can only eliminate message r for type x , as it is dominated by message ℓ . Thus, we write $S_1^1(x) = \{\ell\}$ for the set of messages/signals consistent with rationality for type x .

[Since no strategy of player 2 is eliminated in the first step, it follows that in even (odd) steps only eliminations for player 2 (player 1) are possible.]

2. The optimal behavior of the receiver depends on his belief system

$$\mu_2 = (\mu_2(\cdot|\emptyset), \mu_2(\cdot|\ell), \mu_2(\cdot|r)),$$

where $\mu_2(\cdot|\emptyset)$ is the initial belief about the sender’s type-message pair, and for each $m \in \{\ell, r\}$, if $\mu_2(\Theta \times \{m\}|\emptyset) > 0$, then belief $\mu_2(\cdot|m)$ after observing message m is derived from the initial belief by conditioning (in other words, μ_2 satisfies the chain rule). At the second step of the elimination procedure, the initial belief is assumed to assign probability 1 to the type-message pairs that survived the first step:

$$\mu_2\left(\bigcup_{\theta_1 \in \Theta_1} \{\theta_1\} \times S_1^1(\theta_1)|\emptyset\right) = 1.$$

The same applies to each belief $\mu_2(\cdot|m)$ provided that $m \in S_1^1(\theta_1)$ for some $\theta_1 \in \Theta_1$. Thus, here we have $\mu_2((x, r)|r) = 0$. In words, by strong belief in the sender’s rationality, after observing message r the receiver concludes that the sender is not of type x —this is an instance of forward-induction reasoning. Given this, action e is never a best reply. Hence,

$$S_2^2 = \{a.c, b.c, a.d, b.d\},$$

where, for example, $a.c$ denotes the strategy choosing a after ℓ and c after r .

3. For type y , action r is not a best reply to any belief over S_2^2 . Thus, $S_1^3(y) = \{\ell\}$.

4. Every belief system of the receiver must now assign probability 1 to type z after message r . Thus, $S_2^4 = \{a.c, b.c\}$.

5. Type z now expects to obtain 0 from r and at least 1 from ℓ . Thus, $S_1^5(z) = \{\ell\}$.

No remaining strategy of the receiver can be eliminated. So, we obtain

$$\forall \theta_1 \in \Theta_1, \quad S_1^\infty(\theta_1) = \{\ell\}, \quad S_2^\infty = \{a.c, b.c\}.$$

It follows that the strongly rationalizable paths are (ℓ, a) and (ℓ, b) for every state (sender’s type).

Now consider the following restrictions on the exogenous beliefs of the receiver.⁶ Let Δ_2 collect the belief systems μ_2 that initially assign probability 1 to type z , that is,

$$\mu_2(\{z\} \times S_1|\emptyset) = 1.$$

Strong Δ -Rationalizability is given by the following steps:

$\Delta, 1.$ As above, message r is eliminated for type x , so we write $S_1^{\Delta,1}(x) = S_1^1(x) = \{\ell\}$. But now, some strategies of the receiver are also eliminated. By the chain rule, every belief system $\mu_2 \in \Delta_2$ assigns probability 1 to z given ℓ , if $\mu_2(\{(z, \ell)\}|\emptyset) > 0$, and/or given r , if $\mu_2(\{(z, r)\}|\emptyset) > 0$. Thus, the receiver best replies with a after ℓ and/or with c after r : $S_2^{\Delta,1} = \{a.c, b.c, a.d, a.e\}$.

$\Delta, 2.$ As in strong rationalizability, action e is never a best reply given r ; hence, strategy $a.e$ of the receiver is eliminated: $S_2^{\Delta,2} = \{a.c, b.c, a.d\}$. Moreover, for type z , r is dominated by ℓ with respect to strategies in $S_2^{\Delta,1}$; so, $S_1^{\Delta,2}(z) = \{\ell\}$.

$\Delta, 3.$ For type y , r is dominated by ℓ over $S_2^{\Delta,2}$; thus, $S_1^{\Delta,3}(y) = \{\ell\}$. Moreover, every belief system of the receiver must now assign probability 1 to type y given message r ; thus, $S_2^{\Delta,3} = \{a.d\}$.

We pinned down one strategy for each type of each player:

$$\forall \theta_1 \in \Theta_1, \quad S_1^{\Delta,\infty}(\theta_1) = \{\ell\}, \quad S_2^{\Delta,\infty} = \{a.d\}.$$

The strongly Δ -rationalizable path is (ℓ, a) for every payoff state. Consistently with our main result (Theorem 1), it is one of the two strongly rationalizable paths. Note, however, that the strongly Δ -rationalizable reaction of the receiver to r is d , whereas the strongly rationalizable one was c . This is related to a different order of elimination of type-message pairs of the sender: With the belief restriction, (z, r) is eliminated before (y, r) , inducing the receiver to be certain of y and choose d after r ; without restriction, we have the opposite order of eliminations and d after r is ruled out.

In the example, as we will prove in full generality, strong directed rationalizability with restrictions on initial beliefs about types refines strong rationalizability in terms of paths, for each possible payoff state (profile of types). However, the strongly Δ -rationalizable strategy of the receiver is *not* strongly rationalizable, because the implications about off-path behavior change nonmonotonically after introducing the belief restrictions.

The reason is that *strong belief is not monotone* in the following sense: strong belief in a proposition (e.g., the sender is rational *and* her beliefs satisfy a given restriction) does *not* imply strong belief in a logically weaker proposition (e.g., the sender is just rational). Thus, strong directed rationalizability does not refine strong rationalizability in terms of strategies. For this reason, our path monotonicity result cannot be proven with a straightforward induction argument.

2.2. Robust Implementation, Static Mechanisms

To set the stage, we explain the conceptual connection between robust implementation and rationalizability, focusing first on static mechanisms. Consider an **economic environment** \mathcal{E} with asymmetric information. There are a set I of agents and a set Y of economic outcomes (possibly, lotteries). The (expected) value to player i of outcome y is $v_i(\theta, y)$,

⁶Since there is only one type of receiver, we do not consider restrictions on the sender's beliefs.

where $\theta = (\theta_i)_{i \in I} \in \Theta = \times_{i \in I} \Theta_i$ is a payoff state (also called “state of nature”) and θ_i is i ’s private information about θ , or i ’s “**payoff type**.”

Agents hold hierarchical beliefs about each other’s payoff types, which can be represented by means of a **type structure** \mathcal{T} à la Harsanyi (1967–1968). In words, \mathcal{T} captures what belief hierarchies are commonly believed possible, given some exogenous contextual restrictions on beliefs. Without contextual restrictions, \mathcal{T} is the universal type structure containing all the collectively coherent belief hierarchies (e.g., Mertens and Zamir (1985)).

A planner (she) wants to implement a social choice function (SCF) f , associating each payoff state θ with a desirable outcome $y = f(\theta) \in Y$.⁷ With this purpose, she commits to make the agents interact according to a mechanism \mathcal{M} , that is, some commonly known set of rules that yield a set Z of possible paths of play and an outcome function $g : Z \rightarrow Y$. In static mechanisms, $Z = A$ is just the set of possible action/message profiles; in the subclass of direct mechanisms, Z is isomorphic to Θ , because messages report agents’ types. Triple $\Gamma^b = (\mathcal{M}, \mathcal{E}, \mathcal{T})$ describes a situation of strategic interaction called “**Bayesian game**.” In the traditional full implementation problem, it is assumed that the planner knows both \mathcal{E} and \mathcal{T} . The adopted solution concept (e.g., Bayesian equilibrium, or rationalizability) yields, for each state $\theta \in \Theta$, a set $Z^{\Gamma^b}(\theta)$ of possible paths of play. A mechanism \mathcal{M} **fully implements** SCF f if, for every θ , $f(\theta)$ is the only possible outcome, that is, $g(Z^{\Gamma^b}(\theta)) = \{f(\theta)\}$.⁸ However, the planner often ignores the contextual features represented by type structure \mathcal{T} . If she deems all type structures possible, in compliance with Wilson’s doctrine, a natural notion of **robust full implementation** requires that the same mechanism \mathcal{M} fully implements SCF f for all Bayesian games Γ^b based on $(\mathcal{M}, \mathcal{E})$, that is, across all type structures \mathcal{T} (see Wilson (1987) and Bergemann and Morris (2009, 2012)). Since this paper is only concerned with different forms of full implementation, from now on we will omit the adjective “full.”

Robust implementation is conceptually related to **rationalizability**, that is, the solution concept characterizing the behavioral implications of *Rationality and Common Belief in Rationality* (RCBR).⁹ On the one hand, not relying on the assumption that players’ endogenous beliefs about each other’s behavior serendipitously coordinate on a Bayesian equilibrium is in itself a form of robustness in the spirit of Wilson’s doctrine. On the other hand, it has been observed that *the state-dependent outcomes consistent with Bayesian equilibrium across all type structures are precisely those allowed by a version of rationalizability for games with payoff uncertainty*—aka “belief-free rationalizability”—that applies to structure $(\mathcal{M}, \mathcal{E})$, that is, to a description of the game that does not specify interactive beliefs about payoff types.¹⁰ In particular, *restricting attention to static* (e.g., direct) *mechanisms*, robust Bayesian-equilibrium implementation is equivalent to implementation with respect to rationalizability for games with payoff uncertainty. Analogously, *robust* implementation with respect to rationalizability for Bayesian games is equivalent to

⁷We limit our attention to social choice *functions*. Similar considerations apply to social choice correspondences associating a set of outcomes $F(\theta) \subseteq Y$ to each payoff state θ .

⁸Partial implementation relies on equilibrium analysis and requires instead that $g(\mathbf{z}(\cdot)) = f(\cdot)$ for at least one equilibrium selection $\mathbf{z}(\cdot)$ from equilibrium correspondence $Z^{\Gamma^b}(\cdot)$.

⁹See, for example, Battigalli and Siniscalchi (2002), Battigalli (2003), and the relevant references therein. Note that here “**rationality**” means only expected utility maximization given whatever *subjective* beliefs a player holds about exogenous uncertainty and coplayers’ behavior. Every other restriction on behavior is the result of additional assumptions on interactive beliefs.

¹⁰In “all type structures,” we include those that violate the common prior assumption, although a version of this equivalence result holds if attention is restricted to type structures with a common *marginal* prior on payoff types. See Battigalli and Siniscalchi (2003) and Bergemann and Morris (2017).

implementation with respect to rationalizability for games with payoff uncertainty.¹¹ The intuition for this result is relatively straightforward: (probability-1) belief is a **monotone** operator, that is, believing a weak proposition (large event) is easier than believing a logically stronger proposition (smaller event included in the former one). It follows by an induction argument that common belief in rationality and in contextual restrictions on exogenous interactive beliefs (which yields rationalizability in Bayesian games) implies mere common belief in rationality. Since “no restriction” is a particular kind of contextual restriction, the robustness result follows (cf. Ziegler, 2022). With this, we refer to robust implementation with static mechanisms also as “implementation under RCBR.”

Finally, we are going to consider a weaker form of “virtual implementation,” or **v-implementation**, that only requires to approximate the desired outcome $f(\theta)$ with an arbitrary degree of precision (see Abreu and Matsushima (1992) and Bergemann and Morris (2009)). Clearly, robust v-implementation is easier to achieve than robust implementation. But Bergemann and Morris (2009) show that—within the domain of static mechanisms—even this form of implementation under RCBR is hard when valuations depend on the types of others, as illustrated by the following example.

EXAMPLE 2: A single good must be allocated to one of many agents through a static mechanism with monetary transfers. Each agent/player i values the good

$$v_i(\theta_i, \theta_{-i}) = \theta_i + \gamma \sum_{j \neq i} \theta_j \quad (\gamma \geq 0),$$

where θ_i is private information of i and belongs to a finite set of payoff types Θ_i that satisfies $\{0, 1\} \subseteq \Theta_i \subseteq [0, 1]$. As i 's valuation also depends on θ_{-i} , players have interdependent valuations for the good. The degree of interdependence is increasing in γ . It turns out that, for $\gamma > \frac{1}{|I|-1}$, only constant social choice functions can be v-implemented under RCBR with static mechanisms.

2.3. Robust Implementation, Sequential Mechanisms

Allowing for sequential mechanisms may expand the set of robustly implementable SCFs, but there are different versions of rationalizability for sequential games, characterizing the behavioral implications of different specifications of “common belief in rationality.”¹² The weakest one, aka “weak rationalizability” or “initial rationalizability,” relies on the assumption of *Rationality and Common Initial Belief in Rationality* (RCIBR), see Battigalli (2003)). Therefore, we refer to (robust v-) implementation with respect to this version of rationalizability as “implementation under RCIBR.”

Since initial (probability-1) belief is monotone, the aforementioned results for static mechanisms extend to sequential mechanisms, a (weak) version of perfect Bayesian equilibrium, and implementation under RCIBR. However, since weak rationalizability typi-

¹¹Ollár and Penta (2017, 2023) study rationalizable implementation subject to some natural belief restrictions. This can be interpreted as a form of partial robustness. See Artemov, Kunimoto, and Serrano (2013) and Propositions 4.2–4.3 in Battigalli and Siniscalchi (2003).

¹²Where “rationality” is now meant in the *sequential* sense of subjective expected utility maximization *conditional on* observations about previous moves.




FIGURE 1.—Game tree of a sequential mechanism.

cally allows for a large set of outcomes, it is unlikely that relevant SCFs can be implemented under RCIBR. For instance, in Example 2, even allowing for sequential mechanisms, only constant SCFs can be implemented under RCIBR if $\gamma > \frac{1}{|I|-1}$.¹³

As mentioned in the Introduction and intuitively explained in Section 2.1, a stronger and more interesting notion of rationalizability for sequential games captures a form of forward induction (FI) reasoning, as it characterizes the behavioral implications of *Rationality and Common Strong Belief in Rationality* (RCSBR). The simplest version of rationalizability capturing RCSBR in incomplete-information environments is **strong rationalizability** for games with payoff uncertainty, a kind of “belief-free strong rationalizability” (Battigalli and Siniscalchi (2002)). Therefore, we refer to implementation with respect to strong rationalizability as “implementation under RCSBR.”

Clearly, strong rationalizability refines weak/initial rationalizability. Thus, allowing for sequential mechanisms, v-implementation under RCSBR might significantly expand the set of v-implementable SCFs. Indeed, considering a discretized environment, Müller (2016) shows precisely this. We illustrate this for a simple specification of the economic environment of Example 2.

EXAMPLE 3: Let $I = \{\text{Ann, Bob, Cora}\}$, $\Theta_i = \{0, 1\}$ for each $i \in I$, $\gamma = 2/3$ and consider a planner who wants

- to assign the good with equal probability to one of the players i with $\theta_i = 1$ (high type), if any, and to keep the good otherwise ($q_i(\theta) = |\{j : \theta_j = 1\}|^{-1}$ if $\theta_i = 1$, $q_i(\theta) = 0$ if $\theta_i = 0$);
- to extract most (90%) of the expected value from each high type; thus, low types should pay nothing, a high type of i should pay $-t_i = 0.9$ if there are no other high types, 0.75 if there is one more high type, and 0.7 if all types are high.

As explained above, this nonconstant SCF cannot be robustly implemented by static mechanisms. Consider now the following sequential mechanism.

- *Game tree*: Ann, Bob, and Cora (in this order) sequentially send a message/report in $\{0, 1\}$ with perfect information about previous moves, with the partial exception of Bob, who can send a message in $\{0, 1, \hat{1}\}$ if Ann reports 1; see Figure 1.

- *Outcome function*: after a sequence of three reports in $\{0, 1\}$, the outcome function mimics the SCF (assuming truthful reporting); for a sequence where Bob sends message

¹³See Müller (2016) and Müller (2020).

$\hat{1}$, let

$$(q; t) = (0.98, 0.02, 0; -1.4, -0.015, 0) \quad \text{if } (1, \hat{1}, 0),$$

$$(q; t) = (0.49, 0.02, 0.49; -0.7, -0.03, -0.7) \quad \text{if } (1, \hat{1}, 1).$$

With this, low types prefer to report 0, unless they believe that both coplayers are high types but (will) report 0, whereas high types prefer to report 1 if they believe that there are at least as many coplayers’ high types as coplayers’ high reports. Moreover, after Ann reports 1, the high type of Bob prefers message $\hat{1}$ over 0 if he believes that then Cora will report truthfully, regardless of whether Ann is truly of type 1 or not. This is because, after message $\hat{1}$, he can obtain the good with small positive probability, but at a “discounted price.”

Here, we provide a brief, intuitive illustration of the steps of strong rationalizability; the full procedure with complete arguments can be found in the Supplemental Appendix (Battigalli and Catonini (2026)). Let A0, A1, B0, . . . denote players with their type (Ann of type 0, etc.). **Step 1:** C0 reports 0 at each node (sequence of messages) except (0, 0) (the payments are just too high for the low type, even if the coplayers are believed to be of high type). Analogously, B0 eliminates report 1 at node (1). By contrast, C1 reports 1 at (0, 0). **Step 2** only involves B1 and relies on the additional message $\hat{1}$. Since—according to the first step—A0 may report 1, then B1 does not necessarily prefer to report 1 over 0, as the transfer would exceed the expected value in case Ann’s type is 0 and Cora reports truthfully. However, no matter his beliefs, B1 does prefer to send message $\hat{1}$ instead of 0, because the discounts and C0’s truthful reporting makes it worth pursuing the good even if Ann’s type is 0. Moreover, given that C0 will not report 1, B1 prefers to report 1 over 0 if Ann reports 0. So, no matter Ann’s report, we can rule out report 0 for B1. **Step 3:** Because of this, A0 reports 0. Moreover, by forward induction, C0 reports 0 also at (0, 0), realizing that Bob’s type must be low. **Step 4** is triggered by the *forward-induction* consideration that Ann’s type must be high if she reported 1. With this, B1 and C1 report 1 afterward. **Step 5:** B0 reports 0 at (1) as C1 would not report 0 after (1, $\hat{1}$). Thus, message $\hat{1}$ ends up off path. **Further steps:** Next, A1 rules out report 0. This triggers the last two steps of forward induction, which induce B0 and C1 to report truthfully if Ann reports 0.

The conclusion of strong rationalizability is that every type of every player reports truthfully, and thus the social choice function is implemented. But to get there, we had to depart from a mere sequential-revelation mechanism and introduce a third message for Bob.¹⁴ This message is eventually discarded, and hence, as in Example 1, we have an off-path information set for the last mover.

But is v -implementation under RCSBR *robust*? In other words, suppose agents’ interactive exogenous beliefs about each other’s payoff types satisfy some contextual restrictions represented by a type structure \mathcal{T} . Then their behavior should satisfy *strong rationalizability for the Bayesian game* $\Gamma^b = (\mathcal{M}, \mathcal{E}, \mathcal{T})$. Robustness would require that the given SCF f is v -implementable with respect to strong rationalizability in Bayesian games across all type structures \mathcal{T} . Unfortunately, we cannot replicate the aforementioned inductive argument based on the monotonicity of probability-1 (initial) belief, because *strong belief is*

¹⁴In more complex economic environments, one also needs to give up on perfect information. See the canonical mechanisms of Müller (2016).

not monotone: As illustrated by Example 1, while at the beginning of the game it is easier to believe a weak proposition such as “my coplayers are rational” than a stronger one such as “my coplayers are rational and their exogenous beliefs satisfy the contextual restrictions,” there typically are more observations consistent with the weaker proposition and, therefore, more instances in which strong belief requires assigning probability 1 to this proposition, making it more difficult to strongly believe it. When contextual considerations (e.g., social norms) also shape endogenous beliefs about behavior, unlike Example 1, it is easy to show that the set of induced *paths of play* is nonmonotone with respect to such contextual restrictions (see Battigalli and Friedenberg (2012)).

By contrast, we prove that *the set of state-dependent strongly rationalizable paths of play is (nonempty and) monotone with respect to restrictions on exogenous beliefs*. With this, we can also prove that *v-implementation under RCSBR is robust*: Fix an SCF $f : \Theta \rightarrow Y$. Let $\Gamma = (\mathcal{M}, \mathcal{E})$ denote the game with payoff uncertainty (or “belief-free” game) induced by mechanism \mathcal{M} with outcome function $g : Z \rightarrow Y$ in environment \mathcal{E} , and let $\theta \mapsto \mathcal{Z}^\Gamma(\theta)$ denote the strongly rationalizable-paths correspondence. Suppose that, for all states θ , $g(\mathcal{Z}^\Gamma(\theta)) \approx \{f(\theta)\}$ to an arbitrary degree of precision. Now suppose that the relevant hierarchies of initial beliefs on the payoff-relevant uncertainty are represented by a particular type structure \mathcal{T} . It is without loss of generality to write types as $t_i = (\theta_i, e_i)$ where coordinate e_i affects hierarchical exogenous beliefs, but does not affect payoffs, whereas θ_i can affect both payoffs and beliefs. Appending \mathcal{T} to $(\mathcal{M}, \mathcal{E})$ gives a sequential Bayesian game $\Gamma^b = (\mathcal{M}, \mathcal{E}, \mathcal{T})$. In our analysis, we make the transition from game with payoff uncertainty Γ to Bayesian game Γ^b in two steps. First, we “duplicate” types by replacing each set Θ_i of payoff types with a set $T_i = \Theta_i \times E_i$, and we note that the solution concept is invariant to such duplications: a pair (θ_i, s_i) is strongly rationalizable if and only if $((\theta_i, e_i), s_i)$ is strongly rationalizable in the “belief-free” game with duplicated types for each $e_i \in E_i$. Next, we obtain a type structure \mathcal{T} by adding belief maps $(\beta_i : T_i \rightarrow \Delta(T_{-i}))_{i \in I}$ to such game with duplicated types, which corresponds to specific restrictions on exogenous beliefs in this game: for each type $t_i = (\theta_i, e_i)$, the set of possible exogenous beliefs is the singleton $\{\beta_i(t_i)\}$. With this, our theorem implies that, for all Bayesian games Γ^b obtained by appending a type structure to Γ , $\emptyset \neq \mathcal{Z}^{\Gamma^b}(\theta, e) \subseteq \mathcal{Z}^\Gamma(\theta)$, where $\mathcal{Z}^{\Gamma^b}(\theta, e)$ is the set of strongly rationalizable paths at state (θ, e) in Γ^b . Therefore, $g(\mathcal{Z}^{\Gamma^b}(\theta, e)) \approx \{f(\theta)\}$ for all such games Γ^b and states (θ, e) to an arbitrary degree of precision.¹⁵

3. PRELIMINARIES¹⁶

In this section, we formally introduce the basic incomplete-information framework (3.1), systems of beliefs and sequential best replies (3.2), and the adopted solution concept, strong directed rationalizability (3.3).

3.1. Multistage Games With Payoff Uncertainty

We consider *finite* multistage games with *payoff uncertainty* and *observed actions*. The latter assumption simplifies our notation, but our analysis can be seamlessly extended to games with imperfectly observed actions.¹⁷

¹⁵In the Supplemental Appendix, we illustrate this result in the context of Example 3.

¹⁶The formalism is based on Battigalli, Catonini, and De Vito (2025), Chapters 9 and 15

¹⁷See Section 7, where we also consider extensions to infinite games.

There is a set of players I and each $i \in I$ has a set of potentially available actions A_i . Let $A = \times_{i \in I} A_i$ denote the set of action profiles and $A^{<\mathbb{N}_0}$ the set of finite sequences of such profiles (including the empty sequence \emptyset). A subset of $A^{<\mathbb{N}_0}$ is a **tree** with root \emptyset (the empty sequence) if it is closed under the “prefix-of” precedence relation \preceq (note that \emptyset is a prefix of every sequence). The rules of the game yield a tree $\bar{H} \subseteq A^{<\mathbb{N}_0}$ of possible sequences, called **histories**, and a feasibility correspondence $h \mapsto \mathcal{A}(h) = \{a \in A : (h, a) \in \bar{H}\}$ such that (1) $\mathcal{A}(h) = \times_{i \in I} A_i(h)$ and (2) $\mathcal{A}(h) = \emptyset$ implies $A_i(h) = \emptyset$ for every $i \in I$. The set of terminal histories—or possible paths of play—is $Z = \{z \in \bar{H} : \mathcal{A}(h) = \emptyset\}$, and the set of nonterminal histories is $H = \bar{H} \setminus Z$. Nonterminal histories are publicly observed as soon as they realize.

Each player i *knows* the true value of a payoff-relevant parameter θ_i , called the **payoff type** of i , whereas the set Θ_i of possible values of θ_i is common knowledge. The parameterized payoff function of player i is $u_i : \Theta \times Z \rightarrow \mathbb{R}$, where $\Theta = \times_{i \in I} \Theta_i$ is the set of all possible type profiles, or payoff states. Payoff uncertainty is represented by the dependence of u_i on θ . When convenient, we write $u_{i,\theta} : Z \rightarrow \mathbb{R}$ for the section of u_i at payoff state θ .¹⁸ Thus, a multistage game with payoff uncertainty and observed actions is given by $\Gamma = \langle I, \bar{H}, (\Theta_i, u_i)_{i \in I} \rangle$, where all the featured sets are finite. If $|\mathcal{A}_i(h)| > 1$, then player i is active at nonterminal history h . If $|\mathcal{A}_i(h)| = 1$, player i is inactive and the unique element of $\mathcal{A}_i(h)$ can be thought of as a waiting action. If there is only one active player for each $h \in H$, then Γ features perfect (albeit incomplete) information, that is, there are no simultaneous moves and (by the observed-actions assumption) past moves are perfectly observed. In the analysis of examples, we omit to mention the waiting actions of inactive players.

We interpret each function u_i as the composition of a parameterized **utility function** $v_i : \Theta \times Y \rightarrow \mathbb{R}$ (Y is the relevant space of outcomes) and an **outcome function** $g : Z \rightarrow Y$ specified by the rules of the game: $u_i(\theta, z) = v_i(\theta, g(z))$.

From these primitives, we can derive a set of **strategies** $S_i = \times_{h \in H} \mathcal{A}_i(h)$ for each player i . Let $S = \times_{i \in I} S_i$ and $S_{-i} = \times_{j \neq i} S_j$. Note that we take an *interim perspective*: the game starts with some exogenously given payoff state θ (e.g., representing players’ traits), imperfectly and asymmetrically known by the players. Thus, strategies only describe how behavior depends on previous moves. Let $\zeta : S \rightarrow Z$ denote the **path function** that associates each strategy profile $s = (s_i)_{i \in I} \in S$ with the induced path $z = \zeta(s)$. With this, we define the (parameterized) strategic-form payoff function of player i as

$$U_i : \Theta \times S \longrightarrow \mathbb{R} \\ (\theta, s) \mapsto u_i(\theta, \zeta(s))$$

Finally, for each $h \in \bar{H}$,

$$S(h) = S_i(h) \times S_{-i}(h) = \{(s_i, s_{-i}) \in S : h \preceq \zeta(s)\}$$

denotes the set of all strategy profiles inducing h .

The primitive and derived elements are respectively summarized by the following table:

¹⁸Since the profile of payoff types $\theta = (\theta_i)_{i \in I}$ determines each payoff function $u_{i,\theta} : Z \rightarrow \mathbb{R}$, we are implicitly assuming that there is distributed knowledge of the payoff state. This assumption is made only to simplify the notation.

| Symbol | Terminology |
|-------------------------------------------------------------------------|------------------------------------------------|
| $i \in I$ | players |
| $a_i \in A_i, a \in A = \times_{i \in I} A_i$ | actions of i , action profiles |
| $h \in \bar{H} \subseteq A^{<\mathbb{N}_0}$ | histories (\bar{H} is a tree) |
| $\mathcal{A}_i(h) (\mathcal{A}(h) = \times_{i \in I} \mathcal{A}_i(h))$ | feasible actions (action profiles) given h |
| $z \in Z$ | terminal histories, or paths of play |
| $H = \bar{H} \setminus Z$ | nonterminal histories |
| $\theta_i \in \Theta_i$ | payoff types of i |
| $\theta \in \Theta = \times_{i \in I} \Theta_i$ | payoff states |
| $u_i : \Theta \times Z \rightarrow \mathbb{R}$ | (parameterized) payoff function of i |
| $s_i \in S_i = \times_{h \in H} \mathcal{A}_i(h)$ | strategies of i |
| $s \in S = \times_{i \in I} S_i = S_i \times S_{-i}$ | strategy profiles |
| $s \in S(h) = S_i(h) \times S_{-i}(h)$ | strategy profiles inducing h |
| $\zeta : S \rightarrow Z$ | path function |
| $U_i : \Theta \times S \rightarrow \mathbb{R}$ | (param.) strategic-form payoff function of i |

3.2. Beliefs and Best Replies

We model the beliefs of each player i as the play unfolds by means of **conditional probability systems** (CPSs): observed history h reveals that the set of possible type-strategy profiles of the coplayers is $\Theta_{-i} \times S_{-i}(h)$; thus, we consider arrays of conditional beliefs $\mu_i = (\mu_i(\cdot | \Theta_{-i} \times S_{-i}(h)))_{h \in H}$ over such profiles. The set of CPSs of player i , denoted $\Delta^H(\Theta_{-i} \times S_{-i})$, is the subset of arrays of beliefs $\mu_i \in (\Delta(\Theta_{-i} \times S_{-i}))^H$ such that, for every $h \in H$, $\mu_i(\Theta_{-i} \times S_{-i}(h) | \Theta_{-i} \times S_{-i}(h)) = 1$ and the *chain rule* holds, that is, for all $h, h' \in H$ and $E \subseteq \Theta_{-i} \times S_{-i}(h')$,

$$S_{-i}(h') \subseteq S_{-i}(h) \\ \Rightarrow \mu_i(E | \Theta_{-i} \times S_{-i}(h)) = \mu_i(E | \Theta_{-i} \times S_{-i}(h')) \mu_i(\Theta_{-i} \times S_{-i}(h') | \Theta_{-i} \times S_{-i}(h)).$$

Note that $h \preceq h'$ implies $S_{-i}(h') \subseteq S_{-i}(h)$, but the converse is not true because histories also represent behavior of player i . From now on, we use the simplified notation $\mu_i = (\mu_i(\cdot | h))_{h \in H}$

We will consider type-dependent **restrictions on players' exogenous beliefs** (i.e., initial beliefs about the types of others), represented by subsets of probability measures: for all $i \in I$ and $\theta_i \in \Theta_i$, $\hat{\Delta}_{i, \theta_i} \subseteq \Delta(\Theta_{-i})$. With this, we introduce profiles $\Delta = (\Delta_{i, \theta_i})_{i \in I, \theta_i \in \Theta_i}$ of type-dependent subsets of CPSs such that, for all $i \in I$ and $\theta_i \in \Theta_i$,

$$\Delta_{i, \theta_i} = \left\{ \mu_i \in \Delta^H(\Theta_{-i} \times S_{-i}) : \text{marg}_{\Theta_{-i}} \mu_i(\cdot | \emptyset) \in \hat{\Delta}_{i, \theta_i} \right\}.$$

We represent the behavior of a rational player i of type θ_i by means of a (weak) **sequential best-reply** correspondence $\mu_i \mapsto r_{i, \theta_i}(\mu_i)$ defined as follows. Let $\mathcal{H}_i(s_i) = \{h \in H : s_i \in S_i(h)\}$ denote the set of nonterminal histories that can occur if s_i is played. With this,

$$r_{i, \theta_i}(\mu_i) = \left\{ \bar{s}_i \in S_i : \forall h \in \mathcal{H}_i(\bar{s}_i), \bar{s}_i \in \arg \max_{s_i \in S_i(h)} \mathbb{E}_{\mu_i(\cdot | h)}(U_i(\theta_i, s_i, \cdot)) \right\}.$$

By standard dynamic programming arguments, $r_{i, \theta_i}(\mu_i) \neq \emptyset$ for all payoff types θ_i and CPSs μ_i (see the analysis and discussion in Battigalli, Catonini, and Manili (2023)).

Fix a CPS μ_i and a type θ_i . For each strategy \bar{s}_i and history $h \in \mathcal{H}_i(\bar{s}_i)$, we say that \bar{s}_i is a **continuation best reply** to $\mu_i(\cdot|h) \in \Delta(\Theta_{-i} \times S_{-i}(h))$ for θ_i if, for every $s_i \in S_i(h)$,

$$\mathbb{E}_{\mu_i(\cdot|h)}(U_i(\theta_i, \bar{s}_i, \cdot)) \geq \mathbb{E}_{\mu_i(\cdot|h)}(U_i(\theta_i, s_i, \cdot)).$$

Thus, \bar{s}_i is a (weak) sequential best reply to μ_i for θ_i if \bar{s}_i is a continuation best reply to $\mu_i(\cdot|h)$ for θ_i at every $h \in \mathcal{H}_i(\bar{s}_i)$.

3.3. Strong (Directed) Rationalizability

As informally explained in the **Introduction**, our forward-induction analysis hinges on the notion of “strong belief.” For each event $E_{-i} \subseteq \Theta_{-i} \times S_{-i}$ (e.g., that coplayers’ behavior is consistent with rationality), we say that a CPS μ_i **strongly believes** E_{-i} if μ_i assigns probability 1 to E_{-i} as long as E_{-i} is not contradicted by observation:

$$\forall h \in H, \quad E_{-i} \cap (\Theta_{-i} \times S_{-i}(h)) \neq \emptyset \quad \Rightarrow \quad \mu_i(E_{-i}|h) = 1.$$

We assume that players are *rational* and that the restrictions on exogenous beliefs are **transparent**, that is, the belief restrictions hold and there is common belief of this fact conditional on every nonterminal history. Moreover, we assume that players *strongly believe* that

- the coplayers are rational and the restrictions are transparent;
- the coplayers are rational, the restrictions are transparent, and the coplayers strongly believe that everyone else is rational and that the restrictions are transparent;
- and so on.

In brief, we assume *rationality, transparency of the belief restrictions, and common strong belief thereof*.

The previous hypotheses can be made formal in the language of epistemic game theory. As shown by **Battigalli and Prestipino (2013)**, the behavioral implications of these epistemic hypotheses are characterized by **Strong Directed Rationalizability** (**Battigalli (2003)**, **Battigalli and Siniscalchi (2003)**).¹⁹

Fix a profile $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of subsets of CPSs (see Section 3.2). Also, for each player $i \in I$ and event $E_{-i} \subseteq \Theta_{-i} \times S_{-i}$, let $\Delta_{\text{sb}}^H(E_{-i})$ denote the set of CPSs μ_i that **strongly believe** E_{-i} , and let $R_i^{\Delta,0} = \Theta_i \times S_i$. Then, for each $n > 0$, define the set of **strongly Δ - n -rationalizable** type-strategy pairs of i as

$$R_i^{\Delta,n} = \left\{ (\theta_i, s_i) : \exists \mu_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(R_{-i}^{\Delta,m}) \cap \Delta_{i,\theta_i}, s_i \in r_{i,\theta_i}(\mu_i) \right\}.$$

With this, the set of strongly Δ - n -rationalizable strategies for θ_i is the section at θ_i of $R_i^{\Delta,n}$,

$$S_i^{\Delta,n}(\theta_i) = \{s_i \in S_i : (\theta_i, s_i) \in R_i^{\Delta,n}\},$$

and the set of strongly Δ - n -rationalizable strategy profiles at state θ is

$$S^{\Delta,n}(\theta) = \times_{i \in I} S_i^{\Delta,n}(\theta_i).$$

¹⁹These articles use the term “(strong) Δ -rationalizability.” Recall that we use “(strong) directed rationalizability” to refer to the correspondence that associates each profile of belief restrictions Δ with the corresponding strongly rationalizable behavior, so that Δ “directs” the resulting behavior.

Finally, let

$$R_i^{\Delta,\infty} = \bigcap_{n>0} R_i^{\Delta,n}, \quad R^{\Delta,\infty} = \times_{i \in I} R_i^{\Delta,\infty}$$

denote the set of strongly Δ -rationalizable type-strategy pairs of i and profiles of such pairs, and let

$$S_i^{\Delta,\infty}(\theta_i) = \{s_i \in S_i : (\theta_i, s_i) \in R_i^{\Delta,\infty}\}, \quad S^{\Delta,\infty}(\theta) = \times_{i \in I} S_i^{\Delta,\infty}(\theta_i).$$

Recalling that the sequential best-reply correspondence is nonempty valued and noting that mere restrictions on exogenous beliefs cannot contradict the restrictions on beliefs about type-dependent behavior implied by strategic reasoning, one can prove by induction the following result.

LEMMA 1—cf. Battigalli (2003): *Since Δ represents restrictions on exogenous beliefs, for each $\theta \in \Theta$, the set of strongly Δ -rationalizable strategy profiles is nonempty: $S^{\Delta,\infty}(\theta) \neq \emptyset$.*

When there are no actual belief restrictions, that is, when each Δ_{i,θ_i} is the set $\Delta^H(\Theta_{-i} \times S_{-i})$ of all CPSs of i , Strong Δ -Rationalizability boils down to **Strong Rationalizability** (Pearce (1984), Battigalli (1997)), which characterizes the behavioral implications of *Rationality and Common Strong Belief in Rationality* (Battigalli and Siniscalchi (2002)). We omit the superscript Δ to denote Strong Rationalizability: R_i^∞ (R_i^n) is the set of strongly (n -)rationalizable pairs of i , $S_i^\infty(\theta_i)$ ($S_i^n(\theta_i)$) is the set of strongly (n -)rationalizable strategies of θ_i , and $S^\infty(\theta)$ ($S^n(\theta)$) is the set of profiles of strongly (n -) rationalizable strategies at payoff state θ .

A path (terminal history) $z \in Z$ is strongly Δ -rationalizable if there exists some strongly Δ -rationalizable profile (θ, s) such that $\zeta(s) = z$. Let $\mathcal{Z}(R^{\Delta,\infty})$ denote the set of **strongly Δ -rationalizable paths**. The set of strongly Δ -rationalizable paths **at payoff state θ** is $\zeta(S^{\Delta,\infty}(\theta))$.

The signaling game informally analyzed in Section 2.1 illustrates the formalism and concepts introduced in this section.

EXAMPLE 4: Consider again the signaling game of Example 1. Game Γ is a two-stage game with perfect information, with

$$\begin{aligned} \Theta_1 &= \{x, y, z\}, & \Theta_2 &= \{\bar{\theta}_2\} \text{ (a singleton),} \\ H &= \{\emptyset, (\ell), (r)\}, & Z &= \{(\ell, a), (\ell, b), (r, c), (r, d), (r, e)\}, & \bar{H} &= H \cup Z, \\ \mathcal{A}_1(\emptyset) &= \{\ell, r\}, & \mathcal{A}_2(\ell) &= \{a, b\}, & \mathcal{A}_2(r) &= \{c, d, e\}, \end{aligned}$$

and the type-dependent payoff functions $u_i : \Theta \times Z \rightarrow \mathbb{R}$ described by the following tables:

| | | |
|----------------------------------------------------|-----|-----|
| $u_1(\cdot, \ell, \cdot), u_2(\cdot, \ell, \cdot)$ | a | b |
| $\theta_1 = x$ | 3 1 | 1 0 |
| $\theta_1 = y$ | 1 0 | 1 1 |
| $\theta_1 = z$ | 3 1 | 1 0 |

| | | | |
|----------------------------------------------|-----|-----|-----|
| $u_1(\cdot, r, \cdot), u_2(\cdot, r, \cdot)$ | c | d | e |
| $\theta_1 = x$ | 0 0 | 0 0 | 0 1 |
| $\theta_1 = y$ | 0 0 | 0 1 | 3 0 |
| $\theta_1 = z$ | 0 1 | 2 0 | 2 0 |

Since player 2 is uninformed and inactive in the first stage, $\Theta_2 \times S_2$ is isomorphic to S_2 and $\Delta^H(\Theta_2 \times S_2)$ is isomorphic to $\Delta(S_2)$ (by the chain rule and $S_2(\ell) = S_2(r) = S_2$). We intuitively explained in Example 1 how strong directed rationalizability works in this game.

Thus, we only list below the formal result for each step using the notation introduced above. Without belief restrictions, we have

$$\begin{aligned}
 R_1^1 &= \{(x, \ell), (y, \ell), (y, r), (z, \ell), (z, r)\} \text{ (thus, } S_1^1(x) = \{\ell\}), R_2^1 = S_2; \\
 R_1^2 &= R_1^1, R_2^2 = \{a, b\} \times \{c, d\}; \\
 R_1^3 &= \{(x, \ell), (y, \ell), (z, \ell), (z, r)\}, \text{ (thus, } S_1^3(y) = \{\ell\}), R_2^3 = R_2^2; \\
 R_1^4 &= R_1^3, R_2^4 = \{a.c, b.c\}; \\
 R_1^5 &= \Theta_1 \times \{\ell\} \text{ (thus, } S_1^5(\theta_1) = \{\ell\} \text{ for all } \theta_1), R_2^5 = R_2^4; \\
 R_1^\infty &= \Theta_1 \times \{\ell\}, R_2^\infty = \{a.c, b.c\}; \zeta(S^{\Delta, \infty}(\theta)) = \{\ell\} \times \{a, b\} \text{ for all } \theta.
 \end{aligned}$$

Since Θ_2 is a singleton, we can only have restrictions on the exogenous beliefs of player 2. Formalizing Example 1, let

$$\Delta_2 = \{\mu_2 \in \Delta^H(\Theta_1 \times S_1) : \mu_2(\{z\} \times S_1 | \emptyset) = 1\}$$

denote the set of CPSs that initially assign probability 1 to type $\theta_1 = z$. With this, strong Δ -rationalizability yields

$$\begin{aligned}
 R_1^{\Delta, 1} &= \{(x, \ell), (y, \ell), (y, r), (z, \ell), (z, r)\}, R_2^{\Delta, 1} = \{a.c, b.c, a.d, a.e\}; \\
 R_1^{\Delta, 2} &= \{(x, \ell), (y, \ell), (y, r), (z, \ell)\}, R_2^{\Delta, 2} = \{a.c, b.c, a.d\}; \\
 R_1^{\Delta, 3} &= \Theta_1 \times \{\ell\}, R_2^{\Delta, 3} = \{a.d\}; \\
 R_1^{\Delta, \infty} &= \Theta_1 \times \{\ell\}, R_2^{\Delta, \infty} = \{a.d\}, \zeta(S^{\Delta, \infty}(\theta)) = \{(\ell, a)\} \text{ for all } \theta \in \Theta.
 \end{aligned}$$

Thus, $\mathcal{Z}(R^{\Delta, \infty}) \subset \mathcal{Z}(R^\infty)$; but $R_2^{\Delta, \infty} \not\subseteq R_2^\infty$, actually $R_2^{\Delta, \infty} \cap R_2^\infty = \emptyset$.

4. MAIN THEOREM

We show that *when we consider only restrictions on exogenous beliefs the set of strongly Δ -rationalizable paths is monotone in Δ* , despite the nonmonotonicity of strong belief.

Because it suffices for our application to implementation theory, here we just focus on the comparison between some profile Δ of subsets of CPSs that only restrict exogenous beliefs, and the case of no restrictions ($\Delta_{i, \theta_i} = \Delta^H(\Theta_{-i} \times S_{-i})$ for all i and θ_i , that is, strong rationalizability). Thus, we prove that for any fixed profile of restrictions on exogenous beliefs Δ the set of strongly Δ -rationalizable paths is contained in the set of strongly rationalizable paths. It will be clear that the proof can be easily adapted to obtain the more general path monotonicity claim.²⁰

THEOREM 1: *Fix a profile $\Delta = (\Delta_{i, \theta_i})_{i \in I, \theta_i \in \Theta_i}$ of restrictions on exogenous beliefs. Then, for all steps $n \geq 0$ and states $\theta \in \Theta$, $\emptyset \neq \zeta(S^{\Delta, n}(\theta)) \subseteq \zeta(S^n(\theta))$, that is, for each $(\theta, s) \in R^{\Delta, \infty} \neq \emptyset$, there exists $s' \in S$ such that $(\theta, s') \in R^\infty$ and $\zeta(s) = \zeta(s')$.*

²⁰Moreover, the proof shows (see footnote 23) the path equivalence of strong Δ -rationalizability and an elimination procedure that coincides with strong rationalizability until convergence, and then introduces the Δ -restrictions and continues until convergence (cf. [Catonini \(2019\)](#)). We thank an anonymous referee for this observation.

The assumption that the belief restrictions only apply to exogenous beliefs is tight. In the literature, there are many examples of strong directed rationalizability with restrictions on initial beliefs about coplayers' *behavior* yielding nonstrongly-rationalizable outcomes (see, e.g., Battigalli and Friedenberg (2012) and Catonini (2019)). In the Supplemental Appendix, we provide an analogous example with restrictions on *noninitial* (hence, conditional) beliefs about coplayers' types.

Given that, as we saw in Examples 1 and 4, the two elimination procedures may induce completely disjoint off-path behaviors, proving path monotonicity is hard. Our proof is based on a kind of double induction argument.²¹

4.1. Proof of Theorem 1

Nonemptiness follows from Lemma 1. Here, we only focus on path inclusion. Since comparing directly strong rationalizability and strong Δ -rationalizability is difficult, we construct a sequence of elimination procedures that gradually transform strong Δ -rationalizability into strong rationalizability, and we prove step-by-step path inclusion between each pair of consecutive, "similar" procedures.

Let K be the number of steps that it takes for strong rationalizability to converge: $R^{K-1} \subset R^K = R^\infty$ (\subset denotes *strict* inclusion). Note that K is well-defined because the game is finite. For each $k = 0, \dots, K$, we introduce Procedure k , which performs the *first* k steps of elimination *without* belief restrictions and the *following* steps *with* the belief restrictions. Thus, procedure 0 coincides with strong Δ -rationalizability, while the first K steps of procedure K coincide with strong rationalizability. Hence, the path-inclusions between procedure 0 and procedure 1, procedure 1 and procedure 2, and so on up to procedure K , prove the theorem.

Now we define formally such elimination procedures, denoted by $((X_k^n)_{n=0}^\infty)_{k=0}^K$. If everything is strongly rationalizable, there is nothing to prove; thus, suppose that strong rationalizability deletes some pair (θ_i, s_i) for at least one player i , so that $K > 0$.

As anticipated, for $k = 0$, we have strong Δ -rationalizability:

$$(X_0^n)_{n=0}^\infty = (R^{\Delta,n})_{n=0}^\infty.$$

For each $k = 1, \dots, K$, define $((X_{k,i}^n)_{i \in I})_{n=0}^\infty$ as follows. Let $X_k^0 = \Theta \times S$.

For all $n \in \{1, \dots, k\}$ and $i \in I$,

$$X_{k,i}^n = \left\{ (\theta_i, s_i) \in \Theta_i \times S_i : \exists \mu_i \in \bigcap_{m=0}^{n-1} \Delta_{sb}^H(X_{k,-i}^m), s_i \in r_{i,\theta_i}(\mu_i) \right\}. \tag{1}$$

Thus, for $k > 0$, steps $n = 1, \dots, k$ of procedure k coincide with strong rationalizability: $X_k^n = R^n$ for $n \leq k$.

For all $n > k$ and $i \in I$, let

$$X_{k,i}^n = \left\{ (\theta_i, s_i) \in \Theta_i \times S_i : \exists \mu_i \in \bigcap_{m=0}^{n-1} \Delta_{sb}^H(X_{k,-i}^m) \cap \Delta_{i,\theta_i}, s_i \in r_{i,\theta_i}(\mu_i) \right\}. \tag{2}$$

²¹The techniques we use have common elements with the techniques used by Perea (2018, 2024) and Catonini (2020) in complete-information games to prove, respectively, an *order independence* result for strong rationalizability and an *outcome monotonicity* result for directed rationalizability with respect to initial belief restrictions about the path of play. In particular, like Perea (2018), we decompose the problem of comparing two very different elimination procedures into a chain of pairwise comparisons between more similar procedures, and the proof of a key claim (Claim 4 in the proof of Theorem 1) draws on Catonini (2020).

Thus, procedure k deviates from strong rationalizability from step $n = k + 1$ onwards, because it starts imposing the Δ -restrictions on justifying beliefs only from step $k + 1$.

It follows that, as anticipated, $(X_k^n)_{n=0}^\infty$ is an elimination procedure which coincides with strong rationalizability $(R^n)_{n=0}^\infty$ for the first K steps, so obtaining the strongly rationalizable profiles, but then proceeds to (possibly) delete more profiles by adding the Δ -restrictions. More generally, no procedure needs to converge by step K (although some may converge at an earlier step), but—for our purpose—we can focus on the first K steps of all procedures.

We are going to prove that, for each step of elimination n , *the set of θ -dependent paths that are consistent with step n weakly expands as k increases*, which implies the thesis. To do so, we proceed in this order: first we fix $k \in \{1, \dots, K\}$ and consider procedure $k - 1$ and procedure k ; then we prove the path inclusion between the two procedures at every step of elimination n by induction on n .

First, we provide an intuition of how we exploit the similarity between the two consecutive procedures and how the assumption of exogenous restrictions makes their comparison possible. From this intuition, we will derive the two-fold inductive hypothesis for the formal proof. To simplify notation, we drop the indexes $k - 1$ and k of the two procedures and we call them “ P ” and “ Q ”: $((P_i^n)_{i \in I})_{n=0}^\infty = ((X_{k-1,i}^n)_{i \in I})_{n=0}^\infty$ and $((Q_i^n)_{i \in I})_{n=0}^\infty = ((X_{k,i}^n)_{i \in I})_{n=0}^\infty$. We are also going to apply the notation “ $\cdot|_{\hat{H}}$ ” to (profiles of) strategies or type-strategy pairs in order to restrict the domain of strategies to a subset of histories \hat{H} . Furthermore, for any subset $X \subseteq \Theta \times S$, we let

$$\mathcal{H}(X) = \{h \in H : \exists(\theta, s) \in X, h < \zeta(s)\}$$

denote the set of nonterminal histories that realize for some $(\theta, s) \in X$. With this, for any $X_{-i} \subseteq \Theta_{-i} \times S_{-i}$, to ease notation we also let

$$\mathcal{H}(X_{-i}) = \mathcal{H}(\Theta_i \times S_i \times X_{-i})$$

denote the set of nonterminal histories that realize for some $(\theta_{-i}, s_{-i}) \in X_{-i}$ and $(\theta_i, s_i) \in \Theta_i \times S_i$.

P and Q coincide with strong rationalizability for steps $n \in \{1, \dots, k - 1\}$ and depart at step $n = k$. From now on, let $n = k + 1$; this will bring us to formulate the induction hypothesis of the formal proof with the appropriate indexes.

At step $n - 1 = k$, P adopts the belief restrictions and Q does not, so

$$P^{n-1} \subseteq Q^{n-1}. \tag{3}$$

At step $n = k + 1$ both P and Q adopt the restrictions, but P imposes strong belief in smaller strategy sets and, therefore, *along the paths consistent with these sets*, it remains more restrictive:

$$P^n|_{\mathcal{H}(P^{n-1})} \subseteq Q^n|_{\mathcal{H}(P^{n-1})}. \tag{4}$$

At step $n + 1 = k + 2$, the comparison becomes more complex.

First: Is this step of procedure P still more restrictive than Q regarding beliefs at histories in $\mathcal{H}(P^{n-1})$ about the coplayers’ types and moves at those histories, as expression (4) seems to suggest?

The answer is yes, but only thanks to the assumption that restrictions only concern exogenous beliefs. Restrictions on the beliefs about the endogenous/strategic uncertainty could allow player i to believe in some $(\theta_{-i}, s_{-i}) \in P_{-i}^n$, but not in any counterpart

$(\theta_{-i}, s'_{-i}) \in Q^n_{-i}$ with $s_{-i}|_{\mathcal{H}(P^{n-1})} = s'_{-i}|_{\mathcal{H}(P^{n-1})}$. The role of restricting only the *initial* beliefs is more subtle. Strong belief in P^n_{-i} and in Q^n_{-i} may induce, by forward-induction reasoning, different beliefs about θ_{-i} at some history $h' \in (\mathcal{H}(P^n_{-i}) \cap \mathcal{H}(Q^n_{-i})) \setminus \mathcal{H}(P^{n-1})$. If there were restrictions on such beliefs at h' , it could well be that some of the beliefs derived from Q^n_{-i} would be incompatible with the restrictions. Via the chain rule, this could also rule out some beliefs at some $h \in \mathcal{H}(P^{n-1})$ such that $h < h'$.

But this is not the end of the story. Strong belief in Q^n_{-i} may be more restrictive, or “differently restrictive,” compared with strong belief in P^n_{-i} regarding behavior outside of $\mathcal{H}(P^{n-1})$. This is because the inclusion of equation (4) is restricted to $\mathcal{H}(P^{n-1})$. Thus, strong belief in Q^n_{-i} may rule out some belief about the reactions of the coplayers to a deviation of i from $\mathcal{H}(P^{n-1})$, which is instead allowed by strong belief in P^n_{-i} . Example 4 may help to understand this point.²² Belief in $R_2^4 = \{a.c, b.c\}$ imposes belief in reaction c after a deviation from the unique on-path signal induced by $R_1^{\Delta,3} = \Theta_1 \times \{\ell\}$. By contrast, belief in $R_2^{\Delta,4} = \{a.d\}$ imposes belief in reaction d . With this, it is conceivable that there might be a deviation from one of the paths consistent with P^{n+1} that player i expects to lead out of $\mathcal{H}(P^{n-1})$ and always be strictly profitable under strong belief in Q^n_{-i} . This is why it is hard to prove that

$$P^{n+1}|_{\mathcal{H}(P^n)} \subseteq Q^{n+1}|_{\mathcal{H}(P^n)}. \tag{5}$$

What guarantees that such a deviation does not exist? We are going to argue that $\mathcal{H}(P^{n-1}) \supseteq \mathcal{H}(Q^n)$, so that no strategy in $Q_i^{n+1} \subseteq Q_i^n$ (i.e., no strategy that player i could ever find profitable at step $n + 1$ of procedure Q) leads out of $\mathcal{H}(P^{n-1})$ (actually, of $\mathcal{H}(P^n) \subseteq \mathcal{H}(P^{n-1})$) if the coplayers follow strategies in Q^n_{-i} , as strongly believed by i at step $n + 1$.

Here is where the similarity between the two procedures comes into play: $\mathcal{H}(P^{n-1}) \supseteq \mathcal{H}(Q^n)$ is a reverse inclusion compared to the path inclusion we want to prove, but with procedure Q one step ahead of procedure P . Thus, to see why the inclusion holds, we must flip the roles of the two procedures and start from the trivial observation that, since $Q^{n-1} \subseteq Q^{n-2}$ and Q and P coincide up to step $n - 2$,

$$Q^{n-1} \subseteq P^{n-2}.$$

Next, we consider step n of Q and step $n - 1$ of P . Both steps use the belief restrictions, as Q introduces the restrictions only one step later than P . Thanks to this similarity, we can argue as above (cf. equation (4)) to obtain

$$Q^n|_{\mathcal{H}(Q^{n-1})} \subseteq P^{n-1}|_{\mathcal{H}(Q^{n-1})}. \tag{6}$$

Thus, since $\mathcal{H}(Q^{n-1}) \supseteq \mathcal{H}(Q^n)$, we have $\mathcal{H}(P^{n-1}) \supseteq \mathcal{H}(Q^n)$, as we wanted to show.

Proving (6) was easy because we could rely on the inclusion $Q^{n-1} \subseteq P^{n-2}$, which is stated for complete strategies. But to continue and prove

$$P^{n+2}|_{\mathcal{H}(P^{n+1})} \subseteq Q^{n+2}|_{\mathcal{H}(P^{n+1})},$$

we need $\mathcal{H}(P^n) \supseteq \mathcal{H}(Q^{n+1})$, that is, we need

$$Q^{n+1}|_{\mathcal{H}(Q^n)} \subseteq P^n|_{\mathcal{H}(Q^n)}, \tag{7}$$

²²The example compares directly strong Δ -rationalizability and strong rationalizability, which cannot formally take the role of procedures P and Q , but it still displays the possible relationship between P and Q that we are illustrating.

and to prove this we run into the same complications we had for (5). However, recall that we were able to prove (5) after showing that $\mathcal{H}(P^{n-1}) \supseteq \mathcal{H}(Q^n)$; we can prove (7) in the same way, with the roles of the two procedures flipped, because $\mathcal{H}(Q^{n-1}) \supseteq \mathcal{H}(P^{n-1})$ by (3).

At this point, considering any $n \geq k$, it should be clear that if we take induction hypotheses of the kind

$$P^n|_{\mathcal{H}(P^{n-1})} \subseteq Q^n|_{\mathcal{H}(P^{n-1})}, \tag{8}$$

$$Q^n|_{\mathcal{H}(Q^{n-1})} \subseteq P^{n-1}|_{\mathcal{H}(Q^{n-1})}, \tag{9}$$

then we can use them to prove the next iteration of (9), namely

$$Q^{n+1}|_{\mathcal{H}(Q^n)} \subseteq P^n|_{\mathcal{H}(Q^n)}, \tag{10}$$

and we can use (8) and (10) to prove the next iteration of (8).

Now we formulate this two-fold induction hypothesis for the formal proof. For every $n \geq k$,

IH1(n) for every $i \in I$ and $(\theta_i, s_i) \in X_{k-1,i}^n$, there is $\hat{s}_i^{(\theta_i, s_i)} \in S_i$ such that $(\theta_i, \hat{s}_i^{(\theta_i, s_i)}) \in X_{k,i}^n$ and $\hat{s}_i^{(\theta_i, s_i)}(h) = s_i(h)$ for all $h \in \mathcal{H}(X_{k-1,i}^{n-1})$ (thus, step n of Procedure $k - 1$ path-refines step n of Procedure k);

IH2(n) for every $i \in I$ and $(\theta_i, s_i) \in X_{k,i}^n$, there is $\tilde{s}_i^{(\theta_i, s_i)} \in S_i$ such that $(\theta_i, \tilde{s}_i^{(\theta_i, s_i)}) \in X_{k-1,i}^{n-1}$ and $\tilde{s}_i^{(\theta_i, s_i)}(h) = s_i(h)$ for all $h \in \mathcal{H}(X_k^{n-1})$ (thus, step n of procedure k path-refines step $n - 1$ of procedure $k - 1$).

For $n = K$, IH1 implies that, for each $(\theta, s) \in X_{k-1,i}^K$, there exists $s' \in S$ such that $(\theta, s') \in X_k^K$ and $\zeta(s) = \zeta(s')$. Since k is arbitrary in $\{1, \dots, K\}$, this implies that for each $(\theta, s) \in X_0^K \supseteq R^{\Delta, \infty}$, there exists $s' \in S$ such that $(\theta, s') \in X_K^K = R^\infty$ and $\zeta(s) = \zeta(s')$, that is, strong Δ -rationalizability path-refines strong rationalizability.²³

The rest of this section is devoted to proving IH1 and IH2 by way of induction, following the strategy we outlined above. The formal proofs of Claims 1–4 stated below are deferred to the [Appendix](#).

Basis steps

IH2($n = k$) comes from the observation that, by inspection of (1), $X_k^k \subseteq X_k^{k-1} = R^{k-1} = X_{k-1,i}^{k-1}$; IH1($n = k$) comes from (for all $i \in I$)

$$\begin{aligned} X_{k-1,i}^k &= \left\{ (\theta_i, s_i) \in \Theta_i \times S_i : \exists \mu_i \in \bigcap_{m=0}^{k-1} \Delta_{\text{sb}}^H(X_{k-1,-i}^m) \cap \Delta_{i,\theta_i}, s_i \in r_{i,\theta_i}(\mu_i) \right\} \\ &\subseteq \left\{ (\theta_i, s_i) \in \Theta_i \times S_i : \exists \mu_i \in \bigcap_{m=0}^{k-1} \Delta_{\text{sb}}^H(X_{k,-i}^m), s_i \in r_{i,\theta_i}(\mu_i) \right\} = X_{k,i}^k, \end{aligned}$$

where the first equality holds by (2), the last equality holds by (1), and the inclusion follows from the fact that only the first set features the belief restrictions and that by (1), $X_{k-1,-i}^m = R_{-i}^m = X_{k,-i}^m$ for all $m = 0, \dots, k - 1$.

²³Fixing $K' \geq K$ such that $X_0^{K'} = X_0^{K'+1}$ and $X_K^{K'} = X_K^{K'+1}$, the same argument based on IH1 yields the path inclusion between procedure 0 (strong Δ -rationalizability) and procedure K , which first performs all the steps of strong rationalizability, and then introduces the Δ -restrictions. The analogous argument that starts from $X_K^{K'+K}$ and uses IH2 yields the opposite path inclusion. So, the two procedures are actually path equivalent.

Inductive Steps

The proofs of the two inductive steps, $\text{IH1}(n)\text{-IH2}(n)\Rightarrow\text{IH1}(n + 1)$ and $\text{IH1}(n)\text{-IH2}(n)\Rightarrow\text{IH2}(n + 1)$, are essentially identical, because both procedures $(\mathbf{X}_{k-1}^n)_{n=0}^\infty$ and $(\mathbf{X}_k^n)_{n=0}^\infty$ are defined by (2) at each step $n > k$. We start from the proof of $\text{IH1}(n)\text{-IH2}(n)\Rightarrow\text{IH2}(n + 1)$. We relegate the proof of $\text{IH1}(n)\text{-IH2}(n)\Rightarrow\text{IH1}(n + 1)$, which uses the previously obtained $\text{IH2}(n + 1)$, to the Supplemental Appendix.

Inductive Step, Part IH2

Suppose $\text{IH1}(n)\text{-IH2}(n)$ hold. We must show that $\text{IH2}(n + 1)$ holds. Fix $i \in I$ and $(\theta_i, s_i) \in \mathbf{X}_{k,i}^{n+1}$. We are going to show the existence of a CPS $\tilde{\mu}_i^{(\theta_i, s_i)} \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(\mathbf{X}_{k-1, -i}^m) \cap \Delta_{i, \theta_i}$ and of a strategy $\tilde{s}_i^{(\theta_i, s_i)} \in r_{i, \theta_i}(\tilde{\mu}_i^{(\theta_i, s_i)}) \subseteq \mathbf{X}_{k-1, i}^n$ such that $\tilde{s}_i^{(\theta_i, s_i)}(h) = s_i(h)$ for all $h \in \mathcal{H}(\mathbf{X}_k^n)$. To ease notation, in what follows we do not make explicit the dependence of either the CPS or the strategy on the fixed pair (θ_i, s_i) , so we write $\tilde{s}_i = \tilde{s}_i^{(\theta_i, s_i)}$ and $\tilde{\mu}_i = \tilde{\mu}_i^{(\theta_i, s_i)}$. Since the choice of $i \in I$ and $(\theta_i, s_i) \in \mathbf{X}_{k,i}^{n+1}$ is arbitrary, this will prove $\text{IH2}(n + 1)$.

The construction of $\tilde{\mu}_i$ and \tilde{s}_i will be based on four claims for which we provide formal proofs in the Appendix; here, before each claim, we only provide the main ingredients of its proof.

By definition of $\mathbf{X}_{k,i}^{n+1}$ (see equation (2)), there is some $\mu_i \in \bigcap_{m=0}^n \Delta_{\text{sb}}^H(\mathbf{X}_{k, -i}^m) \cap \Delta_{i, \theta_i}$ such that $s_i \in r_{i, \theta_i}(\mu_i)$.

Using $\text{IH2}(n)$, we can construct a CPS $\tilde{\mu}_i$ for step n of procedure $k - 1$ that mimics μ_i along the paths that are consistent with step n of procedure k . Consistently with notation used for sets of nonterminal histories and in Example 4, for any $\mathbf{X} \subseteq \Theta \times S$, we let

$$\mathcal{Z}(\mathbf{X}) = \{z \in Z : \exists(\theta, s) \in \mathbf{X}, \zeta(s) = z\}$$

denote the set of possible paths given \mathbf{X} .

CLAIM 1: *There exists $\tilde{\mu}_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(\mathbf{X}_{k-1, -i}^m) \cap \Delta_{i, \theta_i}$ such that, for every $h \in \mathcal{H}(\mathbf{X}_k^n) \cap \mathcal{H}_i(s_i)$,*

$$\forall(\theta_{-i}, z) \in \Theta_{-i} \times \mathcal{Z}(\mathbf{X}_k^n), \quad \tilde{\mu}_i(\{\theta_{-i}\} \times S_{-i}(z)|h) = \mu_i(\{\theta_{-i}\} \times S_{-i}(z)|h). \quad (11)$$

Furthermore, $\text{IH2}(n)$ implies that the histories along those paths, $\mathcal{H}(\mathbf{X}_k^n)$, are also consistent with step $n - 1$ of procedure $k - 1$.

CLAIM 2: $\mathcal{H}(\mathbf{X}_k^n) \subseteq \mathcal{H}(\mathbf{X}_{k-1}^{n-1})$.

In what follows, we will also use the following implication of standard dynamic programming arguments.²⁴

CLAIM 3: *Fix a subset of histories \tilde{H} such that, for every $h \in \tilde{H}$, s_i is a continuation best reply to $\tilde{\mu}_i(\cdot|h)$ for θ_i . There exists $\tilde{s}_i \in r_{i, \theta_i}(\tilde{\mu}_i)$ such that $\tilde{s}_i(h) = s_i(h)$ for every $h \in \tilde{H}$.*

Claim 2 allows to apply $\text{IH1}(n)$ and say that every sequential best reply \tilde{s}_i to $\tilde{\mu}_i$, which survives step n of procedure $k - 1$, has a counterpart \tilde{s}_i' that survives step n of procedure k and mimics \tilde{s}_i at each $h \in \mathcal{H}(\mathbf{X}_k^n) \cap \mathcal{H}_i(\tilde{s}_i') = \mathcal{H}(\mathbf{X}_k^n) \cap \mathcal{H}_i(\tilde{s}_i)$. Now note that, by equation

²⁴We provide such arguments in the Appendix; see Lemma 3.

(11) and the fact that μ_i strongly believes $X_{k,-i}^n$, every strategy s'_i that does not leave the paths induced by profiles in X_k^n yields the same expected payoff under $\tilde{\mu}_i(\cdot|h)$ and $\mu_i(\cdot|h)$ for every $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s'_i)$. Obviously, $s_i, \tilde{s}_i \in \text{proj}_{S_i} X_{k,i}^n$ do not leave those paths, and since \tilde{s}_i mimics \tilde{s}'_i as described above, \tilde{s}_i does not leave those paths either. But then, for each $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i) \cap \mathcal{H}_i(\tilde{s}_i)$, the fact that s_i and \tilde{s}_i are continuation best replies to (respectively) μ_i and $\tilde{\mu}_i$ at h implies that they are also continuation best replies (respectively) to $\tilde{\mu}_i$ and μ_i at h . To extend this claim to every $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$, we need to make sure that h is also reached by some sequential best reply \tilde{s}_i to $\tilde{\mu}_i$; for this, we just need an inductive application of Claim 3, from the initial history and moving downwards.

CLAIM 4: For each $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$, strategy s_i is a continuation best reply to $\tilde{\mu}_i(\cdot|h)$ for θ_i .

By Claim 3 with $\tilde{H} = \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$ and Claim 4, there exists $\tilde{s}_i \in r_{i,\theta_i}(\tilde{\mu}_i)$ such that $\tilde{s}_i(h) = s_i(h)$ for all $h \in \mathcal{H}(X_k^n)$. (For each $h \in \mathcal{H}(X_k^n) \setminus \mathcal{H}_i(s_i)$, since $h \notin \mathcal{H}_i(\tilde{s}_i)$, we can always set $\tilde{s}_i(h) = s_i(h)$ because we use the weak notion of sequential best reply which only refers to histories consistent with the candidate strategy.) From equation (2), it follows that $\{\theta_i\} \times r_{i,\theta_i}(\tilde{\mu}_i) \subseteq X_{k-1,i}^n$. Thus, $\tilde{s}_i \in X_{k-1,i}^n$.

5. BAYESIAN GAMES

In the game with payoff uncertainty Γ , players' types θ parameterize the payoff functions to express incomplete and asymmetric information about them. Yet, the previous analysis does not prevent the parameters from containing payoff-irrelevant components; that is, the analysis remains valid if, for some player i and some types $\theta'_i \neq \theta''_i$, we have $u_j(\theta'_i, \theta_{-i}, z) = u_j(\theta''_i, \theta_{-i}, z)$ for all $j \in I, \theta_{-i} \in \Theta_{-i}$, and $z \in Z$. However, we want to introduce such payoff-irrelevant components explicitly, in the following way. An **elaboration**²⁵ of $\Gamma = \langle I, (\Theta_i, A_i, \mathcal{A}_i(\cdot), u_i)_{i \in I} \rangle$ is a structure

$$\Gamma^e = \langle I, (T_i, A_i, \mathcal{A}_i(\cdot), u_i^e)_{i \in I} \rangle$$

such that, for every player $i \in I, T_i = \Theta_i \times E_i$, where E_i is a finite nonempty set, $u_i^e : (\times_{j \in I} T_j) \times Z \rightarrow \mathbb{R}$, and

$$u_i^e((\theta_j, e_j)_{j \in I}, z) = u_i((\theta_j)_{j \in I}, z)$$

for all $(\theta_j, e_j)_{j \in I} \in \times_{j \in I} T_j$ and $z \in Z$. In words, each type $t_i = (\theta_i, e_i)$ is made of the payoff-relevant component θ_i and of a payoff-irrelevant component e_i .

We are going to use the new types $(T_i)_{i \in I}$ as parts of a type structure *à la* Harsanyi (1967–1968). Hence, we assign to each type t_i a probability measure $\beta_i(t_i)$ over the coplayers' types T_{-i} , so that t_i is ultimately associated with a hierarchy of beliefs about the payoff-relevant state θ : the first-order belief is the marginal of $\beta_i(t_i)$ over Θ_{-i} ; the second-order belief is the push-forward of $\beta_i(t_i)$ through the maps

$$(\theta_j, t_j)_{j \neq i} \in \Theta_{-i} \times T_{-i} \mapsto (\theta_j, \text{marg}_{\Theta_{-j}} \beta_j(t_j))_{j \neq i} \in (\Theta_j \times \Delta(\Theta_{-j}))_{j \neq i};$$

²⁵The term “elaboration” was introduced by Fudenberg, Kreps, and Levine (1988) with a related, but different meaning: They added payoff types to define incomplete-information perturbations, whereas we add a payoff-irrelevant component to existing payoff types.

and so forth. A **Bayesian elaboration** of $\Gamma = \langle I, (\Theta_i, A_i, \mathcal{A}_i(\cdot), u_i)_{i \in I} \rangle$ is obtained from adding the profile of belief maps $(\beta_i : T_i \rightarrow \Delta(T_{-i}))_{i \in I}$ to an elaboration:

$$\Gamma^b = \langle I, (T_i, A_i, \mathcal{A}_i(\cdot), u_i^b, \beta_i)_{i \in I} \rangle,$$

where $u_i^b = u_i^c$ for each $i \in I$. Note that an elaboration is essentially the same as the original game with payoff uncertainty when each set E_i is a singleton $\{\bar{e}_i\}$, so that Θ and T are isomorphic (in an obvious sense). In this particular case, a Bayesian elaboration is also called “simple Bayesian game” and it adds to Γ a particular kind of profile of type-dependent restrictions on exogenous beliefs: recalling that we let $\hat{\Delta}_{i, \theta_i} \subseteq \Delta(\Theta_{-i})$ denote the restricted set of initial marginal beliefs of type θ_i of player i about coplayers’ types, we have that $\hat{\Delta}_{i, \theta_i} = \{\beta_i(\theta_i, \bar{e}_i)\}$ is a singleton for all i and θ_i .

We can define **strong rationalizability for an elaboration** Γ^c as we did for Γ , with each set Θ_i replaced by T_i : for each $i \in I$, $R_i^{c,0} = T_i \times S_i$, and for each $n \in \mathbb{N}$,

$$R_i^{c,n} = \left\{ (t_i, s_i) : \exists \mu_i \in \bigcap_{m=0}^{n-1} \Delta_{sb}^H(R_{-i}^{c,m}), s_i \in r_{i,t_i}^c(\mu_i) \right\}, \quad \text{with}$$

$$r_{i,t_i}^c(\mu_i) = \left\{ \bar{s}_i : \forall h \in \mathcal{H}_i(\bar{s}_i), \bar{s}_i \in \arg \max_{s_i \in S_i(h)} \mathbb{E}_{\mu_i(\cdot|h)}(u_i^c(t_i, \cdot, \zeta(s_i, \cdot))) \right\}$$

for every CPS $\mu_i \in \Delta^H(T_{-i} \times S_{-i})$. Of course, by taking the sections of these sets at any given type, we obtain the strongly n -rationalizable strategies for that type:

$$S_i^{c,n}(t_i) = \{s_i \in S_i : (t_i, s_i) \in R_i^{c,n}\}.$$

The following lemma formalizes the idea that the payoff-irrelevant component of types does not affect strong rationalizability.

LEMMA 2: *Fix any elaboration Γ^c of Γ . For all $i \in I$, $n \in \mathbb{N}_0$, and $(\theta_i, e_i) \in T_i$, $S_i^{c,n}(\theta_i, e_i) = S_i^n(\theta_i)$.*

Now require that the belief system (CPS) μ_i that justifies a pair (t_i, s_i) be consistent with $\beta_i(t_i)$ at the outset. In this way, we define **strong rationalizability for a Bayesian elaboration** Γ^b .²⁶ for each $i \in I$, $R_i^{b,0} = T_i \times S_i$, and for each $n \in \mathbb{N}$,

$$R_i^{b,n} = \left\{ (t_i, s_i) : \exists \mu_i \in \bigcap_{m=0}^{n-1} \Delta_{sb}^H(R_{-i}^{b,m}), \text{marg}_{T_{-i}} \mu_i(\cdot|\emptyset) = \beta_i(t_i), s_i \in r_{i,t_i}^b(\mu_i) \right\},$$

where $r_{i,t_i}^b(\mu_i) = r_{i,t_i}^c(\mu_i)$ (defined above) for each $\mu_i \in \Delta^H(T_{-i} \times S_{-i})$, because $u_i^b = u_i^c$. The set of strongly n -rationalizable strategies for type t_i in Γ^b is the section

$$S_i^{b,n}(t_i) = \{s_i \in S_i : (t_i, s_i) \in R_i^{b,n}\}.$$

Strong rationalizability for a Bayesian elaboration is tightly related to strong directed rationalizability for the original game with payoff uncertainty. The equivalence is obvious

²⁶In static (one-stage) Bayesian games, this solution concept is equivalent to interim correlated rationalizability (Dekel, Fudenberg, and Morris (2007)).

for a simple Bayesian game, where each T_i is isomorphic to Θ_i (thus set $T_i = \Theta_i$), and for each θ_i , $\beta_i(\theta_i)$ can be taken as the unique initial belief allowed by $\hat{\Delta}_{i,\theta_i}$. Hence, a corollary of Theorem 1 is that *for every $\theta \in \Theta$, the (nonempty) set of strongly rationalizable paths of any (finite) simple Bayesian game based on a given (finite) multistage game with payoff uncertainty is included in the set of strongly rationalizable paths of the latter*. For a nonsimple Bayesian elaboration Γ^b of Γ , one can perform an analogous exercise after defining an ancillary game with payoff uncertainty $\hat{\Gamma}$ with type sets $\hat{\Theta}_i = T_i$ in place of Θ_i for all $i \in I$. With this, strong rationalizability in Γ^b coincides with strong Δ -rationalizability in $\hat{\Gamma}$ with $\hat{\Delta}_{i,t_i} = \{\beta_i(t_i)\}$ for all $i \in I$ and $t_i \in \hat{\Theta}_i$; strong rationalizability in Γ coincides with strong rationalizability in $\hat{\Gamma}$ because $\hat{\Gamma}$ is an elaboration of Γ , and thus Lemma 2 applies; the two things combined, via Theorem 1, yield the following result (the proof is omitted).

THEOREM 2: *Fix any Bayesian elaboration Γ^b of Γ . Then, for every $n > 0$, for each $(\theta, e) \in T$, $\emptyset \neq \zeta(S^{b,n}(\theta, e)) \subseteq \zeta(S^n(\theta))$, that is, for each $(\theta, e, s) \in R^{b,\infty} \neq \emptyset$, there exists $s' \in S$ such that $(\theta, s') \in R^\infty$ and $\zeta(s) = \zeta(s')$.*

6. ROBUST IMPLEMENTATION

We consider a classical mechanism design setting, which we formalize as follows. Fix an **economic environment**

$$\mathcal{E} = \langle I, Y, (\Theta_i, v_i)_{i \in I} \rangle,$$

where Y —a subset of a Euclidean space—is an outcome space and each $v_i : \Theta \times Y \rightarrow \mathbb{R}$ is a parameterized utility function. A special case of interest for the outcome space is a space of lotteries: $Y = \Delta(X)$, where X is a finite set of deterministic outcomes. In this case, $v_i(\theta, y)$ has to be interpreted as the vNM expected utility of lottery y given payoff state θ . The economic environment collects the outcomes that the designer can assign to players and their preferences for such outcomes. A **multistage mechanism** (with observed actions) is a game form

$$\mathcal{M} = \langle I, \bar{H}, g \rangle,$$

where $g : Z \rightarrow Y$ is an outcome function defined on the set of terminal histories determined by the game tree \bar{H} . Thus, the mechanism specifies the rules of the game that determine the outcome. A pair $(\mathcal{E}, \mathcal{M})$ yields a game with payoff uncertainty

$$\Gamma(\mathcal{E}, \mathcal{M}) = \langle I, \bar{H}, (\Theta_i, (u_{i,\theta} = v_{i,\theta} \circ g)_{\theta \in \Theta})_{i \in I} \rangle,$$

which contains both the rules of the game and the payoffs associated with the terminal histories: $u_{i,\theta}(z) = v_{i,\theta}(g(z))$ for all $\theta \in \Theta$ and $z \in Z$. Finally, we introduce a **social choice function** $f : \Theta \rightarrow Y$, representing the outcome the designer would want to realize as a function of players' types.

We are interested in the possibility of *implementing*, or at least *virtually implementing*, the social choice function; that is, we look for a mechanism where players of any types θ will always reach a terminal history z so that $g(z) = f(\theta)$, or at least $g(z) \approx f(\theta)$ in a sense to be made precise. Of course, the θ -dependent predicted path depends on the adopted solution concept. Following Müller (2016), we adopt strong rationalizability and we focus on virtual implementation (v-implementation). Everything in the analysis is also valid for “exact” implementation.

DEFINITION 1: Social choice function f is **v-implementable under strong rationalizability** (in environment \mathcal{E}) if, for every $\varepsilon > 0$, there exists a multistage mechanism \mathcal{M} such that in game with payoff uncertainty $\Gamma(\mathcal{E}, \mathcal{M})$, for every $\theta \in \Theta$ and $s \in S^\infty(\theta) \neq \emptyset$, $\|g(\zeta(s)) - f(\theta)\| < \varepsilon$.²⁷

Bergemann and Morris (2009) introduce the notion of *robust implementation*, which requires the mechanism to implement the social choice function for any exogenous restrictions on players’ collectively coherent hierarchies of beliefs about types, such as the existence of a common prior. As anticipated in the Introduction, in a static setting, one can show that implementation under rationalizability for static games with payoff uncertainty is robust, since—by monotonicity of probability-1 belief—introducing a type structure that restricts players’ belief hierarchies can only reduce the set of their rationalizable strategies. As shown in Example 1, this is not true for strong rationalizability in sequential games, due to the nonmonotonicity of strong belief. For this reason, it was an open question whether Müller’s (2016) notion of implementation is robust in the sense of Bergemann and Morris (2009).

DEFINITION 2: Social choice function $f : \Theta \rightarrow Y$ is **robustly v-implementable under strong rationalizability** (in environment \mathcal{E}) if, for every $\varepsilon > 0$, there exists a multistage mechanism \mathcal{M} such that, in every Bayesian elaboration $\Gamma^b(\mathcal{E}, \mathcal{M})$ of the game with payoff uncertainty $\Gamma(\mathcal{E}, \mathcal{M})$, for all $t = (\theta, e) \in T$ and $s \in S^{b,\infty}(t) \neq \emptyset$, $\|g(\zeta(s)) - f(\theta)\| < \varepsilon$.

In light of Theorem 2, we can give a positive answer to the open question.

COROLLARY 1: Fix a finite economic environment \mathcal{E} and an SCF $f : \Theta \rightarrow Y$. If f is v-implementable under strong rationalizability, then f is also robustly v-implementable under strong rationalizability.

PROOF: Suppose that f is v-implementable under strong rationalizability and let \mathcal{M} be a mechanism such that, in game with payoff uncertainty $\Gamma(\mathcal{E}, \mathcal{M})$, for all $\theta \in \Theta$ and $s \in S^\infty(\theta)$, $\|g(\zeta(s)) - f(\theta)\| < \varepsilon$. Take any Bayesian elaboration $\Gamma^b(\mathcal{E}, \mathcal{M})$ of $\Gamma(\mathcal{E}, \mathcal{M})$. By Theorem 2, for all $(\theta, e) \in \Theta \times E = T$ and $s \in S^{b,\infty}(\theta, e)$, $\emptyset \neq \zeta(S^{b,\infty}(\theta, e)) \subseteq \zeta(S^\infty(\theta))$. It follows that, for all $t = (\theta, e) \in T$ and $s \in S^{b,\infty}(t) \neq \emptyset$, $\|g(\zeta(s)) - f(\theta)\| < \varepsilon$. Q.E.D.

7. DISCUSSION AND EXTENSIONS

In this section, we consider some limitations of our analysis and we discuss possible extensions and related conceptual issues.

Imperfectly Observed Actions. Our results extend to finite sequential games with imperfectly observed actions, as long as perfect recall holds, letting nonterminal histories $h \in H$ be replaced with information sets $h_i \in H_i$ for each player i . Indeed, perfect recall allows to preserve the key elements of our analysis: dynamic consistency of subjective expected utility maximization, and the factorization of the sets of strategy profiles consistent with any given information set h_i as $S(h_i) = S_i(h_i) \times S_{-i}(h_i)$. However, from the perspective of mechanism design, perfect recall as defined in traditional game theory is a *hybrid*

²⁷In the definition, we require that $S^\infty(\theta) \neq \emptyset$ so that the “for all...” condition does not hold vacuously. In fact, we know from Lemma 1 that $S^\infty(\theta) \neq \emptyset$ for all $\theta \in \Theta$.

| Symbol | Terminology |
|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| $(\mu_i(\cdot h))_{h \in H} \in \Delta^H(\Theta_{-i} \times S_{-i})$ | conditional probability systems (CPSs) |
| $\Delta_{i,\theta_i} \subseteq \Delta(\Theta_{-i})$ | restricted set of exogenous beliefs for i of type θ_i |
| $\Delta_{i,\theta_i} \subseteq \Delta^H(\Theta_{-i} \times S_{-i})$ | restricted set of CPSs for i of type θ_i |
| $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ | profile of (transparent) belief restrictions |
| $\mathcal{H}_i(s_i) = \{h \in H : s_i \in S_i(h)\}$ | set of nonterminal histories consistent with s_i |
| $\mu_i \mapsto r_{i,\theta_i}(\mu_i)$ | (weak) sequential best-reply correspondence for θ_i |
| $\Delta_{i,\text{sb}}(E_{-i}), E_{-i} \subseteq \Theta_{-i} \times S_{-i}$ | set of CPSs of i that strongly believe E_{-i} |
| $R_i^{\Delta,n} \subseteq \Theta_i \times S_i$ | set of strongly Δ - n -rationalizable pairs of i |
| $S_i^{\Delta,n}(\theta_i) = \{s_i \in S_i : (\theta_i, s_i) \in R_i^{\Delta,n}\}$ | set of strongly Δ - n -rationalizable strategies for θ_i |
| $\mathcal{Z}(X) = \{z : \exists(\theta, s) \in X, z = \zeta(s)\}$ | set of paths allowed by elements of $X \subseteq \Theta \times S$ |
| $\mathcal{H}(X) = \{h : \exists(\theta, s) \in X, h < \zeta(s)\}$ | set of nonterminal histories allowed by $X \subseteq \Theta \times S$ |
| $t_i = (\theta_i, e_i) \in T_i = \Theta_i \times E_i$ | types à la Harsanyi |

property of information partitions that should be “unpacked,” separating the information reaching players as per the rules specified by the mechanism from the mnemonic abilities of the agents playing the game, which are personal traits just like their preferences. As shown in Battigalli and Generoso (2024), such separation is both possible and conceptually useful: information partitions can be *derived* from primitive elements describing the rules of the game on the one hand, and mnemonic abilities on the other hand. Perfect recall of information partitions obtains if either (1) the relevant agents have perfect memory, a personal feature, or (2) the game rules are such that moving players are always reminded of the signals that previously reached them and the actions they took. In both cases, information sets h_i correspond to *personal histories* of signals received and actions taken by i . By (2), perfect recall can be “enforced” by the designer.

Infinite Type Sets and Infinite Horizon. Although we consider finite multistage games with incomplete information, the analysis of strong rationalizability in Battigalli (2003) allows for a continuum of types and an infinite horizon, provided that some regularity assumptions hold, for example, that type sets are compact metric spaces, feasible actions sets $\mathcal{A}_i(h)$ are finite for all $h \in H$ and $i \in I$, and payoff functions are continuous in the obvious product topology. However, the proof of Theorem 1 relies on the fact that, in a finite game, the procedure of elimination of type-strategy pairs (θ_i, s_i) ends after finitely many steps. Nonetheless, we conjecture that our results can be extended to games with infinite type sets and an infinite horizon that satisfy the aforementioned regularity properties. Intuitively, for each step n , IH1 can be used to show the path inclusion between strong Δ - n -rationalizability and strong n -rationalizability; a limit argument can then be applied by compactness and continuity.

APPENDIX

This section contains ancillary results and the proofs omitted from the main body of the paper (with the exception of the detailed proof of inductive step IH1 in the proof of Theorem 1, which is contained in the Supplemental Appendix). We also include here a table summarizing elements of the analysis introduced in Sections 3, 4, and 5.

A.1. Dynamic Programming and Forward Consistency

We use the following dynamic programming result. First, recall from Section 3.2 that a strategy \bar{s}_i is a continuation best reply (from h) to conditional belief $\mu_i(\cdot|h) \in \Delta(\Theta_{-i} \times$

$S_{-i}(h)$) for type θ_i if, for every $s_i \in S_i(h)$,

$$\mathbb{E}_{\mu_i(\cdot|h)}(U_i(\theta_i, \bar{s}_i, \cdot)) \geq \mathbb{E}_{\mu_i(\cdot|h)}(U_i(\theta_i, s_i, \cdot)).$$

LEMMA 3: Fix a CPS μ_i , a type θ_i , and a strategy s_i . If, for every $h \in \mathcal{H}_i(s_i)$, there exists a continuation best reply $s'_i \in S_i(h)$ to $\mu_i(\cdot|h)$ for θ_i such that $s'_i(h) = s_i(h)$, then s_i is a sequential best reply to μ_i for θ_i , that is, $s_i \in r_{i,\theta_i}(\mu_i)$.

PROOF: We prove this result by contraposition. Suppose that $s_i \notin r_{i,\theta_i}(\mu_i)$. We need to show that there is some $\bar{h} \in \mathcal{H}_i(s_i)$ such that, for every $s'_i \in S_i(\bar{h})$, if $s'_i(\bar{h}) = s_i(\bar{h})$, then s'_i is not a continuation best reply to $\mu_i(\cdot|\bar{h})$ for θ_i . Let $\mathcal{H}_i^D(s_i, \mu_i)$ denote the nonempty set of histories $h \in \mathcal{H}_i(s_i)$ such that s_i is not a continuation best reply to $\mu_i(\cdot|h)$. Since the game is finite, $\mathcal{H}_i^D(s_i, \mu_i)$ has at least one maximal element \bar{h} , that is, $\bar{h} \in \mathcal{H}_i^D(s_i, \mu_i)$ is not a strict prefix of any other $h \in \mathcal{H}_i^D(s_i, \mu_i)$. Since $\bar{h} \in \mathcal{H}_i^D(s_i, \mu_i)$, there is some $\bar{s}_i \in S_i(\bar{h})$ such that

$$\mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, \bar{s}_i, \cdot)) > \mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, s_i, \cdot)). \quad (12)$$

Pick any $s'_i \in S_i(\bar{h})$ such that $s'_i(\bar{h}) = s_i(\bar{h})$ (this includes $s'_i = s_i$). To take care of the possibility that $(\bar{h}, (s_i(\bar{h}), a_{-i})) \in Z$ for some a_{-i} and to ease notation, for all z such that $\mu_i(\Theta_{-i} \times S_{-i}(z)|\bar{h}) > 0$ and all $(\theta_{-i}, s_{-i}) \in \Theta_{-i} \times S_{-i}(z)$, write

$$\mu_i(\theta_{-i}, s_{-i}|z) = \frac{\mu_i(\theta_{-i}, s_{-i}|\bar{h})}{\mu_i(\Theta_{-i} \times S_{-i}(z)|\bar{h})},$$

$$\mathbb{E}_{\mu_i(\cdot|z)}(U_i(\theta_i, s'_i, \cdot)) = \sum_{\theta_{-i} \in \Theta_{-i}} \mu_i(\{\theta_{-i}\} \times S_{-i}(z)|z) u_i(\theta_i, \theta_{-i}, z).$$

With this, letting $\bar{A}_{-i} = \{a_{-i} : \mu_i(\Theta_{-i} \times S_{-i}(\bar{h}, a_{-i})|\bar{h}) > 0\}$, the following decomposition holds:

$$\mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, s'_i, \cdot)) = \sum_{a_{-i} \in \bar{A}_{-i}} \mu_i(\Theta_{-i} \times S_{-i}(\bar{h}, a_{-i})|\bar{h}) \mathbb{E}_{\mu_i(\cdot|(\bar{h}, (s_i(\bar{h}), a_{-i})))}(U_i(\theta_i, s'_i, \cdot)).$$

By maximality of \bar{h} in $\mathcal{H}_i^D(s_i, \mu_i)$, s_i is a continuation best reply to each $\mu_i(\cdot|h)$ with $h = (\bar{h}, (s_i(\bar{h}), a_{-i})) \in H$. Since such h is the immediate follower of \bar{h} selected by action profile $(s_i(\bar{h}), a_{-i})$, it holds that $s_i \in S_i(h)$. Since $s'_i(\bar{h}) = s_i(\bar{h})$, $s'_i \in S_i(h)$ as well. Thus,

$$\mathbb{E}_{\mu_i(\cdot|(\bar{h}, (s_i(\bar{h}), a_{-i})))}(U_i(\theta_i, s_i, \cdot)) \geq \mathbb{E}_{\mu_i(\cdot|(\bar{h}, (s_i(\bar{h}), a_{-i})))}(U_i(\theta_i, s'_i, \cdot))$$

for all $a_{-i} \in \bar{A}_{-i}$ (the other action profiles in $\mathcal{A}_{-i}(\bar{h})$ do not affect expected payoff calculations). It follows that

$$\mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, s_i, \cdot)) \geq \mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, s'_i, \cdot)). \quad (13)$$

Equations (12) and (13) combined yield

$$\mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, \bar{s}_i, \cdot)) > \mathbb{E}_{\mu_i(\cdot|\bar{h})}(U_i(\theta_i, s'_i, \cdot)),$$

so s'_i is not a continuation best reply to $\mu_i(\cdot|\bar{h})$.

Q.E.D.

The omitted parts of the proof of Theorem 1 require to construct CPSs that strongly believe some key events. It turns out that it is simpler to construct a “forward-consistent belief system” (Battigalli, Catonini, and Manili (2023)) with such features and then claim the existence of a CPS that preserves them. A **forward-consistent belief system** is an array of beliefs $\hat{\mu}_i = (\hat{\mu}_i(\cdot|h))_{h \in H} \in (\Delta(\Theta_{-i} \times S_{-i}))^H$ such that, for every $h \in H$, $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h)|h) = 1$ and the *forward chain rule* holds: for all $h, h' \in H$ and $E_{-i} \subseteq \Theta_{-i} \times S_{-i}(h')$,

$$h \preceq h' \implies \hat{\mu}_i(E_{-i}|h) = \hat{\mu}_i(E_{-i}|h')\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h).$$

The forward chain rule is weaker than the chain rule, because as noted in Section 3.2, $S_{-i}(h') \subseteq S_{-i}(h)$ does not imply $h \preceq h'$. The definition of “strong belief” for a forward-consistent belief system is the same as for a CPS: belief system $\hat{\mu}_i$ **strongly believes** E_{-i} if

$$\forall h \in H, \quad E_{-i} \cap (\Theta_{-i} \times S_{-i}(h)) \neq \emptyset \implies \hat{\mu}_i(E_{-i}|h) = 1.$$

For the transformation of forward-consistent belief systems into CPSs, we rely on the following result.

LEMMA 4—Battigalli, Catonini, and Manili (2023): *Fix a strategy s_i and a forward-consistent belief system $\hat{\mu}_i$ that strongly believes $E_{-i}^1, \dots, E_{-i}^{n-1}$, where $E_{-i}^{n-1} \subseteq \dots \subseteq E_{-i}^1$. Then there is a CPS $\tilde{\mu}_i$ that strongly believes $E_{-i}^1, \dots, E_{-i}^{n-1}$ such that $\tilde{\mu}_i(\cdot|h) = \hat{\mu}_i(\cdot|h)$ for all $h \in \mathcal{H}_i(s_i)$.*

A.2. Omitted Parts of the Proof of Theorem 1

A.2.1. Proof of Claim 1

We construct an array of beliefs $\hat{\mu}_i = (\hat{\mu}_i(\cdot|h))_{h \in H}$ such that, for each $h \in H$:

F0. $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h)|h) = 1$;

F1. for all h' such that $h < h'$,

$$\forall E \subseteq \Theta_{-i} \times S_{-i}(h'), \quad \hat{\mu}_i(E|h')\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) = \hat{\mu}_i(E|h); \quad (14)$$

F2. for all $m = 0, \dots, n-1$, if $h \in \mathcal{H}(X_{k-1,-i}^m)$, then $\hat{\mu}_i(X_{k-1,-i}^m|h) = 1$;

F3. $\text{marg}_{\Theta_{-i}}\hat{\mu}_i(\cdot|\emptyset) = \text{marg}_{\Theta_{-i}}\mu_i(\cdot|\emptyset)$;

F4. if $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$,

$$\forall (\theta_{-i}, z) \in \Theta_{-i} \times \mathcal{Z}(X_k^n), \quad \hat{\mu}_i(\{\theta_{-i}\} \times S_{-i}(z)|h) = \mu_i(\{\theta_{-i}\} \times S_{-i}(z)|h). \quad (15)$$

By F0 and F1, $\hat{\mu}_i$ is a forward-consistent belief system. By F2, it strongly believes $X_{k-1,-i}^1, \dots, X_{k-1,-i}^{n-1}$. Hence, by Lemma 4, there exists a CPS $\tilde{\mu}_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(X_{k-1,-i}^m)$ such that $\tilde{\mu}_i(\cdot|h) = \hat{\mu}_i(\cdot|h)$ for all $h \in \mathcal{H}_i(s_i)$. By $\tilde{\mu}_i(\cdot|\emptyset) = \hat{\mu}_i(\cdot|\emptyset)$, F3, and $\mu_i \in \Delta_{i,\theta_i}$, we get $\tilde{\mu}_i \in \Delta_{i,\theta_i}$. Finally, for every $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$, $\tilde{\mu}_i(\cdot|h) = \hat{\mu}_i(\cdot|h)$, and F4 yield (11).

Now we start with the construction. By IH2(n), for every $(\theta_{-i}, s_{-i}) \in X_{k,-i}^n$, there exists a profile $(\tilde{s}_j^{(\theta_j, s_j)})_{j \neq i} \in S_{-i}$ such that $(\theta_j, \tilde{s}_j^{(\theta_j, s_j)})_{j \neq i} \in X_{k-1,-i}^{n-1}$ and, for each $j \neq i$, $\tilde{s}_j^{(\theta_j, s_j)}(h) = s_j(h)$ for all $h \in \mathcal{H}(X_k^{n-1})$. With this, define a map $\tilde{\eta} : \Theta_{-i} \times S_{-i} \rightarrow \Theta_{-i} \times S_{-i}$ as follows:

$$\forall (\theta_{-i}, s_{-i}) \in (\Theta_{-i} \times S_{-i}), \quad \tilde{\eta}(\theta_{-i}, s_{-i}) = \begin{cases} (\theta_j, \tilde{s}_j^{(\theta_j, s_j)})_{j \neq i} & \text{if } (\theta_{-i}, s_{-i}) \in X_{k,-i}^n, \\ (\theta_{-i}, s_{-i}) & \text{otherwise.} \end{cases}$$

For each $h \in \mathcal{H}(X_k^n)$, define $\hat{\mu}_i(\cdot|h)$ as the $\tilde{\eta}$ -push-forward (image measure) of $\mu_i(\cdot|h)$. For future reference, observe that

$$\hat{\mu}_i(X_{k-1,-i}^{n-1}|h) = \mu_i(\tilde{\eta}^{-1}(X_{k-1,-i}^{n-1})|h) \geq \mu_i(X_{k,-i}^n|h) = 1, \tag{16}$$

where the first equality holds by construction, the inequality holds by $\tilde{\eta}(X_{k,-i}^n) \subseteq X_{k-1,-i}^{n-1}$, and the last equality holds by strong belief in $X_{k,-i}^n$. Now define

$$\tilde{H} = \{h \in H \setminus \mathcal{H}(X_k^n) : \exists \bar{h} \in \mathcal{H}(X_k^n), \bar{h} < h, \hat{\mu}_i(\Theta_{-i} \times S_{-i}(h)|\bar{h}) > 0\}.$$

For each $h \in \tilde{H}$, let $p^*(h)$ denote the longest $\bar{h} < h$ with $\bar{h} \in \mathcal{H}(X_k^n)$ such that $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h)|\bar{h}) > 0$, and derive $\hat{\mu}_i(\cdot|h)$ by conditioning $\hat{\mu}_i(\cdot|p^*(h))$. To conclude the construction, fix $\bar{\mu}_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(X_{k-1,-i}^m)$, and for each $h \in H \setminus (\mathcal{H}(X_k^n) \cup \tilde{H}) =: \hat{H}$, let $\hat{\mu}_i(\cdot|h) = \bar{\mu}_i(\cdot|h)$.

First, we show that $\hat{\mu}_i$ satisfies F2. For each $h \in \mathcal{H}(X_k^n)$, equation (16) yields $\hat{\mu}_i(X_{k-1,-i}^{n-1}|h) = 1$. For each $h \in \tilde{H}$, equation (16) yields $\hat{\mu}_i(X_{k-1,-i}^{n-1}|p^*(h)) = 1$, from which $\hat{\mu}_i(X_{k-1,-i}^{n-1}|h) = 1$ follows by construction. For each $h \in \hat{H}$ and $m = 0, \dots, n-1$, if $h \in \mathcal{H}(X_{k-1,-i}^m)$, $\hat{\mu}_i(X_{k-1,-i}^m|h) = 1$ follows from $\hat{\mu}_i(\cdot|h) = \bar{\mu}_i(\cdot|h)$ and $\bar{\mu}_i \in \Delta_{\text{sb}}^H(X_{k-1,-i}^m)$.

Next, we show that, for every $h \in \mathcal{H}(X_k^n)$ and $(\theta_{-i}, h') \in \Theta_{-i} \times (\mathcal{H}(X_k^n) \cup \mathcal{Z}(X_k^n))$,

$$\hat{\mu}_i(\{\theta_{-i}\} \times S_{-i}(h')|h) = \mu_i(\{\theta_{-i}\} \times S_{-i}(h')|h), \tag{17}$$

which yields: condition (15) when $h' \in \mathcal{Z}(X_k^n)$, thus F4; F3 when h and h' coincide with the initial history; and, for future reference,

$$\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) = \mu_i(\Theta_{-i} \times S_{-i}(h')|h). \tag{18}$$

By construction, we have

$$\hat{\mu}_i(\{\theta_{-i}\} \times S_{-i}(h')|h) = \mu_i(\tilde{\eta}^{-1}(\{\theta_{-i}\} \times S_{-i}(h'))|h).$$

We need to show that

$$\tilde{\eta}^{-1}(\{\theta_{-i}\} \times S_{-i}(h')) = \{\theta_{-i}\} \times S_{-i}(h'). \tag{19}$$

Fix first $s_{-i} \in S_{-i}$ such that $(\theta_{-i}, s_{-i}) \in \tilde{\eta}^{-1}(\{\theta_{-i}\} \times S_{-i}(h'))$. Then there exists $s'_{-i} \in S_{-i}(h')$ such that $\tilde{\eta}(\theta_{-i}, s_{-i}) = (\theta_{-i}, s'_{-i})$. By definition of $\tilde{\eta}$, either $s'_{-i} = s_{-i}$, or $s_{-i}(\tilde{h}) = s'_{-i}(\tilde{h})$ for each $\tilde{h} \in \mathcal{H}(X_k^{n-1})$, so in particular for each $\tilde{h} < h'$, given that $h' \in \mathcal{H}(X_k^n) \cup \mathcal{Z}(X_k^n)$. Hence, $s'_{-i} \in S_{-i}(h')$ implies $s_{-i} \in S_{-i}(h')$, that is, $(\theta_{-i}, s_{-i}) \in \{\theta_{-i}\} \times S_{-i}(h')$. Now fix $s_{-i} \in S_{-i}(h')$. Let $(\theta_{-i}, s'_{-i}) = \tilde{\eta}(\theta_{-i}, s_{-i})$. By definition of $\tilde{\eta}$, either $s'_{-i} = s_{-i}$, or $s'_{-i}(\tilde{h}) = s_{-i}(\tilde{h})$ for each $\tilde{h} \in \mathcal{H}(X_k^{n-1})$, so in particular for each $\tilde{h} < h'$, given that $h' \in \mathcal{H}(X_k^n) \cup \mathcal{Z}(X_k^n)$. Hence, $s_{-i} \in S_{-i}(h')$ implies $s'_{-i} \in S_{-i}(h')$, which means $(\theta_{-i}, s_{-i}) \in \tilde{\eta}^{-1}(\{\theta_{-i}\} \times S_{-i}(h'))$.

Finally, we show that $\hat{\mu}_i$ satisfies F0 and F1. For each $h \in \mathcal{H}(X_k^n)$, since $\mu_i(\Theta_{-i} \times S_{-i}(h)|h) = 1$, equation (18) with $h' = h$ yields F0. For each $h \in \tilde{H}$, F0 follows by conditioning. For each $h \in \hat{H}$, F0 holds by $\hat{\mu}_i(\cdot|h) = \bar{\mu}_i(\cdot|h)$.

For F1, equation (14) holds if $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) = 0$, because then $\hat{\mu}_i(E|h) = 0$, so suppose that $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) > 0$.

Case 1: $h \in \hat{H}$. Then $h' \in \hat{H}$, too. Hence, $\hat{\mu}_i(\cdot|h) = \bar{\mu}_i(\cdot|h)$ and $\hat{\mu}_i(\cdot|h') = \bar{\mu}_i(\cdot|h')$, so $\hat{\mu}_i$ inherits (14) from $\bar{\mu}_i$, which is a CPS.

Case 2: $h \in \tilde{H}$. Then $\hat{\mu}_i(\cdot|h)$ is derived from $\hat{\mu}_i(\cdot|p^*(h))$ by conditioning. By $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) > 0$, we have $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|p^*(h)) > 0$, hence $h' \in \tilde{H}$, too, and $p^*(h) = p^*(h')$. Thus, $\hat{\mu}_i(\cdot|h')$ is derived from $\hat{\mu}_i(\cdot|p^*(h))$, too, and (14) follows.

Case 3: $h \in \mathcal{H}(X_k^n)$. If $h' \in \mathcal{H}(X_k^n)$, let $\bar{h} = h'$, otherwise, by $\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) > 0$, $h' \in \tilde{H}$, and in this case let $\bar{h} = p^*(h')$. Thus, $\bar{h} \in \mathcal{H}(X_k^n)$. For each $E \subseteq \Theta_{-i} \times S_{-i}(\bar{h})$, by construction of $\hat{\mu}_i$ and equation (18), we get

$$\hat{\mu}_i(E|\bar{h})\hat{\mu}_i(\Theta_{-i} \times S_{-i}(\bar{h})|h) = \mu_i(\tilde{\eta}^{-1}(E)|\bar{h})\mu_i(\Theta_{-i} \times S_{-i}(\bar{h})|h).$$

Equation (19) implies that $\tilde{\eta}^{-1}(E) \subseteq \Theta_{-i} \times S_{-i}(\bar{h})$; so, since μ_i is a CPS, we have

$$\mu_i(\tilde{\eta}^{-1}(E)|\bar{h})\mu_i(\Theta_{-i} \times S_{-i}(\bar{h})|h) = \mu_i(\tilde{\eta}^{-1}(E)|h),$$

and $\mu_i(\tilde{\eta}^{-1}(E)|h) = \hat{\mu}_i(E|h)$ by construction of $\hat{\mu}_i$. So,

$$\hat{\mu}_i(E|\bar{h})\hat{\mu}_i(\Theta_{-i} \times S_{-i}(\bar{h})|h) = \hat{\mu}_i(E|h). \tag{20}$$

If $\bar{h} = h'$, we are done. Otherwise, for each $E' \subseteq \Theta_{-i} \times S_{-i}(h')$, we have

$$\begin{aligned} \hat{\mu}_i(E'|h')\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) &= \frac{\hat{\mu}_i(E'|p^*(h'))}{\hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|p^*(h'))} \hat{\mu}_i(\Theta_{-i} \times S_{-i}(h')|h) \\ &= \hat{\mu}_i(E'|p^*(h'))\hat{\mu}_i(\Theta_{-i} \times S_{-i}(p^*(h'))|h) \\ &= \hat{\mu}_i(E'|h), \end{aligned}$$

where the first equality holds by definition of $\hat{\mu}_i(E'|h')$ and the second and third equalities follow from equation (20) with $\bar{h} = p^*(h')$ and $E = \Theta_{-i} \times S_{-i}(h')$ for the second equality, $E = E'$ for the third.

A.2.2. Proof of Claim 2

Fix $\hat{s} \in \text{proj}_S X_k^n$. By IH2(n), there exists $\hat{s}' \in \text{proj}_S X_{k-1}^{n-1}$ such that $\hat{s}'(\tilde{h}) = \hat{s}(\tilde{h})$ for every $\tilde{h} \in \mathcal{H}(X_k^{n-1}) \supseteq \mathcal{H}(X_k^n)$. It follows that $\zeta(\hat{s}) = \zeta(\hat{s}') \in \mathcal{Z}(X_{k-1}^{n-1})$.

A.2.3. Proof of Claim 3

Construct \tilde{s}_i as follows. For each $h \in \tilde{H}$, let $\tilde{s}_i(h) = s_i(h)$. For each $h \in H \setminus \tilde{H}$, let $\tilde{s}_i(h) = s'_i(h)$ for some continuation best reply s'_i to $\tilde{\mu}_i(\cdot|h)$ for θ_i . It follows from Lemma 3 that $\tilde{s}_i \in r_{i,\theta_i}(\tilde{\mu}_i)$.

A.2.4. Proof of Claim 4

First, note that $\mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$ is closed with respect to prefixes (predecessors): for each $h \in \mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$, each prefix $h' \prec h$ belongs to $\mathcal{H}(X_k^n) \cap \mathcal{H}_i(s_i)$. So, suppose by way of induction that Claim 4 holds for every $h' \prec h$, which is vacuously true if $h = \emptyset$. Then, setting $\tilde{H} = \{h' \in H : h' \prec h\}$, Claim 3 guarantees the existence of some $\tilde{s}_i \in r_{i,\theta_i}(\tilde{\mu}_i)$ such that $\tilde{s}_i(h') = s_i(h')$ for every $h' \prec h$, thus $\tilde{s}_i \in S_i(h)$.

First, we need to show that $\zeta(\tilde{s}_i, \tilde{s}_{-i}) \in \mathcal{Z}(X_k^n)$ for every $(\theta_{-i}, \tilde{s}_{-i}) \in \text{supp} \tilde{\mu}_i(\cdot|h)$. So, fix $(\theta_{-i}, \tilde{s}_{-i}) \in \text{supp} \tilde{\mu}_i(\cdot|h)$. Note that $\{\theta_i\} \times r_{i,\theta_i}(\tilde{\mu}_i) \subseteq X_{k-1,i}^n$, and hence $\tilde{s}_i \in \text{proj}_{S_i} X_{k-1,i}^n$. So,

by IH1(n) there exists $\tilde{s}_i \in \text{proj}_{S_i} X_{k,i}^n$ such that $\tilde{s}'_i(h) = \tilde{s}_i(h)$ for every $h \in \mathcal{H}(X_{k-1}^{n-1})$.²⁸ Fix $(\theta_{-i}, \tilde{s}_{-i}) \in \tilde{\eta}^{-1}((\theta_{-i}, \tilde{s}_{-i})) \subseteq X_{k,-i}^n$ —it exists because $\tilde{\mu}_i(\cdot|h) = \hat{\mu}_i(\cdot|h)$ and $\hat{\mu}_i(\cdot|h)$ is the $\tilde{\eta}$ -push-forward of $\mu_i(\cdot|h)$ (see the proof of Claim 1). By definition of $\tilde{\eta}$, $\zeta(\tilde{s}'_i, \tilde{s}'_{-i}) \in \mathcal{Z}(X_k^n)$. For every $\tilde{h} < \zeta(\tilde{s}'_i, \tilde{s}'_{-i})$, we have $\tilde{h} \in \mathcal{H}(X_k^n) \subseteq \mathcal{H}(X_k^{n-1})$, hence $\tilde{s}_{-i}(\tilde{h}) = \tilde{s}'_{-i}(\tilde{h})$ by definition of $\tilde{\eta}$. Claim 2 gives $\mathcal{H}(X_k^n) \subseteq \mathcal{H}(X_{k-1}^{n-1})$, therefore $\tilde{s}_i(\tilde{h}) = \tilde{s}'_i(\tilde{h})$ as well. It follows that $\zeta(\tilde{s}_i, \tilde{s}_{-i}) = \zeta(\tilde{s}'_i, \tilde{s}'_{-i}) \in \mathcal{Z}(X_k^n)$.

For each $(\theta_{-i}, z) \in \Theta_{-i} \times \mathcal{Z}(X_k^n)$, the probability of (θ_{-i}, z) induced by \tilde{s}_i and $\tilde{\mu}_i(\cdot|h)$ (resp., $\mu_i(\cdot|h)$) is 0, if $\tilde{s}_i \notin S_i(z)$, and it amounts to $\tilde{\mu}_i(\{\theta_{-i}\} \times S_{-i}(z)|h)$ (resp., $\mu_i(\{\theta_{-i}\} \times S_{-i}(z)|h)$) otherwise. Then, by equation (11), \tilde{s}_i induces the same probability over each $(\theta_{-i}, z) \in \Theta_{-i} \times \mathcal{Z}(X_k^n)$ under $\tilde{\mu}_i(\cdot|h)$ and under $\mu_i(\cdot|h)$, hence the same distribution over $\Theta_{-i} \times \mathcal{Z}$, because the probability induced by \tilde{s}_i and $\tilde{\mu}_i(\cdot|h)$ over $\Theta_{-i} \times (\mathcal{Z} \setminus \mathcal{Z}(X_k^n))$ is zero: as we have previously shown, for each $(\theta_{-i}, \tilde{s}_{-i}) \in \text{supp} \tilde{\mu}_i(\cdot|h)$, $\zeta(\tilde{s}_i, \tilde{s}_{-i}) \in \mathcal{Z}(X_k^n)$. The same conclusion can be reached for s_i in the same way, after observing that for each $(\theta_{-i}, s_{-i}) \in \text{supp} \mu_i(\cdot|h)$, since $(\theta_i, s_i, \theta_{-i}, s_{-i}) \in X_k^n$, we have $\zeta(s_i, s_{-i}) \in \mathcal{Z}(X_k^n)$. So, call $\pi^{\tilde{s}_i}$ and π^{s_i} the unique expected payoffs induced by, respectively, (θ_i, \tilde{s}_i) and (θ_i, s_i) under both beliefs $(\mu_i(\cdot|h)$ and $\tilde{\mu}_i(\cdot|h))$. Since \tilde{s}_i and s_i are continuation best replies for θ_i to, respectively, $\tilde{\mu}_i(\cdot|h)$ and $\mu_i(\cdot|h)$, we have $\pi^{\tilde{s}_i} \geq \pi^{s_i}$ and $\pi^{s_i} \geq \pi^{\tilde{s}_i}$. Hence, $\pi^{s_i} = \pi^{\tilde{s}_i}$. But then, also s_i is a continuation best reply for θ_i to $\tilde{\mu}_i(\cdot|h)$.

A.3. Proof of Lemma 2

The statement is trivially true for $n = 0$. Suppose by way of induction that it is true for each $m < n$; fix $i \in I$ and $(\theta_i, e_i) \in T_i = \Theta_i \times E_i$ arbitrarily. Let $\tilde{s}_i \in S_i^n(\theta_i)$. Then there is a CPS $\mu_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(R_{-i}^m)$ such that $\tilde{s}_i \in r_{i,\theta_i}(\mu_i)$. Define $\mu_i^c \in (\Delta(T_{-i} \times S_{-i}))^H$ as follows: for all $h \in H$, $s_{-i} \in S_{-i}(h)$, $(\theta_{-i}, e_{-i}) \in T_{-i}$,

$$\mu_i^c(\theta_{-i}, e_{-i}, s_{-i}|h) = \frac{1}{|E_{-i}|} \mu_i(\theta_{-i}, s_{-i}|h).$$

It can be checked that μ_i^c is a CPS, that is, $\mu_i^c \in \Delta^H(T_{-i} \times S_{-i})$. Furthermore, since $\mu_i(\cdot|h) = \text{marg}_{\Theta_{-i} \times S_{-i}(h)} \mu_i^c(\cdot|h)$ for each $h \in H$, and the e_j -component of the type of each player $j \in I$ is payoff irrelevant, $\tilde{s}_i \in r_{i,(\theta_i, e_i)}(\mu_i^c)$. Finally, the aforementioned marginalization relationship between μ_i and μ_i^c and the inductive hypothesis imply that $\mu_i^c \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(R_{-i}^{c,m})$. Therefore, $\tilde{s}_i \in S_i^{c,n}(\theta_i, e_i)$. Conversely, suppose that $\tilde{s}_i \in S_i^{c,n}(\theta_i, e_i)$. Then there is a CPS $\mu_i^c \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(R_{-i}^{c,m})$ such that $\tilde{s}_i \in r_{i,(\theta_i, e_i)}(\mu_i^c)$. Define $\mu_i \in (\Delta(\Theta_{-i} \times S_{-i}))^H$ as $\mu_i(\cdot|h) = \text{marg}_{\Theta_{-i} \times S_{-i}(h)} \mu_i^c(\cdot|h)$ for each $h \in H$. It can be checked that μ_i is a CPS, that is, $\mu_i \in \Delta^H(\Theta_{-i} \times S_{-i})$. Similar to the previous argument, since the e_j -component of the type of each player $j \in I$ is payoff irrelevant, $\tilde{s}_i \in r_{i,\theta_i}(\mu_i)$. Furthermore, the marginalization relationship between μ_i and μ_i^c and the inductive hypothesis imply that $\mu_i \in \bigcap_{m=0}^{n-1} \Delta_{\text{sb}}^H(R_{-i}^m)$.

REFERENCES

Abreu, Dilip and Hitoshi Matsushima (1992), “Virtual Implementation in Iteratively Undominated Strategies: Complete Information.” *Econometrica*, 60, 993–1008. [0506,0511]
 Artemov, Georgy, Takashi Kunimoto, and Roberto Serrano (2013), “Robust Virtual Implementation. Toward a Reinterpretation of the Wilson Doctrine.” *Journal of Economic Theory*, 148, 424–447. [0511]

²⁸This is the only passage where we use IH1(n) at full power, namely where it is important (to then apply Claim 3) that IH1(n) involves all the histories in $H(X_{k-1}^{n-1})$ and not just those in $H(X_{k-1}^n)$.

- Battigalli, Pierpaolo (1997), “On Rationalizability in Extensive Games.” *Journal of Economic Theory*, 74, 40–61. [0506,0507,0518]
- Battigalli, Pierpaolo (2003), “Rationalizability in Infinite, Dynamic Games of Incomplete Information.” *Research in Economics*, 57, 1–38. [0510,0511,0517,0518,0529]
- Battigalli, Pierpaolo and Emiliano Catonini (2024), “The Epistemic Spirit of Divinity.” *Journal of Economic Theory*, 222, 105903. [0506]
- Battigalli, Pierpaolo and Emiliano Catonini (2026), “Supplement to ‘Monotonicity and Robust Implementation Under Forward-Induction Reasoning.’” *Econometrica Supplemental Material*, 94, <https://doi.org/10.3982/ECTA23709>. [0513]
- Battigalli, Pierpaolo and Amanda Friedenberg (2012), “Forward-Induction Reasoning Revisited.” *Theoretical Economics*, 7, 57–98. [0506,0514,0520]
- Battigalli, Pierpaolo and Nicolo’ Generoso (2024), “Information Flows and Memory in Games.” *Games and Economic Behavior*, 145, 356–376. [0529]
- Battigalli, Pierpaolo and Andrea Prestipino (2013), “Transparent Restrictions on Beliefs and Forward-Induction Reasoning in Games With Asymmetric Information.” *The B.E. Journal of Theoretical Economics (Contributions)*, 13 (1), 1–53. [0517]
- Battigalli, Pierpaolo and Marciano Siniscalchi (2002), “Strong Belief and Forward-Induction Reasoning.” *Journal of Economic Theory*, 106, 356–391. [0505,0506,0510,0512,0518]
- Battigalli, Pierpaolo and Marciano Siniscalchi (2003), “Rationalization and Incomplete Information.” *Advances in Theoretical Economics*, 3 (1), Art. 3. [0506,0510,0511,0517]
- Battigalli, Pierpaolo, Emiliano Catonini, and Nicodemo De Vito (2025), “Game Theory: Analysis of Strategic Thinking.” Typescript. [0514]
- Battigalli, Pierpaolo, Emiliano Catonini, and Julien Manili (2023), “Belief Change, Rationality, and Strategic Reasoning in Sequential Games.” *Games and Economic Behavior*, 142, 527–551. [0507,0516,0531]
- Bergemann, Dirk and Stephen Morris (2009), “Robust Virtual Implementation.” *Theoretical Economics*, 4, 45–88. [0505,0506,0510,0511,0528]
- Bergemann, Dirk and Stephen Morris (2012), “An Introduction to Robust Mechanism Design.” *Foundations and Trends in Microeconomics*, 3, 169–230. [0505,0510]
- Bergemann, Dirk and Stephen Morris (2017), “Belief-Free Rationalizability and Informational Robustness.” *Games and Economic Behavior*, 104, 744–759. [0510]
- Catonini, Emiliano (2019), “Rationalizability and Epistemic Priority Orderings.” *Games and Economic Behavior*, 114, 101–117. [0519,0520]
- Catonini, Emiliano (2020), “On Non-Monotonic Strategic Reasoning.” *Games and Economic Behavior*, 120, 209–224. [0520]
- Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2007), “Interim Correlated Rationalizability.” *Theoretical Economics*, 2, 15–40. [0526]
- Fudenberg, Drew, David Kreps, and David K. Levine (1988), “On the Robustness of Equilibrium Refinements.” *Journal of Economic Theory*, 44, 354–380. [0525]
- Glazer, Jacob and Motty Perry (1996), “Virtual Implementation in Backwards Induction.” *Games and Economic Behavior*, 15, 27–32. [0506]
- Harsanyi, John (1967–1968), “Games of Incomplete Information Played by Bayesian Players. Parts I, II, III.” *Management Science*, 14, 159–182, 320–334, 486–502. [0506,0510,0525]
- Kohlberg, Elon (1990), “Refinement of Nash Equilibrium: The Main Ideas.” In *Game Theory and Applications* (Tatsuro Ichiishi, Abraham Neyman, and Yair Tauman, eds.), 3–45, Academic Press, San Diego. [0506]
- Mertens, Jean-Francois and Shmuel Zamir (1985), “Formulation of Bayesian Analysis for Games With Incomplete Information.” *International Journal of Game Theory*, 14, 1–29. [0506,0510]
- Moore, John and Rafael Repullo (1988), “Subgame Perfect Implementation.” *Econometrica*, 56, 1191–1220. [0506]
- Müller, Christoph (2016), “Robust Virtual Implementation Under Common Strong Belief in Rationality.” *Journal of Economic Theory*, 162, 407–450. [0505,0506,0512,0513,0527,0528]
- Müller, Christoph (2020), “Robust Implementation in Weakly Perfect Bayesian Strategies.” *Journal of Economic Theory*, 189, 105038. [0512]
- Ollár, Mariann and Antonio Penta (2017), “Full Implementation and Belief Restrictions.” *American Economic Review*, 107, 2243–2277. [0511]
- Ollár, Mariann and Antonio Penta (2023), “A Network Solution to Robust Implementation: The Case of Identical but Unknown Distributions.” *Review of Economic Studies*, 90, 2517–2554. [0511]
- Osborne, Martin and Ariel Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge. [0506]
- Pearce, David (1984), “Rationalizable Strategic Behavior and the Problem of Perfection.” *Econometrica*, 52, 1029–1050. [0507,0518]

- Perea, Andres (2018), “Why Forward Induction Leads to the Backward Induction Outcome: A New Proof for Battigalli’s Theorem.” *Games and Economic Behavior*, 110, 120–138. [0520]
- Perea, Andres (2024), “More Reasoning, Less Outcomes: A Monotonicity Result for Reasoning in Dynamic Games.” EPICENTER W.P. 32, Maastricht University. [0520]
- Wilson, Robert (1987), “Game-Theoretic Analyses of Trading Processes.” In *Advances in Economic Theory, Fifth World Congress* (Truman Bewley, ed.), Vol. 1, 33–70, Cambridge University Press, New York. [0506, 0510]
- Ziegler, Gabriel (2022), “Informational Robustness of Common Belief in Rationality.” *Games and Economic Behavior*, 132, 592–597. [0511]

Co-editor Marina Halac handled this manuscript.

Manuscript received 25 November, 2024; final version accepted 15 November, 2025; available online 18 November, 2025.

All authors assume responsibility for all aspects of the paper.