



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games

P. Battigalli, E. Catonini, G. Lanzani, and M. Marinacci

Working Paper n. 607

This Version: September, 2017

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games*

P. Battigalli,^a E. Catonini,^b G. Lanzani,^c and M. Marinacci^a

^aDepartment of Decision Sciences and IGIER - Università Bocconi

^bNational Research University Higher School of Economics - ICEF

^cDepartment of Economics - Massachusetts Institute of Technology

Abstract

We consider a game with sequential moves played by agents who are randomly drawn from large populations and matched. We assume that, when players are uncertain about the strategy distributions of the opponents, preferences over actions at any information set admit a smooth-ambiguity representation in the sense of Klibanoff, Marinacci, and Mukerji (*Econometrica*, 2005). This may induce dynamically inconsistent preferences and calls for an appropriate definition of sequential best response. We take this into account in our analysis of self-confirming equilibrium (SCE) and rationalizable SCE in sequential games with feedback played by agents with non-neutral ambiguity attitudes. Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (*Amer. Econ. Rev.*, 2015) show that the set of SCE's of a simultaneous-move game with feedback expands as ambiguity aversion increases. We show by example that SCE in a sequential game is not equivalent to SCE applied to the strategic form of such game, and that the previous monotonicity result does not extend to general sequential games. Still, we provide sufficient conditions under which the monotonicity result holds for (rationalizable) SCE.

KEYWORDS: Sequential games with feedback, smooth ambiguity, self-confirming equilibrium, rationalizable self-confirming equilibrium.

1 Introduction

When a game is played recurrently, and the learning dynamic has reached a rest point, each agent chooses a best reply to his subjective belief, which may be incorrect, but is

*We thank Sarah Auster, Roberto Corrao, Enrico De Magistris, Nicodemo De Vito, Thomas J. Sargent and the participants to conferences and seminar presentations at Bielefeld University, Hebrew University of Jerusalem, CISEI-2017 (Anacapri), the 6th Workshop on Stochastic Methods in Game Theory (Erice), and ESEM 2017 (Lisbon) for useful comments. Financial support from the European Research Council (grant 324219) and Guido Cazzavillan Scholarship is gratefully acknowledged.

confirmed by the evidence available to him. A profile of strategies and beliefs with this property is a **self-confirming equilibrium** (henceforth, SCE; see Fudenberg and Levine 1993).¹ The standard definition of SCE assumes that agents are subjective expected utility maximizers, i.e., that they are ambiguity neutral.² Yet, a large body of empirical evidence supports the ambiguity aversion hypothesis. This is particularly relevant when agents only have scarce evidence of opponents' behavior and face therefore strategic uncertainty. Following Battigalli, Cerreia-Vioglio, Maccheroni and Marinacci (2015, henceforth BCMM), we analyze the SCE concept in games played by ambiguity averse agents. Unlike BCMM, who essentially restrict their attention to simultaneous-move games, we consider games with sequential moves, represented in extensive form. We also analyze a notion of rationalizable SCE that captures sophisticated strategic reasoning. In a sequential game, agents have evidence, at best, of how opponents play on the equilibrium path, but no evidence of how they would play off the path. Thus, sequential games constitute a natural context for self-confirming equilibrium analysis. In the rest of this introduction, we describe how this paper adds to the previous literature in general and to BCMM in particular.

BCMM analyze SCE in simultaneous-move population games played recurrently by agents with non-neutral attitudes toward ambiguity, which is the imperfect quantifiability of the relevant risks. Specifically, agents are assumed to have smooth-ambiguity preferences in the sense of Klibanoff, Marinacci, and Mukerji (2005, henceforth KMM). This decision model is flexible and analytically convenient for game theoretic applications. First, it separates ambiguity attitudes, a stable personal trait like risk attitudes, from the perception of uncertainty, which is a property of subjective beliefs affected by the game situation; second, it provides a parameterization of ambiguity aversion analogous to the parameterization of risk aversion, which simplifies comparative static exercises.

SCE is defined as follows: Let I denote the set of player roles (e.g., buyer and seller) and let S_i denote the set of pure strategies of any agent playing in role i . A profile of strategy distributions $\sigma^* = (\sigma_i^*)_{i \in I} \in \times_{i \in I} \Delta(S_i)$ is a self-confirming equilibrium if, for each $i \in I$ and each $s_i^* \in S_i$ with $\sigma_i^*(s_i^*) > 0$, there is a belief μ_i about the strategy distributions σ_{-i} of the opponents that justifies s_i^* as a KMM-best response and is consistent with the long-run distribution of ex post observations for i generated by s_i^* and σ_{-i}^* (e.g., the (s_i^*, σ_{-i}^*) -induced distribution of terminal nodes). Since the distribution of observations may not reveal the true underlying distribution of strategies σ_{-i}^* , agents may be uncertain about it.

Assuming that the own-payoff relevant consequences are observed by each agent after each play, BCMM prove a **monotonicity result**: higher ambiguity aversion entails a larger

¹A version of this equilibrium concept was first put forward by Battigalli (1987) and called “conjectural equilibrium” (see also Battigalli and Guaitoli, 1988). Part of the literature on this topic maintained the same terminology. Note that Fudenberg and Levine (1993) assume that agents observe ex-post the path of play, whereas Battigalli (1987) and most of the papers on conjectural equilibrium consider more general hypotheses about feedback. See the discussion in Battigalli et al. (2015) and the references therein.

²See Cerreia-Vioglio et al. (2013) and the survey by Marinacci (2015).

set of equilibria. Intuitively, for each agent, the strategy played repeatedly in equilibrium yields known risks, because the agent observes his long-run distribution of payoffs, while deviations are untested and may be perceived as ambiguous; therefore, higher ambiguity aversion penalizes deviations, but not the equilibrium choice. This monotonicity result implies that greater ambiguity aversion entails less predictability of strategies in the long-run, because the set of possible steady states is larger.³

The scope of BCMM’s analysis is essentially limited to simultaneous-move games and, possibly, games played in strategic form, such as experimental games played in the lab with the so-called “strategy method,”⁴ because it is well known that ambiguity aversion may make preferences over strategies dynamically inconsistent (e.g., Siniscalchi 2011). As a consequence, there would be incentives to make covert commitments if such commitment moves were available.⁵ Thus, the fact that agents in sequential games cannot (irreversibly) choose strategies—because they just choose actions at decision nodes—must be faced and dealt with explicitly. We assume that agents are sophisticated: they understand their future, contingent incentives; thus, they choose actions in early stages predicting that such incentives determine their actions in later stages, that is, they plan by “folding back” given their subjective beliefs about other players; hence they execute “**unimprovable**” strategies. Since unimprovable strategies may be different from best replies in the normal form of the game, the definition of equilibrium due to BCMM cannot be applied to sequential games. Note that this is unrelated to how agents react to completely unexpected events. The confirmed-beliefs condition of SCE implies that agents cannot be surprised on the equilibrium path; furthermore, the SCE concept does not rely on complete information, nor does it capture strategic reasoning; hence, it does not model how agents predict the reactions of others to unexpected moves. In this paper, we also analyze reactions to unexpected moves and strategic reasoning about such reactions.

Two questions naturally arise: First, how does SCE defined on the extensive form relate to SCE in the normal form of the same game? Second, does the comparative ambiguity aversion result extend to games with sequential moves? To elaborate on the first question, fix a sequential game Γ represented in extensive form with a set of paths (terminal nodes) Z and a specification of players’ **feedback** $f = (f_i)_{i \in I}$, that is, what they can observe at the end of each play; formally, each f_i is a function defined on Z . Then, we can derive the normal (or strategic) form $(G, F) = \mathcal{N}(\Gamma, f)$, where G is given by the normal-form payoff functions, and each player i ’s feedback f_i about the path is replaced by a corresponding normal-form feedback F_i defined on the set $S = \times_{i \in I} S_i$ of pure strategy profiles (e.g., each

³A similar result can be proved for comparative risk aversion in the particular case of *pure* SCE (see the lecture notes of Battigalli, 2017), and for comparative risk or ambiguity aversion in the case of rationalizable strategies (Battigalli et al., 2016a). However, the intuition and proof in the case of rationalizability are very different from the SCE case.

⁴See the survey by Brandts and Charness (2011) and the references therein.

⁵Unlike *overt* commitment, covert commitment moves are not observed by other players. The strategic advantages of overt commitment are well known at least since Schelling (1960) and do not depend on dynamic inconsistency as traditionally defined in decision theory.

player only observes his monetary payoff, which is a function of $s \in S$). Under subjective expected utility maximization, which is dynamically consistent, SCE in the extensive form is realization equivalent to SCE in the normal form. With ambiguity aversion, instead, we show that there may be different, non-nested sets of equilibrium outcomes; in other words, working with the normal form is neither too permissive nor too restrictive, it is just wrong.⁶

Given that the sets of SCE outcomes in the extensive and normal form do not coincide, we cannot rely on the monotonicity result of BCMM to argue that in a game with sequential moves the set of SCE outcomes expands as ambiguity aversion increases. Indeed, we show by example that monotonicity does not hold in general. It is still true that, on the equilibrium path, equilibrium actions entail known risks, while deviations may be perceived as ambiguous. But, if a deviation is followed by other actions of the same player, folding-back planning may require that plans (i.e., predictions) about these actions change with ambiguity aversion, and dynamic inconsistency may lead to an increase in the value of deviations.

Despite this, we can prove the monotonicity result for two special cases: (i) games where, on each path, no player moves more than once, and (ii) pure strategy equilibria of games with no chance moves.⁷ In case (i), the difficulties mentioned above cannot arise; in case (ii) one can show that, loosely speaking, in equilibrium ambiguity aversion is not distinguishable from the risk aversion of expected utility maximizers, who are dynamically consistent. We also prove a related noteworthy result: The set of SCE outcomes with ambiguity neutral agents is always included in the set of SCE outcomes with ambiguity averse agents. This means that the standard version of the SCE concept, by ignoring ambiguity aversion, overestimates the predictability of long-run outcomes in recurrent interactions.

Next, we turn to the well known issue of the impossibility of *overt* (i.e., observable) commitment and how this affects strategic reasoning and equilibrium. Rationality in dynamic games requires that agents choose subjective best replies at *all* information sets, including the unexpected one, given beliefs revised upon observing unexpected moves. We call this condition “**full unimprovability**.” If SCE is not refined so as to capture strategic reasoning, requiring full unimprovability rather than simple unimprovability does not change the set of equilibrium outcomes. The reason is that in an SCE agents may hold wrong beliefs about the reactions of others to non-equilibrium moves, which are necessarily unexpected; hence, they can expect responses that are irrational given the actual payoff functions of co-players. Assuming instead that (some features of) such payoff functions

⁶Fudenberg and Levine (1993) note that SCE is not normal-form invariant, arguing that the normal form is therefore insufficient. But their comment rests on the maintained assumption that players always observe ex-post the path of play, which in a simultaneous-move game is just the actual profile of strategies (actions of the normal form). Therefore, when they compare SCE in the extensive and normal form, they change what players can observe ex post about the behavior of others: only on-path actions in the extensive form, and complete strategies in the normal form.

⁷We call such equilibria “symmetric” because in the population game scenario they represent situations where all agents in the same role play in the same way.

are common knowledge, one can embed strategic reasoning into the SCE concept: if (a) players are rational, (b) their beliefs are confirmed, and (c) there is common belief of (a) and (b), then their strategies form a **rationalizable SCE**. This concept has been analyzed under the assumption of subjective expected utility maximization, that is, neutral ambiguity attitudes (e.g., Rubinstein and Wolinsky 1994, Dekel et al. 1999).⁸ Here we provide an extension for non-neutral ambiguity attitudes. The analysis involves some delicate technical details.

The rest of the paper is organized as follows. Sections 2 and 3 introduce the setup and the smooth ambiguity criterion; Section 4 defines and analyzes unimprovability; Section 5 defines our SCE concept; Section 6 presents comparative results for SCE; Section 7 introduces full unimprovability and shows that such strengthening by itself does not affect SCE outcomes; Section 8 relies on full unimprovability to define rationalizable SCE and provides comparative results for such concept; finally, Section 9 discusses the relevance of some assumptions and equilibrium concepts, and provides hints for the generalization of the analysis. The main text contains some intuitive arguments, but all formal proofs are collected in the Appendix.

2 Framework

We analyze an agent with non-neutral ambiguity attitudes who plays a game with sequential moves. We assume that the commitment technology of this agent is explicitly represented by the rules of the game. Therefore, the agent can control — i.e., irreversibly choose — only his (pure) actions at whatever information set is being reached. We also assume that he is sophisticated and therefore he takes this into account when he plans how to play the game.

We take the point of view of an agent who plays in role $i \in I$ of a *finite extensive-form game* Γ and has *perfect recall*. Let H_i denote the collection of **information sets** of i and let $A_i(h)$ be the set of actions available at $h \in H_i$. We assume for expositional simplicity that $|A_i(h)| \geq 2$ for each $h \in H_i$, where $|X|$ denotes the cardinality of a finite set X . This means that we include in H_i only the information sets where i is active. Let \emptyset denote the root of the game, then $\{\emptyset\} \in H_i$ if and only if i is a first mover.⁹ We endow H_i with the

⁸Battigalli (1987) and Battigalli and Guaitoli (1988) consider a weaker concept of SCE in rationalizable strategies justified by the following assumptions: (a) players are rational, (b) their beliefs are confirmed, and (c') there is common belief of (a) only. To the best of our knowledge, unlike plain SCE concept, there is no learning foundation of rationalizable SCE. On the other hand, we can give a kind of learning foundation of SCE in rationalizable strategies. See our discussion in Section 9.

⁹Our preferred representation of games in extensive form starts from sequences of action profiles, that correspond to the nodes of the game tree (e.g., Chapters 6 and 11 of Osborne and Rubinstein, 1994) and allows for the representation of players' information also at nodes where they are not active, such the root for players who are not first-movers. This affects the way we draw pictures and describe examples, but it is otherwise irrelevant for the analysis of the paper.

weak (respectively, strict) precedence relation \preceq ($<$) inherited from the game tree.¹⁰ The set of **strategies** for player i is $S_i = \times_{h \in H_i} A_i(h)$. For every $s_i \in S_i$ and $h \in H_i$, we let $s_{i,h}$ denote the action specified by s_i at h ; thus, $s_i = (s_{i,h})_{h \in H_i} \in \times_{h \in H_i} A_i(h)$.

We model randomization explicitly as the choice of a randomization device. Therefore it is important to allow for **chance moves** as the moves of a special player denoted by $0 \notin I$. With this, H_0 denotes the collection of information sets of the chance player, $A_0(h)$ is set of chance moves available at $h \in H_0$, and $S_0 = \times_{h \in H_0} A_0(h)$ is the set of “strategies” of the pseudo-player 0. Throughout, we maintain for simplicity the assumption that the *probabilities of chance moves are commonly known*. Such probabilities are specified by a “behavioral strategy” $\beta_0 \in \times_{h \in H_0} \Delta(A_0(h))$, with $\beta_0(a_0|h) > 0$ for every $h \in H_0$ and $a_0 \in A_0(h)$; β_0 induces the “mixed strategy” $\sigma_0 \in \Delta(S_0)$ such that

$$\sigma_0(s_0) = \prod_{h \in H_0} \beta_0(s_{0,h}|h) > 0$$

for every $s_0 \in S_0$. Thus, the outcome distributions respectively induced by β_0 and σ_0 coincide for every strategy profile of the true players (see Kuhn, 1953). Since players are always certain that the “mixed strategy of chance” is σ_0 , we model explicitly only each player i ’s beliefs about the behavior of the true opponents $-i = I \setminus \{i\}$. We let $S = \times_{j \in I} S_j$ denote the set of pure strategy profiles of the true players, whereas $S_{-i} = \times_{j \in I \setminus \{i\}} S_j$ and $S_{0,-i} = S_0 \times S_{-i}$ denote the set of pure strategy profiles of opponents respectively excluding and including chance.

Let Z denote the set of terminal nodes of the game. Every profile (s_0, s) induces a complete path, hence a terminal node, through the **outcome function**

$$\zeta : S_0 \times S \rightarrow Z.$$

Since the definition of ζ is standard, we take it for granted and then define some derived concepts using ζ .

For conceptual clarity, we also include in the description of the extensive form Γ a **consequence function**

$$\gamma : Z \rightarrow C$$

which specifies the material consequence $c = \gamma(z) \in C$ of each terminal node $z \in Z$. For example, we may have $C \subseteq \mathbb{R}^I$ where $c = (c_j)_{j \in I}$ is a consumption allocation or a distribution of monetary payoffs to players. Thus, player i ’s risk attitudes (preferences over objective lotteries of consequences) are represented by a **vNM utility function**

$$v_i : C \rightarrow \mathbb{R}.$$

¹⁰Perfect recall implies that, for all $h, h' \in H_i$, there are nodes (histories) $x \in h$ and $x' \in h'$ such that x precedes x' if and only if every node of h' is preceded by a node of h . With this, we can stipulate that, for all $h, h' \in H_i$, h strictly **precedes** h' , written $h < h'$, if every node of h' is strictly preceded by a node of h . Perfect recall implies that each $h \in H_i$ can have at most one immediate predecessor. The reflexive closure of $<$ is \preceq , an antisymmetric, and transitive relation that makes H_i a directed forest. If $\{\emptyset\} \in H_i$, then (H_i, \preceq) is a directed tree.

To ease notation, we write the **payoff function** of i as

$$u_i = v_i \circ \gamma : Z \rightarrow \mathbb{R}.$$

It is convenient to specify the information about strategies implied by any information set h . First, for any node x , let

$$S_{0,I}(x) = \{(s_0, s) \in S_0 \times S : x \prec \zeta(s_0, s)\}$$

denote the set of pure strategy profiles reaching x . With this, for any subset of nodes h ,

$$S_{0,I}(h) = \bigcup_{x \in h} S_{0,I}(x)$$

is the set of strategy profiles (s_0, s) reaching h ,

$$S_{0,-i}(h) = \text{proj}_{S_{0,-i}} S_{0,I}(h) = \{s_{0,-i} \in S_{0,-i} : \exists (x, s_i) \in h \times S_i, x \prec \zeta(s_i, s_{0,-i})\}$$

is the set of pure strategy profiles of chance and opponents that allow for h , and

$$S_i(h) = \text{proj}_{S_i} S_{0,I}(h) = \{s_i \in S_i : \exists (x, s_{0,-i}) \in h \times S_{0,-i}, x \prec \zeta(s_i, s_{0,-i})\}$$

is the sets of i 's strategies allowing for h .

It is useful to keep in mind that perfect recall implies the following factorization:

$$\forall h \in H_i, S_{0,I}(h) = S_i(h) \times S_{0,-i}(h).$$

Furthermore, it also implies that

$$\forall g, h \in H_i, g \prec h \Rightarrow S_{0,-i}(h) \subseteq S_{0,-i}(g).$$

Intuitively, i obtains finer information about strategies as the play unfolds.

Each agent playing in role i knows that his opponents are drawn at random from large populations $j \in I \setminus \{i\}$ of agents, with each agent playing a pure strategy. The **distribution of pure strategies in population j** is some *unknown* measure $\sigma_j \in \Delta(S_j)$, hence, by random matching, the *objective* probability of facing opponents playing pure strategy profile $s_{-i} = (s_j)_{j \in I \setminus \{i\}}$ is the unknown product¹¹

$$\sigma_{-i}(s_{-i}) = \prod_{j \in I \setminus \{i\}} \sigma_j(s_j).$$

¹¹Statistical independence follows from random matching: For each $i \in I$, let P_i denote the of agents playing in role i , and let $\varsigma_i : P_i \rightarrow S_i$ denote the (measurable) strategy map of population i . If agents are drawn at random from their populations, that is, according to a *uniform* distribution on $\times_{i \in I} P_i$, then the induced distribution on S given $(\varsigma_i)_{i \in I}$ is a product measure.

To ease notation, we identify each profile of distributions $(\sigma_j)_{j \in I \setminus \{i\}}$ with the corresponding product distribution on S_{-i} . We let

$$\Sigma_{-i} = \left\{ \sigma_{-i} \in \Delta(S_{-i}) : \exists (\sigma_j)_{j \in I \setminus \{i\}} \in \times_{j \in I \setminus \{i\}} \Delta(S_j), \sigma_{-i} = \times_{j \in I \setminus \{i\}} \sigma_j \right\}$$

denote the set of these product distributions.¹² We endow Σ_{-i} with the topology inherited from the Euclidean topology on $\mathbb{R}^{S_{-i}}$, which makes it compact, and with the corresponding Borel sigma algebra $\mathcal{B}(\Sigma_{-i})$.

At each point of the game, the agent playing in role i has some **belief** $\mu_i \in \Delta(\Sigma_{-i})$. The belief μ_i that i holds at the beginning of the game is i 's **prior**. For each $\mu_i \in \Delta(\Sigma_{-i})$, we let $p_{\mu_i} \in \Delta(S_{-i})$ denote the **predictive probabilities** implied by μ_i : for each $s_{-i} \in S_{-i}$,

$$p_{\mu_i}(s_{-i}) = \int_{\Sigma_{-i}} \sigma_{-i}(s_{-i}) \mu_i(d\sigma_{-i}).$$

We summarize our notation in the following table and illustrate it with an example. We will refer to this example repeatedly.

Notation	Terminology
Γ	extensive-form game
$i, j \in I$	players ($j = 0 \notin I$ denotes chance)
$h \in H_i$	information sets of i
\preceq ($<$)	(strict) precedence relation of Γ
(H_i, \preceq)	directed forest of information sets of i
$a_i \in A_i(h)$	i 's actions at $h \in H_i$

$s_i \in S_i = \times_{h \in H_i} A_i(h)$	strategies of i
$s \in S$ ($s_{-i} \in S_{-i}$, $s_{0,-i} \in S_{0,-i}$)	strategy profiles (of $-i = I \setminus \{i\}$, of $\{0\} \cup I \setminus \{i\}$)
$S_{0,I}(h)$	strategy profiles (including 0) reaching h
$S_i(h) = \text{proj}_{S_i} S_{0,I}(h)$	strategies of i allowing for h
$S_{0,-i}(h) = \text{proj}_{S_{0,-i}} S_{0,I}(h)$	strategy profiles of 0 and $-i$ allowing for h
$S_{-i}(h) = \text{proj}_{S_{-i}} S_{0,I}(h)$	strategy profiles of $-i$ allowing for h
$\sigma_j \in \Delta(S_j)$	strategy distributions on S_j
$\sigma_{-i} \in \Sigma_{-i} \subset \Delta(S_{-i})$	product distributions on S_{-i}
$\mu_i \in \Delta(\Sigma_{-i})$	beliefs of i
$p_{\mu_i} \in \Delta(S_{-i})$	predictive probabilities implied by μ_i
$z \in Z$	terminal histories/nodes
$\zeta : S_0 \times S \rightarrow Z$	outcome function
$\gamma : Z \rightarrow C$	consequence function
$v_i : C \rightarrow \mathbb{R}$	vNM utility function of i
$u_i = v_i \circ \gamma : Z \rightarrow \mathbb{R}$	payoff function of i

¹²We use symbol \times to denote both the Cartesian product of sets and the product of measures.




Figure 1: Leading Example

Example 1 *The game depicted in Figure 1 is a two-person, common-interest, multistage game where $I = \{1, 2\}$ and 0 is chance. If we identify nodes with histories,¹³ information sets, actions sets and terminal histories/nodes are as follows:*

$$\begin{aligned}
 H_0 &= \{\text{In}\}, \quad H_1 = \{\{\emptyset\}, \{(\text{In}, \text{G})\}\}, \quad H_2 = \{(\text{In}, \text{G})\}, \\
 A_0(\{\text{In}\}) &= \{\text{E}, \text{G}\}, \quad A_2(\{(\text{In}, \text{G})\}) = \{\text{L}, \text{R}\}, \\
 A_1(\{\emptyset\}) &= \{\text{In}, \text{Out}\}, \quad A_1(\{(\text{In}, \text{G})\}) = \{\text{T}, \text{M}, \text{B}\}, \\
 Z &= \{\text{Out}, (\text{In}, \text{E})\} \cup \{(\text{In}, \text{G})\} \times \{\text{T}, \text{M}, \text{B}\} \times \{\text{L}, \text{R}\}.
 \end{aligned}$$

Numbers at terminal histories/nodes, including the boxes in the matrix subgame, give the common payoff of players 1 and 2. The probabilities of chance moves are $\sigma_0(\text{E}) = \sigma_0(\text{G}) = \frac{1}{2}$.

Assume that the agent in role 1 has the following belief:

$$\mu_1(\sigma_2) = \begin{cases} \frac{1}{2} & \text{if } \sigma_2 \in \{\delta_{\text{L}}, \delta_{\text{R}}\}, \\ 0 & \text{otherwise,} \end{cases}$$

where δ_x denotes the Dirac measure supported by x . Intuitively, he thinks that all the agents playing in role 2 “attended the same school” and hence are doing the same thing, but he does not know what. The induced predictive probabilities are $p_{\mu_1}(\text{L}) = p_{\mu_1}(\text{R}) = \frac{1}{2}$. \blacktriangle

¹³See, e.g., Chapter 11 in Osborne and Rubinstein (1994).

3 Smooth-ambiguity preferences over actions

In this section, we take the perspective of an agent, or decision maker, playing in role i , henceforth DM_i , with given beliefs about the behavior of agents in different roles and a well defined plan, or strategy. Specifically, he has a plan s_i specifying the action $s_{i,h} \in A_i(h)$ he expects to take (but he is not committed to take) at each information set $h \in H_i$; he has no randomization technology beyond what is already explicitly represented in the extensive form of the game (see Section 2), and we assume that he is certain about his contingent behavior, i.e., he has a deterministic contingent plan.

Conditional distributions and conditional objective expected utility Let $\Sigma_{-i}(h)$ denote the set of distributions that assign positive probability to $S_{-i}(h)$, that is,

$$\Sigma_{-i}(h) = \{\sigma_{-i} \in \Sigma_{-i} : \sigma_{-i}(S_{-i}(h)) > 0\}.$$

For every $h \in H_i$ and $\sigma_{-i} \in \Sigma_{-i}(h)$, we can compute the objective conditional distribution on the opponents' strategy profiles consistent with h :

$$\forall (s_0, s_{-i}) \in S_{0,-i}(h), \sigma_{0,-i}(s_0, s_{-i}|h) = \frac{\sigma_0(s_0)\sigma_{-i}(s_{-i})}{(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))} \quad (1)$$

(note that that $\sigma_{-i} \in \Sigma_{-i}(h)$ if and only if $(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h)) > 0$, because σ_0 is strictly positive). With this, we can define the **strategic-form vNM conditional expected utility function**

$$U_i(\cdot, \cdot|h) : \begin{array}{ll} S_i(h) \times \Sigma_{-i}(h) & \rightarrow \mathbb{R}, \\ (s_i, \sigma_{-i}|h) & \mapsto \sum_{s_{0,-i} \in S_{0,-i}(h)} \sigma_{0,-i}(s_{0,-i}|h) u_i(\zeta(s_i, s_{0,-i})). \end{array}$$

In words, if i is certain that σ_{-i} is the true objective probability model, then upon observing h (he believes that) his conditional objective expected utility from following strategy $s_i \in S_i(h)$ is $U_i(s_i, \sigma_{-i}|h)$.¹⁴

Plans and replacements Plan s_i yields a continuation on the information sets in H_i following any given $h \in H_i$ (that is, the projection of s_i onto $\times_{\{h' \in H_i : h \preceq h'\}} A_i(h')$). DM_i expects to continue according to this plan, but he knows (by perfect recall) that he has already chosen the actions leading to h , possibly violating s_i , and he considers the consequences of choosing action $a_i \in A_i(h)$, again possibly violating s_i . It is convenient to define

¹⁴One can show that this coincides with the more familiar formula

$$U_i(s_i, \sigma_{-i}|h) = \sum_{x \in h} \mathbb{P}_{s_i, \sigma_{-i}}(x|h) \sum_{z \in Z} \mathbb{P}_{s_i, \sigma_{-i}}(z|x) u_i(z)$$

where $\mathbb{P}_{s_i, \sigma_{-i}}(\cdot|\cdot)$ denotes the probability of reaching a node conditional on an information set, or an earlier node, given by σ_{-i} and the known probabilities of chance moves.

the **replacement** plan $(s_{i|h}, a_i)$ obtained by replacing s_i with the already chosen actions at information sets preceding h and with action a_i at h :

$$(s_{i|h}, a_i)_{h'} = \begin{cases} a_i & \text{if } h' = h, \\ \alpha_i(h', h) & \text{if } h' \prec h, \\ s_{i,h'} & \text{otherwise,} \end{cases}$$

where $\alpha_i(h', h)$ is the action chosen at $h' \prec h$ in order to reach h .¹⁵ Finally, we let $s_{i|h}$ denote the replacement plan obtained when action $s_{i,h}$ is played at h :

$$(s_{i|h})_{h'} = \begin{cases} \alpha_i(h', h) & \text{if } h' \prec h, \\ s_{i,h'} & \text{otherwise.} \end{cases}$$

Action values We assume that DM_i 's preferences over actions, given his beliefs and plan, satisfy the smooth-ambiguity model of KMM: On top of the vNM utility function $v_i : C \rightarrow \mathbb{R}$ specified by game Γ (hence, the payoff function $u_i = v_i \circ \gamma$), we assume that there is a continuous and *strictly increasing second-order utility* function

$$\phi_i : \mathbb{V}_i \rightarrow \mathbb{R},$$

where

$$\mathbb{V}_i = \left[\min_z v_i(\gamma(z)), \max_z v_i(\gamma(z)) \right]$$

is the convex hull of the range of v_i . For every given $h \in H_i$, $\mu_i \in \Delta(\Sigma_{-i}(h))$, and $s_i \in S_i$, DM_i assigns values to actions $a_i \in A_i(h)$ as follows:

$$V_i(a_i|h; s_i, \mu_i, \phi_i) = \phi_i^{-1} \left(\int_{\Sigma_{-i}(h)} \phi_i(U_i((s_{i|h}, a_i), \sigma_{-i}|h)) \mu_i(d\sigma_{-i}) \right). \quad (2)$$

Condition $\mu_i(\Sigma_{-i}(h)) = 1$ suggests that, in eq. (2), we interpret μ_i as the conditional belief of DM_i upon observing h . We postpone the definition of conditional beliefs to Section 4. If μ_i assigns probability 1 to some $\sigma_{-i} \in \Sigma_{-i}(h)$ (that is, $\mu_i = \delta_{\sigma_{-i}} \in \Delta(\Sigma_{-i}(h))$), then

$$V_i(a_i|h; s_i, \mu_i, \phi_i) = \phi_i^{-1}(\phi_i(U_i((s_{i|h}, a_i), \sigma_{-i}|h))) = U_i((s_{i|h}, a_i), \sigma_{-i}|h).$$

Thus, ambiguity attitudes are immaterial when DM_i is certain about the true probability model, because in this case he does not perceive any ambiguity. Note also that (2) boils down to the classical subjective expected utility formula if ϕ_i is linear (ambiguity neutrality), hence equivalent to the identity function $\text{Id}_{\mathbb{V}_i}$. On the other hand, ambiguity aversion is characterized by the concavity of ϕ_i . We emphasize in our notation only the dependence of values of i 's actions on parameter ϕ_i , not on the vNM utility function v_i , because we are going to consider different possible shapes of ϕ_i (in particular, linear and concave) with a fixed v_i .

¹⁵By perfect recall, α_i is well defined.

Example 2 In the game of Figure 1,

$$\begin{aligned} U_1(\text{In.B}, \delta_L | \{(\text{In}, \text{G})\}) &= 36, \\ U_1(\text{In.B}, \delta_R | \{(\text{In}, \text{G})\}) &= 1, \end{aligned}$$

and

$$\begin{aligned} U_1(\text{In.B}, \delta_L | \{\emptyset\}) &= \frac{1}{2} \cdot 36 = 18, \\ U_1(\text{In.B}, \delta_R | \{\emptyset\}) &= \frac{1}{2} \cdot 1 = \frac{1}{2}. \end{aligned}$$

Assume $\phi_1(u_i) = \sqrt{u_i}$ and let μ_1 be the belief of Example 1, that is

$$\mu_1(\sigma_2) = \begin{cases} \frac{1}{2} & \text{if } \sigma_2 \in \{\delta_L, \delta_R\}, \\ 0 & \text{otherwise;} \end{cases}$$

then

$$V_1(\text{B} | \{(\text{In}, \text{G})\}; s_1, \mu_1, \phi_1) = \left(\frac{1}{2} \sqrt{36} + \frac{1}{2} \sqrt{1} \right)^2 = (3.5)^2 = 12.25$$

for each $s_1 \in S_1((\text{In}, \text{G}))$. ▲

4 Conditional beliefs and unimprovability

A prior belief $\mu_i \in \Delta(\Sigma_{-i})$ over co-players' strategy distributions induces a joint belief $\pi_{\mu_i} \in \Delta(S_{0,-i} \times \Sigma_{-i})$ determined by the following equation:

$$\forall (s_0, s_{-i}, E_{-i}) \in S_{0,-i} \times \mathcal{B}(\Sigma_{-i}), \pi_{\mu_i}(\{(s_0, s_{-i})\} \times E_{-i}) = \sigma_0(s_0) \int_{E_{-i}} \sigma_{-i}(s_{-i}) \mu_i(d\sigma_{-i}). \quad (3)$$

Note that¹⁶

$$\pi_{\mu_i}(S_{0,-i} \times E_{-i}) = \mu_i(E_{-i}), \quad (4)$$

and

$$\pi_{\mu_i}(S_0 \times \{s_{-i}\} \times \Sigma_{-i}) = p_{\mu_i}(s_{-i}).$$

Each information set $h \in H_i$ corresponds to the conditioning event $S_{0,-i}(h) \times \Sigma_{-i}(h)$. Let

$$H_i(\mu_i) = \{h \in H_i : \pi_{\mu_i}(S_{0,-i}(h) \times \Sigma_{-i}(h)) > 0\} = \{h \in H_i : p_{\mu_i}(S_{-i}(h)) > 0\}$$

¹⁶One may use beliefs over the product space $S_{0,-i} \times \Sigma_{-i}$ as the primitive object. Of course, the structural assumption that agents know the probabilities of chance moves and that they are randomly matched with opponents drawn large populations implies that only beliefs that admit the representation (3) for some $\mu_i \in \Delta(\Sigma_{-i})$ are admissible. Then we can derive the belief over distributions using (4). For a decision theoretic approach akin to the one discussed in this footnote, with models in place of population distributions, see Cerreia-Vioglio et al. (2013).

denote the subset of information sets of player i that he believes he can reach with positive probability.¹⁷ If $h \in H_i(\mu_i)$, then we can derive the conditional probability of every measurable set $E_{-i} \in \mathcal{B}(\Sigma_{-i})$ of strategy distributions:

$$\begin{aligned}
\mu_i(E_{-i}|h) &= \pi_{\mu_i}(S_{0,-i} \times E_{-i} | S_{0,-i}(h) \times \Sigma_{-i}(h)) \\
&= \frac{\pi_{\mu_i}((S_{0,-i} \times E_{-i}) \cap (S_{0,-i}(h) \times \Sigma_{-i}(h)))}{\pi_{\mu_i}(S_{0,-i}(h) \times \Sigma_{-i}(h))} \\
&= \frac{\pi_{\mu_i}(S_{0,-i}(h) \times (\Sigma_{-i}(h) \cap E_{-i}))}{\sum_{(s_0, s_{-i}) \in S_{0,-i}(h)} \sigma_0(s_0) \int_{\Sigma_{-i}(h)} \sigma_{-i}(s_{-i}) \mu_i(d\sigma_{-i})} \\
&= \frac{\sum_{(s_0, s_{-i}) \in S_{0,-i}(h)} \sigma_0(s_0) \int_{\Sigma_{-i}(h) \cap E_{-i}} \sigma_{-i}(s_{-i}) \mu_i(d\sigma_{-i})}{\sum_{(s_0, s_{-i}) \in S_{0,-i}(h)} \sigma_0(s_0) p_{\mu_i}(s_{-i})} \\
&= \frac{\int_{E_{-i}} (\sigma_0 \times \sigma_{-i})(S_{0,-i}(h)) \mu_i(d\sigma_{-i})}{(\sigma_0 \times p_{\mu_i})(S_{0,-i}(h))}. \tag{5}
\end{aligned}$$

For example, if μ_i has *finite support*:

$$\mu_i(\sigma_{-i}|h) = \frac{\mu_i(\sigma_{-i})(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))}{\sum_{\sigma'_{-i} \in \text{Supp}\mu_i} \mu_i(\sigma'_{-i})(\sigma_0 \times \sigma'_{-i})(S_{0,-i}(h))}.$$

With this, we let

$$V_i(a_i|h; s_i, \mu_i, \phi_i) = \phi_i^{-1} \left(\int_{\Sigma_{-i}(h)} \phi_i(U_i((s_i|h, a_i), \sigma_{-i}|h)) \mu_i(d\sigma_{-i}|h) \right) \tag{6}$$

whenever $h \in H_i(\mu_i)$, and we derive well defined preferences over actions (given μ_i and s_i) only for information sets that are possible according to μ_i . This is what we need for our baseline definition of unimprovability, which is instead silent about choices at information sets deemed unreachable.

Definition 1 A strategy s_i is (μ_i, ϕ_i) -*unimprovable* if

$$\forall h \in H_i(\mu_i), s_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; s_i, \mu_i, \phi_i).$$

Since the game has finite horizon, we can interpret unimprovability as “folding-back optimality”: given (μ_i, ϕ_i) , DM_i derives a contingent plan that prescribes a choice for each information set he deems reachable. He starts from information sets $h \in H_i(\mu_i)$ with no followers in $H_i(\mu_i)$, and to each one of them he assigns an action $s_{i,h}$ that maximizes

¹⁷The equality holds because the probabilities of chance moves are strictly positive.

$V_i(a_i|h; s_i, \mu_i, \phi_i)$. Then he folds back, considering the information sets $h \in H_i(\mu_i)$ such that every follower has no further followers; for every such follower, viz. h' , DM_i predicts that the previously selected maximizing action $s_{i,h'}$ will be chosen; and so on until all the reachable information sets in $H_i(\mu_i)$ have been covered backwards.

Of course, we could define beliefs, and thus impose optimality requirements, also at information sets in $H_i \setminus H(\mu_i)$. For the time being, we are not interested in doing so: the moves of DM_i at each $h \in H_i \setminus H(\mu_i)$ will be immaterial for the equilibrium outcomes (by ex post perfect recall and confirmed beliefs, see Section 5), and impossible to predict for the opponents as long as we do not assume that they know the payoff function of i (cf. Section 7).

From the point of view of an external observer, or of agents in roles different from i , it is impossible to distinguish between two strategies of DM_i that yield the same outcomes independently of the opponents' behavior. This leads to the following notion of equivalence, which will play an important role in comparing self-confirming equilibria for different levels of ambiguity aversion (see Section 6).

Definition 2 (*Kuhn, 1953*) *Two (possibly degenerate) strategy distributions σ_i^* and σ_i are realization-equivalent if they induce the same distribution on terminal nodes, that is,*

$$\forall (z, s_0, s_{-i}) \in Z \times S_0 \times S_{-i}, \quad \sum_{s_i: \zeta(s_0, s_i, s_{-i})=z} \sigma_i^*(s_i) = \sum_{s_i: \zeta(s_0, s_i, s_{-i})=z} \sigma_i(s_i).$$

The set of strategy distributions realization-equivalent to σ_i^* is denoted by $[\sigma_i^*]$.

We let $[s_i^*]$ denote the set strategies (that is, Dirac distributions) realization-equivalent to s_i^* .

Remark 1 *Fix any $\sigma_i, \sigma_i^* \in \Delta(S_i)$; $\sigma_i \in [\sigma_i^*]$ if and only if $\sigma_i([s_i^*]) = \sigma_i^*([s_i^*])$ for every $s_i^* \in S_i$.*

Let $H_i(s_i) = \{h \in H_i : s_i \in S_i(h)\}$ denote the subset of information sets of DM_i that can be reached when s_i is played. Focusing on pure strategies, we obtain the following observation:

Remark 2 (*Theorem 1, Kuhn 1953*) *Fix any $s_i, s_i^* \in S_i$; $s_i \in [s_i^*]$ if and only if s_i and s_i^* are behaviorally equivalent, that is, if and only if $H_i(s_i) = H_i(s_i^*)$ and $s_{i,h} = s_{i,h}^*$ for each $h \in H_i(s_i^*)$.*

Now, suppose that DM_i is ambiguity neutral: $\phi_i = \text{Id}_{V_i}$. Then, by a classical dynamic programming result, unimprovability is equivalent to “global” (ex ante) subjective EU-maximization.¹⁸

¹⁸All the dynamic programming results of this section can be proved by standard folding-back arguments.

Proposition 1 For every strategy $s_i^* \in S_i$ and prior $\mu_i \in \Delta(\Sigma_{-i})$ the following are equivalent:

- (1) $[s_i^*]$ contains a $(\mu_i, \text{Id}_{\mathbb{V}_i})$ -unimprovable strategy,
- (2) $s_i^* \in \arg \max_{s_i \in S_i} \sum_{(s_0, s_{-i})} \sigma_0(s_0) p_{\mu_i}(s_{-i}) u_i(\zeta(s_0, s_i, s_{-i}))$.

We introduce the following strengthening of unimprovability:

Definition 3 A strategy s_i is (μ_i, ϕ_i) -sequentially optimal if

$$\forall h \in H_i(\mu_i), s_i \in \arg \max_{s'_i \in S_i} \phi_i^{-1} \left(\int_{\Sigma_{-i}(h)} \phi_i \left(U_i(s'_i|_h, \sigma_{-i}|_h) \right) \mu_i(d\sigma_{-i}|_h) \right).$$

If DM_i has dynamically inconsistent preferences over strategies, a (μ_i, ϕ_i) -sequentially optimal strategy may not exist, as illustrated in Example 3 below. However, if DM_i is ambiguity neutral (hence, his preferences are dynamically consistent), unimprovability coincides with sequential optimality:

Proposition 2 A strategy s_i is $(\mu_i, \text{Id}_{\mathbb{V}_i})$ -unimprovable if and only if it is $(\mu_i, \text{Id}_{\mathbb{V}_i})$ -sequentially optimal.

Example 3 Consider the game of Figure 1 and the belief μ_1 of Example 1:

$$\mu_1(\sigma_2) = \begin{cases} \frac{1}{2} & \text{if } \sigma_2 \in \{\delta_L, \delta_R\}, \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$H_1(\mu_1) = H_1 = \{\{\emptyset\}, \{\text{In}, \text{G}\}\}.$$

The induced belief on $S_{0,2} \times \Sigma_2$ is

$$\pi_{\mu_1}(s_0, s_2, \sigma_2) = \begin{cases} \frac{1}{4} & \text{if } (s_0, s_2, \sigma_2) \in \{\text{E}, \text{G}\} \times \{(\text{L}, \delta_L), (\text{R}, \delta_R)\}, \\ 0 & \text{otherwise.} \end{cases}$$

For an ambiguity neutral player 1 with belief μ_1 , the value of M and B at (In, G) is $(\frac{1}{2}36 + \frac{1}{2}1) > 9$, so it is higher than the value of T. Therefore, by folding-back, the $(\mu_1, \text{Id}_{\mathbb{V}_i})$ -unimprovable strategies are In.M and In.B .

Now suppose instead that the ambiguity attitudes of player 1 are represented by some strictly concave $\bar{\phi}_1$ with:

$$\begin{aligned} \bar{\phi}_1(u) &= \sqrt{u} \text{ if } 1 \leq u < 36, \\ \bar{\phi}_1\left(\frac{1}{2}\right) &= -1. \end{aligned}$$

At (In, G) , player 1 still prefers M (or B) over T, because:

$$\begin{aligned} V_1(\text{M} | \{(\text{In}, \text{G})\}; s_1, \mu_1; \bar{\phi}_1) &= \bar{\phi}_1^{-1} \left(\frac{1}{2} \bar{\phi}_1(1) + \frac{1}{2} \bar{\phi}_1(36) \right) \\ &= (3.5)^2 > 9 = V_1(\text{T} | \{(\text{In}, \text{G})\}; s_1, \mu_1, \bar{\phi}_1). \end{aligned}$$

Hence, a $(\mu_1, \bar{\phi}_1)$ -unimprovable strategy must prescribe action M or B at (In, G) . But then, it must also prescribe action Out at $\{\emptyset\}$. Indeed, for every strategy s_1 such that $s_{1,(\text{In}, \text{G})} \in \{\text{M}, \text{B}\}$ we have:

$$\begin{aligned} V_1(\text{In} | \{\emptyset\}; s_1, \mu_1, \bar{\phi}_1) &= \bar{\phi}_1^{-1} \left(\frac{1}{2} \cdot \bar{\phi}_1 \left(\frac{1}{2} \cdot 36 \right) + \frac{1}{2} \cdot \bar{\phi}_1 \left(\frac{1}{2} \cdot 1 \right) \right) \\ &= \left(\frac{1}{2} \cdot \sqrt{18} - \frac{1}{2} \right)^2 < 4 = V_1(\text{Out} | \{\emptyset\}; s_1, \mu_1, \bar{\phi}_1). \end{aligned}$$

So, on the one hand, the only $(\mu_1, \bar{\phi}_1)$ -unimprovable strategies are Out.M and Out.B. On the other hand, from the perspective of the agent at the root of Γ , the value of committing to strategy In.T is

$$V_1(\text{In} | \{\emptyset\}; \text{In.T}, \mu_1, \bar{\phi}_1) = \left(\sqrt{\frac{1}{2} \cdot 9} \right)^2 > 4 = V_1(\text{Out} | \{\emptyset\}; \text{Out}.a_1, \mu_1, \bar{\phi}_1), \quad (7)$$

for all $a_1 \in \{\text{T}, \text{M}, \text{B}\}$. So, player 1 would commit to In.T if he only could. \blacktriangle

This example illustrates the well-known dynamic inconsistency of preferences of decision makers with non-neutral attitudes towards ambiguity.^{19,20} To address this problem, we assume that agents are sufficiently sophisticated to understand the incentives they would face at each information set deemed possible, and plan/predict their contingent behavior by folding back. The resulting (μ_i, ϕ_i) -unimprovable strategy (or strategies) represents how agents in role i with belief μ_i predict they would choose at future information sets; such strategy and μ_i yield a value for each action available at the current information set.

Of course, DM_i may be indifferent at some information sets. A further refinement can be obtained for ambiguity averse agents by imposing a consistent-planning condition: whenever DM_i is indifferent at h then he breaks ties according to the preferences at the immediate predecessor of h in H_i . If this does not solve all the indifferences, ties are broken according to the preferences of the twice-removed predecessor, and so on. We omit this refinement for simplicity, and also because we find it arbitrary.

¹⁹See Siniscalchi (2011), for illustrative examples and an in-depth analysis of this issue.

²⁰Note that these dynamic inconsistencies arise as a consequence of the combination of Bayesian updating and non-neutral ambiguity attitudes. Indeed, it is not even obvious from the decision theoretic literature that ambiguity averse players are supposed to update beliefs according to the standard rules of conditional probabilities (see Epstein and Schneider 2007, Hanany and Klibanoff 2009, and Hanany et al. 2017). Instead we take the position that these rules are part of rational cognition, and we stick to them. This position is supported also by works that justify Bayesian updating in an evolutionary perspective, see Blume and Easley (2006) and the references therein.

5 Selfconfirming equilibrium

BCMM analyze a notion of smooth self-confirming equilibrium under the assumption that agents play the strategic form of a game with feedback and ambiguity attitudes, as with the strategy method in lab experiments. Specifically, consider a triple (Γ, f, ϕ) , where Γ is a standard extensive-form game, $f = (f_i : Z \rightarrow M)_{i \in I}$ is a profile of feedback functions such that every f_i describes the message $m \in M$ that player i observes ex post as a function of the terminal node, and $\phi = (\phi_i : \mathbb{V}_i \rightarrow \mathbb{R})_{i \in I}$ is a profile of strictly increasing functions capturing players attitudes toward ambiguity. The **structural strategic feedback function** of i associates each strategy profile (s_0, s) with a corresponding message, that is, $F_i = f_i \circ \zeta : S_0 \times S \rightarrow M$. We let $\hat{F}_i(s_i, \sigma_{-i}) \in \Delta(M)$ denote the **pushforward** distribution of messages induced by strategy s_i and the profile of strategy distributions σ_{-i} , given σ_0 . Specifically:

$$\forall (s_i, \sigma_{-i}, m) \in S_i \times \Sigma_{-i} \times M, \hat{F}_i(s_i, \sigma_{-i})(m) = \sum_{(s_0, s_{-i}) : F_i(s_0, s_{-i}) = m} \sigma_0(s_0) \sigma_{-i}(s_{-i}).$$

To relate to BCMM it is convenient to define the strategic, or normal form of a game with feedback (Γ, f) . The **normal-form** (expected) **payoff function** of player i is $U_i : S \rightarrow \mathbb{R}$ with

$$\forall s \in S, U_i(s) = \sum_{s_0 \in S_0} \sigma_0(s_0) u_i(\zeta(s_0, s))$$

Similarly, we define **normal-form feedback function** $\bar{F}_i : S \rightarrow \bar{M}$ as follows: If strategy profile s is played in the long run, then i observes the *distribution* of messages determined by s and chance probabilities. Therefore, we let $\bar{M} = \Delta(M)$ and

$$\forall (s, m) \in S \times M, \bar{F}_i(s)(m) = \hat{F}_i(s_i, \delta_{s_{-i}})(m) = \sum_{s_0 : F_i(s_0, s) = m} \sigma_0(s_0).$$

With this, the normal form of (Γ, f) is $\mathcal{N}(\Gamma, f) = (S_i, U_i, \bar{F}_i)_{i \in I}$. The equilibrium concept of BCMM applies to $(\mathcal{N}(\Gamma, f), \phi) = (S_i, U_i, \bar{F}_i, \phi_i)_{i \in I}$ under the assumption that each agent in role i covertly commits in advance to a strategy s_i . Here, instead, we analyze an equilibrium concept that is appropriate when agents play (Γ, f, ϕ) with the “direct method” making choices as the play unfolds, and we compare it with the strategic-form concept of BCMM.

Given that the information structure of Γ is assumed to satisfy perfect recall, we maintain the assumption that (Γ, f) satisfies “**ex post perfect recall**”:²¹

²¹See Battigalli et al. (2016b). There, the statement of the ex post perfect recall property is slightly different. The two versions are equivalent for extensive-form representations that specify the information of each player i at each nonterminal node, not only those where i is active (cf. Battigalli and Bonanno, 1999). Otherwise, one should use the statement of this paper.

Assumption (*Ex post perfect recall*) For every player $i \in I$, the augmented collection of information sets that includes the partition of Z induced by f_i ,

$$\hat{H}_i = H_i \cup \{f_i^{-1}(m) : m \in f_i(Z)\},$$

satisfies the perfect recall assumption. Thus, in particular, for all terminal histories $z, z' \in Z$, if there are an information set $h \in H_i$ and a node $x \in h$ such that $x \prec z$ and either z' has no predecessor in h , or $\alpha_i(x, z) \neq \alpha_i(x', z')$ for the predecessor x' of z' in h , then $f_i(z) \neq f_i(z')$.

Furthermore, we also consider (but we do not always assume) the following property of feedback:²²

Definition 4 An extensive-form game with feedback (Γ, f) satisfies **observable payoffs** whenever the payoff of every player only depends on his ex post information:

$$\forall (i, z, z') \in I \times Z^2, f_i(z) = f_i(z') \Rightarrow u_i(z) = u_i(z').$$

In other words, the payoff function is constant on each element of the ex post information partition $\{f_i^{-1}(m) : m \in f_i(Z)\}$. We say that (Γ, f, ϕ) satisfies observable payoffs if (Γ, f) does.

In some examples, we will assume that agents observe the terminal node they reach, that is, $f_i = \text{Id}_Z$. We call this hypothesis **perfect feedback**.

Definition 5 A **self-confirming equilibrium (SCE)** of (Γ, f, ϕ) is a profile of strategy distributions $\bar{\sigma} = (\bar{\sigma}_i)_{i \in I}$ with the following property: For each $i \in I$ and $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$ there is a belief $\mu_{\bar{s}_i} \in \Delta(\Sigma_{-i})$ such that

(rationality) \bar{s}_i is $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable,

(confirmed beliefs) $\mu_{\bar{s}_i} \left(\left\{ \sigma_{-i} \in \Sigma_{-i} : \hat{F}_i(\bar{s}_i, \sigma_{-i}) = \hat{F}_i(\bar{s}_i, \bar{\sigma}_{-i}) \right\} \right) = 1$.

An SCE $\bar{\sigma}$ is a **symmetric SCE (symSCE)** if, for each $i \in I$, there is a pure strategy \bar{s}_i with $\bar{\sigma}_i(\bar{s}_i) = 1$, that is, if all agents in the same population i play the same pure strategy.

The confirmed beliefs condition requires that the belief $\mu_{\bar{s}_i}$ justifying \bar{s}_i exclude all the distributions that are not observationally equivalent to the true one, $\bar{\sigma}_{-i}$. When profiles $\bar{\sigma}$ and $(\mu_{\bar{s}_i})_{i \in I, \bar{s}_i \in \text{Supp} \bar{\sigma}_i}$ satisfy the foregoing SCE conditions, we say that $\bar{\sigma}$ is **justified by confirmed beliefs** $(\mu_{\bar{s}_i})_{i \in I, \bar{s}_i \in \text{Supp} \bar{\sigma}_i}$. The set of (symmetric) self-confirming equilibria of (Γ, f, ϕ) is denoted by $SCE(\Gamma, f, \phi)$ ($\text{symSCE}(\Gamma, f, \phi)$).

²²See Battigalli et al. (2016b) for an in-depth analysis of the properties of feedback and how they affect the SCE set.

As in BCMM, the confirmed beliefs condition says that an agent rules out opponents' strategy distributions that are inconsistent with his "empirical distribution" of observations. More specifically, we consider stability conditions for a profile of strategy distributions in a scenario where agents drawn at random from large populations (corresponding to game roles) play the given game recurrently and learn from their personal experience. Suppose each agent keeps playing the same (pure) strategy for a very long time, and consider an agent in role i who has been playing \bar{s}_i and accumulated a large dataset of personal observations. With probability 1, and in the limit, this dataset is summarized by the frequency distribution of observations generated by his strategy \bar{s}_i and by the actual strategy distributions for the opponents' populations, $\bar{\sigma}_{-i}$, that is, $\hat{F}_i(\bar{s}_i, \bar{\sigma}_{-i})$.²³ Every profile of distributions σ_{-i} that yields the same distribution of observations is empirically indistinguishable from the true one, $\bar{\sigma}_{-i}$, and hence it cannot be objectively rejected.

We first observe that every game with feedback and ambiguity attitudes has an SCE:

Proposition 3 *For every (Γ, f, ϕ) , $SCE(\Gamma, f, \phi) \neq \emptyset$.*

Intuitively, every finite game Γ has a sequential equilibrium $(\bar{\beta}_i)_{i \in I}$ in behavioral strategies. Consider the corresponding mixed strategy profile $(\bar{\sigma}_i)_{i \in I}$ and let $\mu_{\bar{s}_i} = \delta_{\bar{\sigma}_{-i}}$ for every i and $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$. Since no ambiguity is perceived, agents with these beliefs behave as expected utility maximizers. With this, it can be shown that each $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$ is $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable, because $\bar{\sigma}$ corresponds to a sequential equilibrium.

Observe that our definition does not coincide with the one proposed in BCMM, because the rationality assumption of BCMM is given by the ex ante KMM criterion:

$$\bar{s}_i \in \arg \max_{s_i \in S_i} \int_{\Sigma_{-i}} \phi_i(U_i(s_i, \sigma_{-i})) \mu_i(d\sigma_{-i}).$$

This best reply condition is appropriate only in simultaneous moves games, possibly obtained by having agents play the strategic form of a sequential game (cf. BCMM, pp. 665-667). Here, instead, we require agents to maximize the KMM value over actions at every information set they deem reachable. Therefore, the set of self-confirming equilibria *à la* BCMM of a sequential game (Γ, f, ϕ) is $SCE(\mathcal{N}(\Gamma, f), \phi) = SCE\left((S_i, U_i, \bar{F}_i, \phi_i)_{i \in I}\right)$.

Proposition 1 implies that our definition is (realization) equivalent to the one of BCMM when agents are *ambiguity neutral*: Given $\bar{\sigma} \in SCE(\mathcal{N}(\Gamma, f), \text{Id}_{\mathbb{V}_i})$ with associated beliefs $(\mu_{s_i})_{i \in I, s_i \in \text{Supp} \bar{\sigma}_i}$, replace each $s_i \in \text{Supp} \bar{\sigma}_i$ with a realization-equivalent $(\mu_{s_i}, \text{Id}_{\mathbb{V}_i})$ -unimprovable strategy $\hat{s}_i(s_i)$, where $\hat{s}_i(\cdot)$ is a suitably defined map (see Proposition 1); then define σ_i as the pushforward of $\bar{\sigma}_i$ under map $\hat{s}_i(\cdot)$, that is, $\sigma_i = \bar{\sigma}_i \circ \hat{s}_i^{-1}$; the resulting profile σ with associated beliefs $(\mu_{\hat{s}_i(s_i)})_{i \in I, s_i \in \text{Supp} \bar{\sigma}_i}$ satisfies the SCE conditions. To sum up:

²³See Battigalli et al. (2016c) for a learning foundation of self-confirming equilibrium with non-neutral ambiguity attitudes.

Remark 3 Suppose that ϕ_i is linear for each $i \in I$. Then, for every $\bar{\sigma} \in SCE(\mathcal{N}(\Gamma, f), \phi)$ there is some $\sigma \in SCE(\Gamma, f, \phi)$ such that, for each $i \in I$, $\bar{\sigma}_i$ and σ_i are realization-equivalent.

If agents instead are ambiguity averse, the SCE's of a game are not realization-equivalent to SCE's of its strategic-form representation.

Example 4 Consider the game of Figure 1. In Example 3 we considered a belief μ_1 and ambiguity attitudes $\bar{\phi}_1$ of player 1 such that Out.M and Out.B are $(\mu_1, \bar{\phi}_1)$ -unimprovable strategies. It follows that (Out.M, L) (for instance) is a symmetric SCE of the game, where Out.M is justified by belief μ_1 (trivially confirmed for any feedback function) and L is (vacuously) justified by any belief (for any $\bar{\phi}_2$). However, note that inequality (7) implies that (Out.M, L) does not belong to $SCE(\mathcal{N}(\Gamma, f), \bar{\phi})$. Specifically, (7) implies that, for every belief μ_1 and action $a_1 \in \{T, M, B\}$, strategy Out. a_1 is not ex-ante optimal, and thus it does not satisfy the best reply condition of BCMM. \blacktriangle

Comment on knowledge of the game The definition of SCE relies on very weak interpretive assumptions about agents' knowledge of the game: each agent playing in role i has to know only his preferences (v_i, ϕ_i) , the extensive-game form (hence, also $\zeta : S \rightarrow Z$ and $\gamma : Z \rightarrow C$), and his feedback function $f_i : Z \rightarrow M$. Therefore, in a SCE an agent in population i may believe that a positive fraction of agents in population j are implementing strategies that cannot be justified by any belief, given their true preferences and feedback (v_j, ϕ_j, f_j) (possibly unknown to agents in population i). Essentially, SCE is a solution concept for incomplete information games with private values. In Section 8, we analyze a notion of rationalizable SCE that is appropriate when there is common knowledge of (Γ, f, ϕ) . Here, we just note that our definition of SCE is realization equivalent to one that replaces unimprovability with respect to a prior belief μ_i with full unimprovability with respect to a system of conditional beliefs $(\mu_i(\cdot|h))_{h \in H_i}$ (see Proposition 9 in Section 7).

6 Monotonicity of selfconfirming equilibrium

In this section we analyze changes in the set of equilibria when feedback and ambiguity attitudes are modified with respect to some baseline \bar{f} and $\bar{\phi}$ respectively. Say that the feedback profile f is *coarser* than \bar{f} if f_i is \bar{f}_i -measurable²⁴ for each $i \in I$; in other words, for each player i , the partition of Z induced by f_i is coarser than the partition of Z induced by \bar{f}_i . It is quite straightforward to show that *if f is coarser than \bar{f} then $SCE(\Gamma, f, \phi) \supseteq SCE(\Gamma, \bar{f}, \phi)$* , because coarser ex post information makes it easier to satisfy the confirmed beliefs condition (see BCMM).

We would also like to prove an extension for sequential games of the following *monotonicity theorem* of BCMM: under observable payoffs, the SCE correspondence is monotone with

²⁴That is, $f_i = \eta_i \circ \bar{f}_i$ for some $\eta_i : \bar{f}_i(Z) \rightarrow M$.

respect to ambiguity aversion, that is, as ambiguity aversion increases, the set of SCE expands. Specifically, say that (Γ, f, ϕ) features **more ambiguity aversion** than $(\Gamma, f, \bar{\phi})$ if, for each $i \in I$, $\phi_i = \varphi_i \circ \bar{\phi}_i$ for some concave and strictly increasing function φ_i . Say that (Γ, f, ϕ) features **ambiguity aversion** if each ϕ_i is concave. BCMM proved that if (Γ, f) has observable payoffs and (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$, then

$$SCE(\mathcal{N}(\Gamma, f), \phi) \supseteq SCE(\mathcal{N}(\Gamma, f), \bar{\phi}).$$

Therefore, if (Γ, f, ϕ) features **ambiguity aversion**, then

$$SCE(\mathcal{N}(\Gamma, f), \phi) \supseteq SCE(\mathcal{N}(\Gamma, f), \text{Id}_{\mathbb{V}}).$$

As already observed, self-confirming equilibria in the strategic and extensive form of the game are not realization equivalent. Hence, the monotonicity result of BCMM cannot be invoked to obtain an equivalent result for SCE in sequential games, not even in terms of induced outcome distributions. Yet, the core of the argument of BCMM can be adapted to sequential games when all the strategies in the support of an SCE with baseline ambiguity attitudes $\bar{\phi}$ are *sequentially optimal* under the confirmed beliefs that justify them, that is, at every reachable information set h the prescribed continuation strategy is the one that maximizes the value at h (see Definition 3).

Recall that $[\sigma_i]$ (respectively $[s_i]$) is the set of distributions (resp., strategies) realization equivalent to σ_i (resp. s_i), and that $\sigma'_i \in [\sigma_i]$ if and only if $\sigma'_i([s_i]) = \sigma_i([s_i])$ for every s_i .²⁵

Lemma 1 *Suppose that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$ and fix any $\bar{\sigma} \in SCE(\Gamma, f, \bar{\phi})$ justified by the confirmed beliefs $(\mu_{\bar{s}_i})_{i \in I, \bar{s}_i \in \text{Supp} \bar{\sigma}_i}$. Suppose that for each $i \in I$, every $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$ is $(\mu_{\bar{s}_i}, \bar{\phi}_i)$ -sequentially optimal. Then, there exists some $\sigma \in SCE(\Gamma, f, \phi)$ such that, for each $i \in I$, $\sigma_i \in [\bar{\sigma}_i]$ and every $s_i \in \text{Supp} \sigma_i$ is justified by confirmed belief $\mu_{\bar{s}_i}$ for some $\bar{s}_i \in [s_i] \cap \text{Supp} \bar{\sigma}_i$.*

We provide a sketch of proof of Lemma 1 because it helps to understand how dynamic (in)consistency matters for SCE analysis. Fix $\bar{\sigma} \in SCE(\Gamma, f, \bar{\phi})$ and consider a strategy $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$ that is sequentially optimal given $\mu_{\bar{s}_i}$. Since the confirmed-beliefs condition does not depend on ambiguity attitudes, we only have to argue that some $s_i \in [\bar{s}_i]$ is $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable. Fix any h consistent with \bar{s}_i , that is, $h \in H_i(\bar{s}_i)$. By ex post perfect recall, $\mu_{\bar{s}_i}$ assigns probability 1 to the set of distributions σ_{-i} such that $\sigma_{-i}(S_{-i}(h)) = \bar{\sigma}_{-i}(S_{-i}(h))$, because i observes the frequency of h . There are two cases. (1) If $\bar{\sigma}_{-i}(S_{-i}(h)) > 0$, then $\mu_{\bar{s}_i}(\Sigma_{-i}(h)) > 0$ and conditional belief $\mu_{\bar{s}_i}(\cdot|h)$ is determined by Bayes rule. Since payoffs are observable, according to conditional belief $\mu_{\bar{s}_i}(\cdot|h)$ the equilibrium action $\bar{s}_{i,h}$ is unambiguous (that is, it involves known risks), whereas deviations are untested and can be perceived as ambiguous. Thus, keeping continuation plan

²⁵See Definition 2 and Remark 1.

and beliefs fixed, an increase in ambiguity aversion from the baseline $\bar{\phi}_i$ to the more concave ϕ_i decreases the value of deviations without affecting the value of $\bar{s}_{i,h}$. Moreover, by sequential optimality, after each deviation the original continuation plan described by \bar{s}_i is optimal under $\bar{\phi}_i$ given $\mu_{\bar{s}_i}(\cdot|h)$. This implies that any $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable continuation plan makes deviations less attractive under $\bar{\phi}_i$ (if the plan involves further deviations from \bar{s}_i down the road) and hence less attractive than $\bar{s}_{i,h}$ under the higher ambiguity aversion represented by ϕ_i . (2) If $\bar{\sigma}_{-i}(S_{-i}(h)) = 0$, also $\mu_{\bar{s}_i}$ assigns probability 0 to h (that is, $h \notin H_i(\mu_{\bar{s}_i})$) and unimprovability does not impose any optimality requirement on $\bar{s}_{i,h}$ (see Definition 1). Hence, one can find a $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable and realization-equivalent strategy $s_i \in [\bar{s}_i]$. The following example illustrates this intuition.

Example 5 Let Γ' be the game of Figure 1, but with payoff 5 instead of 4 at outcome Out. The symmetric SCE (Out.M, L) of $(\Gamma, f, \bar{\phi})$ of Example 4 is also a symmetric SCE of $(\Gamma', f, \bar{\phi})$ justified by the same confirmed beliefs as in Examples 1 and 4:

$$\mu_1(\sigma_2) = \begin{cases} \frac{1}{2} & \text{if } \sigma_2 \in \{\delta_L, \delta_R\}, \\ 0 & \text{otherwise.} \end{cases}$$

Now, Out.M is not only a $(\mu_1, \bar{\phi}_1)$ -unimprovable strategy, but it is also $(\mu_1, \bar{\phi}_1)$ -sequentially optimal, because at the root the best alternative strategy In.T (see Example 3) yields an unambiguous expected payoff of $4.5 < 5$. Consider now any strictly increasing and concave transformation $\phi_1 = \varphi_1 \circ \bar{\phi}_1$ such that:

$$\begin{aligned} \phi_1(u) &= \bar{\phi}_1(u) = \sqrt{u} & \text{if } u \in [0, 9], \\ \bar{\phi}_1(u) &= 6 > \phi_1(36) > \phi_1(9) = 3. \end{aligned}$$

Specifically, $\phi_1(36)$ is close to 3. In the transformation from $\bar{\phi}_1$ to ϕ_1 , the values of In.M and In.B at the root (given μ_1) decrease, whereas the value of In.T is constant. Then, at the root, In. a_1^ϕ is worse than Out, where a_1^ϕ denotes any action that maximizes player 1's (μ_1, ϕ_1) -value conditional on $\{(In, G)\}$. This implies that Out. a_1^ϕ is (μ_1, ϕ_1) -unimprovable and $(Out.a_1^\phi, L)$ is a symmetric SCE of (Γ', f, ϕ) realization equivalent to (Out.M, L). \blacktriangle

What can go wrong when the SCE strategies are *not* sequentially optimal under the confirmed beliefs that justify them? Take the viewpoint of an agent in population i at some information set $h \in H_i(\mu_i) \cap H(s_i)$, where s_i is the agent's strategy in an SCE of $(\Gamma, \bar{\phi}, f)$ and μ_i is the confirmed belief that justifies it. At h , the agent evaluates a deviation from the equilibrium action $s_{i,h}$ to an alternative action a_i , after which he might play once more at an information set $h' \in H_i(\mu_i)$. Suppose that, given his belief and some action a'_i at h' , the deviation at h has a higher value than the SCE expected payoff. Yet, s_i is not $(\mu_i, \bar{\phi}_i)$ -sequentially optimal, the agent also realizes that his "future self" at h' will play action $s_{i,h'}$ different from a'_i , and this makes the agent prefer $s_{i,h}$ to a_i at h . But, as his

ambiguity aversion increases from $\bar{\phi}_i$ to ϕ_i , the future self of the agent may switch from $s_{i,h'}$ to a'_i at h' for all the confirmed beliefs that justify $s_{i,h}$ under $\bar{\phi}_i$. Then, although for any fixed belief and action at h' the value of a_i compared to $s_{i,h}$ at h decreases (because a_i exposes the agent to ambiguity while $s_{i,h}$ does not, and the agent has become more ambiguity averse), the value of a_i at h under the predicted choice at h' can increase when moving from $\bar{\phi}_i$ to ϕ_i . The following example demonstrates this possibility.

Example 6 Consider the game of Figure 1. Let $\bar{\phi}_1$ and ϕ_1 be the ambiguity attitudes described in Examples 4 and 5. In Appendix 6 we show that, if $\phi(36)$ is sufficiently small, Out does not belong to the set of SCE outcomes of (Γ, ϕ, f) (although it does for $(\Gamma, f, \bar{\phi})$). The intuition is as follows. For an intermediate level of ambiguity aversion, captured by the baseline second-order utility $\bar{\phi}_1$, upon reaching (In, G) player 1 is tempted by actions M and B even under the most pessimistic belief ($\mu_1(\cdot | \{(In, G)\}) = \frac{1}{2}\delta_L + \frac{1}{2}\delta_R$, cf. Lemma 9). At the root, he anticipates this and, scared by the implied ex-ante objective expected reward of $\frac{1}{2}$ under the “bad model” (δ_L if he plans M and δ_R if he plans B), he chooses Out. For a higher level of ambiguity aversion, captured by ϕ_1 , at (In, G) player 1 is tempted by M and B only for sufficiently optimistic beliefs. As a consequence, at the root, he is less worried by the “bad model,” either because now he plans the unambiguous action T for the subgame, or because he plans an ambiguous action and he deems the “bad model” sufficiently unlikely. Therefore, he chooses In at the root. \blacktriangle

Thus, the monotonicity result of BCMM does not extend to all sequential games, even if we restrict our attention to distributions of outcomes.²⁶ Yet, we can use Lemma 1 to show that the monotonicity result holds for classes of games and equilibria of interest.

6.1 No player moves more than once

We say that **no player moves more than once**²⁷ in Γ if, for every $i \in I$ and $z \in Z$, there is at most one information set $h \in H_i$ that contains a predecessor of z . In this class of games, at any information set $h \in H_i$, the agent does not move again after h ; thus, no information set of i is prevented by any strategy of i . In this case the value of an action at $h \in H_i$ does not depend on i 's strategy, unimprovability coincides with sequential optimality,²⁸ strategies (or strategy distributions) are realization-equivalent if and only if they coincide, hence Lemma 1 implies the following result:

Corollary 4 Fix two games with observable payoffs where no player moves more than once, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$.

²⁶This shows that the conjecture informally stated by BCMM (p. 667) is false.

²⁷Perfect information games with this property have been called “simple” (see, e.g., Fudenberg and Levine, 2006). We prefer the more explicative phrase because (i) we are not assuming perfect information and (ii) “simple” is vague and had been used earlier with a different meaning (see Battigalli, 2003).

²⁸Indeed, Proposition 2 can be extended to any ϕ_i in one-move games; yet, Proposition 1 cannot!

Then

$$SCE(\Gamma, f, \phi) \supseteq SCE(\Gamma, f, \bar{\phi}).$$

6.2 Symmetric self-confirming equilibria

Consider now a *symmetric* SCE \bar{s} in a game with observable payoffs and *without chance moves*. We can show that there is a symmetric SCE equilibrium s with the same outcome.²⁹

Theorem 5 *Fix two games with observable payoffs, ambiguity aversion, and without chance moves, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then:*

$$\zeta(\text{symSCE}(\Gamma, f, \phi)) \supseteq \zeta(\text{symSCE}(\Gamma, f, \bar{\phi})).$$

We prove Theorem 5 in the Appendix, here we provide an intuitive argument. The SCE strategies may not be sequentially optimal under the confirmed beliefs that justify them. However, consider alternative beliefs that are supported by Dirac models and give the same predictive beliefs as the original ones. By construction, these beliefs are confirmed by the equilibrium play, and they feature two additional properties. First, they are the most pessimistic beliefs among those that give rise to the same predictive probabilities; thus, by certainty of the equilibrium payoff, they justify a realization equivalent symmetric SCE. This is shown by Lemma 9 in the Appendix and it is based on the following intuition: A belief supported by Dirac models is a “mean-preserving spread” of the objective expected payoffs that the agent deems possible, and the concavity of ϕ_i implies that player i is averse to such spread. Second, absent chance moves, beliefs supported by Dirac models cannot entail dynamic inconsistencies of preferences over strategies; thus, the symmetric SCE’s they justify features sequentially optimal strategies. This is shown by Lemma 8 in the Appendix. The main idea behind this result is that dynamic inconsistency is due to the perception of hedging opportunities that may be optimal ex-ante, but would not be implemented ex-post after (partial) resolution of the uncertainty. But, absent chance moves, an agent whose belief is supported by Dirac models does not perceive such hedging opportunities. With this, we can use Lemma 1 to prove the monotonicity result for equilibrium outcomes.

Absence of chance moves and symmetry of the equilibrium are tight conditions. In Example 6, outcome Out is induced by a symSCE of the game with chance moves $(\Gamma, f, \bar{\phi})$, but it is not induced by any symSCE of the game (Γ, f, ϕ) , which features more ambiguity

²⁹Note that we can identify $\text{symSCE}(\Gamma, f_i, \bar{\phi})$ with a subset of S . Hence, it makes sense to write $\zeta(\text{symSCE}(\Gamma, f_i, \bar{\phi}))$.

aversion. As for the role of symmetry (pure equilibrium), consider the following example.

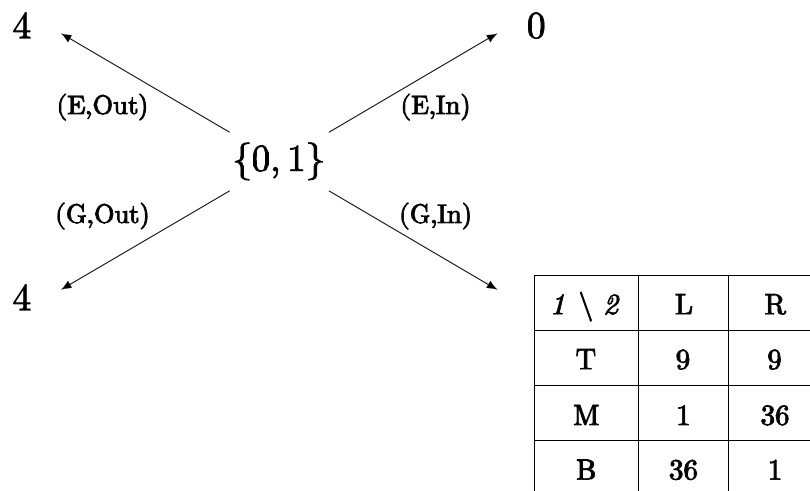


Figure 2: A 3-person common interest game: $I = \{0, 1, 2\}$

Example 7 Let Γ' be the following modification of the game of Figure 1: deviating from our standard notation, here 0 is not chance but rather an actual player, choosing between E and G simultaneously with player 1 at the root. Assuming common interests, all players obtain payoff 4 at both outcomes (E, Out) and (G, Out). See Figure 2. Assume perfect feedback and consider the same $\bar{\phi}_1$ and ϕ_1 of Example 6. The SCE $\bar{\sigma} = (\sigma_0, \text{Out.M, L})$ of $(\Gamma', f, \bar{\phi})$ yields outcomes (E, Out) and (G, Out) with probability $\frac{1}{2}$ (cf. Example 4). By perfect feedback and confirmed beliefs, in any realization-equivalent SCE player 1 is certain of σ_0 . Yet, for the same argument as for Example 6, no (μ_1, ϕ_1) -unimprovable strategy of player 1 prescribes action Out when the marginal of μ_1 on Σ_0 is σ_0 . Thus, no SCE of (Γ', f, ϕ) yields the same outcome distribution as $\bar{\sigma}$. \blacktriangle

6.3 Ambiguity aversion versus ambiguity neutrality

Fix a sequential game with feedback (Γ, f) and let $\text{Id}_{\mathbb{V}} = (\text{Id}_{\mathbb{V}_i})_{i \in I}$ denote the profile of players' identity functions characterizing their neutrality toward ambiguity. We relate the set $\text{SCE}(\Gamma, f, \text{Id}_{\mathbb{V}})$ of SCE's of Γ given feedback functions f and neutral ambiguity attitudes with the set $\text{SCE}(\Gamma, f, \phi)$ with non-neutral ambiguity attitudes ϕ . We start with a preliminary observation:³⁰

Remark 4 For every two-person game with feedback (Γ, f) and every profile ϕ of second-order utility functions, $\text{SCE}(\Gamma, f, \text{Id}_{\mathbb{V}}) \subseteq \text{SCE}(\Gamma, f, \phi)$.

³⁰Cf. footnote 23 of BCMM.

To see this, note that ambiguity neutrality and the convexity of $\Sigma_{-i} = \Delta(S_{-i})$ in two-person games allow to replace the justified beliefs supporting $\bar{\sigma}$ as an SCE of $(\Gamma, f, \text{Id}_{\mathbb{V}})$ with the corresponding Dirac beliefs supported by their predictive measure. Since ambiguity attitudes are immaterial for agents with Dirac beliefs, $\bar{\sigma}$ is also an SCE of (Γ, f, ϕ) . This argument does not hold with $n > 2$ players, because in this case Σ_{-i} is not convex, hence, the Dirac measure supported by the predictive of $\mu_{\bar{\sigma}_i}$ may belong to $\Delta(\Delta(S_{-i})) \setminus \Delta(\Sigma_{-i})$. Nonetheless, we can relate $SCE(\Gamma, f, \text{Id}_{\mathbb{V}})$ and $SCE(\Gamma, f, \phi)$ for a large class of games.

As a corollary of their main monotonicity result, BCMM show that, under observable payoffs, $SCE(\mathcal{N}(\Gamma, f), \text{Id}_{\mathbb{V}})$ is contained in $SCE(\mathcal{N}(\Gamma, f), \phi)$ if each ϕ_i is concave and there are observable payoffs. Even if the main monotonicity result of BCMM does not extend to SCE of sequential games for the entire spectrum of ambiguity attitudes, we are still able to obtain a sequential version of this corollary in terms of induced outcome distributions. By dynamic consistency under ambiguity neutrality (i.e., by Proposition 2), the result about the comparison between ambiguity neutrality and ambiguity aversion is a corollary of Lemma 1. Define the function

$$\begin{aligned} \hat{\zeta} : \Sigma &\rightarrow \Delta(Z) \\ \bar{\sigma} &\mapsto \hat{\zeta}(\bar{\sigma})(z) = \sum_{(s_0, s) : \zeta(s_0, s) = z} \sigma_0(s_0) \cdot \bar{\sigma}(s), \end{aligned}$$

where Σ is the set of product measures on S . This is the pushforward map that gives for every (product) distribution over strategy profiles $\bar{\sigma}$ the corresponding probability distribution $\hat{\zeta}(\bar{\sigma})$ on terminal nodes.

Corollary 6 *Suppose that (Γ, f) has observable payoffs and (Γ, f, ϕ) features ambiguity aversion. Then, the set of SCE distributions over terminal nodes of (Γ, f, ϕ) contains the set of SCE distributions over terminal nodes with ambiguity neutrality:*

$$\hat{\zeta}(SCE(\Gamma, f, \phi)) \supseteq \hat{\zeta}(SCE(\Gamma, f, \text{Id}_{\mathbb{V}})).$$

Given the large body of empirical evidence supporting the ambiguity aversion hypothesis, we conclude that the standard SCE concept, which implicitly assumes neutral ambiguity attitudes, overestimates the predictability of long-run outcomes of learning dynamics.

Finally, note that Corollary 4 and Corollary 6, via Lemma 1, imply that all beliefs that justify $SCE(\Gamma, f, \bar{\phi})$ in one-move games or under ambiguity neutrality (linear $\bar{\phi}$) also justify the corresponding equilibria under ϕ . This is not true for symmetric equilibria in games without chance moves (see the proof of Theorem 5).

7 Conditional probability systems and full unimprovability

So far we studied how ambiguity aversion and the ensuing possibility of dynamic inconsistency (the incentive to covertly commit, if possible) affect SCE analysis. However, we

neglected strategic reasoning based on common knowledge of (some features of) the game. To analyze strategic reasoning in sequential games we have to address an additional and more traditional issue: We need to model how a player thinks that the other agents would react to unexpected moves. Even if players are ambiguity neutral, the analysis of Sections 4-6 is insufficient to address this issue: As a preliminary step, we need to assume that agents have well defined conditional beliefs at all information sets, including the unexpected ones. The conditional beliefs of any given agent at different information sets have to be mutually consistent if we want to preserve/extend the unimprovability principle. Moreover, we want to model players who reason strategically about the game before playing it. For this purpose, we have to consider also their beliefs at the root, even if they are not first movers. Therefore, for every $i \in I$, we will consider the expanded collection of information sets $\bar{H}_i = H_i \cup \{\{\emptyset\}\}$.

To simplify the analysis of this section and the following one we focus on games *without chance moves*. Thus, the outcome function is $\zeta : S \rightarrow Z$, the strategic-form payoff feedback and payoff functions of player i are $F_i = f_i \circ \zeta : S \rightarrow M$ and $U_i = u_i \circ \zeta : S \rightarrow \mathbb{R}$.³¹

To understand the following definition, consider beliefs $\mu_i(\cdot|g), \mu_i(\cdot|h) \in \Delta(\Sigma_{-i})$ at two information sets g and h , so that g precedes h ($g \prec h$), and suppose that h is possible according to $\mu_i(\cdot|g)$, that is, $p_{\mu_i}(S_{-i}(h)|g) > 0$. Then, the conditional belief $\mu_i(\cdot|h)$ can be derived from $\mu_i(\cdot|g)$ in the way prescribed by (5): for all $E_{-i} \in \mathcal{B}(\Sigma_{-i})$,

$$\mu_i(E_{-i}|h) = \frac{\int_{E_{-i} \cap \Sigma_{-i}(g)} \sigma_{-i}(S_{-i}(h)|g) \mu_i(d\sigma_{-i}|g)}{p_{\mu_i}(S_{-i}(h)|g)}.$$

Next note that we can write the required relation between $\mu_i(\cdot|g)$ and $\mu_i(\cdot|h)$ without explicitly stating condition $p_{\mu_i}(S_{-i}(h)|g) > 0$: for all $g, h \in \bar{H}_i$ with $g \prec h$, and $E_{-i} \in \mathcal{B}(\Sigma_{-i})$

$$\mu_i(E_{-i}|h) p_{\mu_i}(S_{-i}(h)|g) = \int_{E_{-i} \cap \Sigma_{-i}(g)} \sigma_{-i}(S_{-i}(h)|g) \mu_i(d\sigma_{-i}|g). \quad (8)$$

Definition 6 A *conditional probability system (CPS) on (Σ_{-i}, \bar{H}_i)* is an array of probability measures

$$\mu_i(\cdot|\cdot) = (\mu_i(\cdot|h))_{h \in \bar{H}_i} \in [\Delta(\Sigma_{-i})]^{\bar{H}_i}$$

such that

- (1) for all $h \in \bar{H}_i$, $\mu_i(\Sigma_{-i}(h)|h) = 1$,
- (2) for all $g, h \in \bar{H}_i$ with $g \prec h$ eq. (8) holds.

Definition 7 A *CPS on (S_{-i}, \bar{H}_i)* is an array of probability measures

$$p_i(\cdot|\cdot) = (p_i(\cdot|h))_{h \in \bar{H}_i} \in [\Delta(S_{-i})]^{\bar{H}_i}$$

³¹The analysis can be extended to games with chance moves, but some of the notation and some of the proofs would be more complex.

such that $(p_i(\cdot|h))_{h \in \bar{H}_i} = (p_{\mu_i(\cdot|h)})_{h \in \bar{H}_i}$ for some CPS $\mu_i(\cdot|\cdot) = (\mu_i(\cdot|h))_{h \in \bar{H}_i}$ on (Σ_{-i}, \bar{H}_i) .

In Definition 7, we define a CPS on (S_{-i}, \bar{H}_i) indirectly as the “predictive” of some CPS on (Σ_{-i}, \bar{H}_i) . We could have given a direct definition: Indeed, $p_i \in [\Delta(S_{-i})]^{\bar{H}_i}$ is a CPS on (S_{-i}, \bar{H}_i) if and only if (1) for every $h \in \bar{H}_i$, $p_i(S_{-i}(h)|h) = 1$ and (2) for all $g, h \in \bar{H}_i$ with $g \prec h$ and all $s_{-i} \in S_{-i}(h)$,

$$p_i(s_{-i}|h)p_i(S_{-i}(h)|g) = p_i(s_{-i}|g). \quad (9)$$

Notation 7 The set of CPS’s on (Σ_{-i}, \bar{H}_i) $[(S_{-i}, \bar{H}_i)]$ is denoted by $\Delta^{\bar{H}_i}(\Sigma_{-i})$ $[\Delta^{\bar{H}_i}(S_{-i})]$.

With this, we can give a stronger definition of unimprovability:

Definition 8 A strategy s_i is **fully** $(\mu_i(\cdot|\cdot), \phi_i)$ -**unimprovable** (where $\mu_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$) if

$$\forall h \in H_i, s_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; s_i, \mu_i, \phi_i).$$

It is well known that, for ambiguity neutral agents, we have a refined dynamic programming result:³²

Proposition 8 For every $s_i^* \in S_i$ and $\mu_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$, the following are true:

(1*) $[s_i^*]$ contains a fully $(\mu_i(\cdot|\cdot), \text{Id}_{V_i})$ -unimprovable strategy if and only if

$$\forall h \in \bar{H}_i(s_i^*), s_i^* \in \arg \max_{s_i \in S_i(h)} \sum_{s_{-i} \in S_{-i}(h)} U_i(s_i, s_{-i}) p_{\mu_i}(s_{-i}|h);$$

(2*) s_i^* is a fully $(\mu_i(\cdot|\cdot), \text{Id}_{V_i})$ -unimprovable strategy if and only if

$$\forall h \in \bar{H}_i, (s_{i|h}^*, s_{i,h}^*) \in \arg \max_{s_i \in S_i(h)} \sum_{s_{-i} \in S_{-i}(h)} U_i(s_i, s_{-i}) p_{\mu_i}(s_{-i}|h).$$

Full-improvability solves the difficulties described above. Intuitively, we can interpret a $(\mu_i(\cdot|\cdot), \phi_i)$ -fully unimprovable strategy s_i^* as the **plan** of an agent in role i , which can be obtained with a “folding back” dynamic programming procedure on the subjective decision tree implied by CPS $\mu_i(\cdot|\cdot)$. In our perspective, the fact that the beliefs of an agent with perfect recall are given by a CPS reflects the **epistemic unity of the agent’s self**: the agent always incorporates new information into his system of knowledge and beliefs in a way that is consistent with his previous beliefs and with the rules of conditional probabilities, even when new information follows previously unexpected information.

Before we move on to model strategic reasoning, we verify that even if SCE is strengthened by requiring full unimprovability, the set of possible outcomes does not change.

³²Definition 8 requires maximization over actions; hence, it considers only the information sets where i is active. The following propositions relate to maximization over strategies; hence, they also include (ex ante) maximization at $\{\emptyset\}$ even if i is not a first mover. Also, recall that, according to our notation, $(s_{i|h}^*, s_{i,h}^*)$ is the minimal modification of s_i^* that makes h reachable and plays like s_i^* at all information sets h' that do not strictly precede h .

Definition 9 A *fully unimprovable self-confirming equilibrium* of (Γ, f, ϕ) is a profile of strategy distributions $(\sigma_i^*)_{i \in I}$ with following property: For every $i \in I$ and $s_i^* \in \text{Supp}\sigma_i^*$ there is a CPS $\mu_{s_i^*}(\cdot|\cdot) \in \Delta^{\hat{H}_i}(\Sigma_{-i})$ such that
 (rationality) s_i^* is fully $(\mu_{s_i^*}(\cdot|\cdot), \phi_i)$ -unimprovable,
 (confirmed beliefs) $\mu_{s_i^*} \left(\left\{ \sigma_{-i} \in \Sigma_{-i} : \hat{F}_i(s_i^*, \sigma_{-i}) = \hat{F}_i(s_i^*, \sigma_{-i}^*) \right\} \mid \{\emptyset\} \right) = 1$.

Note that the confirmed beliefs condition refers only to the initial beliefs $\mu_{s_i^*}(\cdot|\{\emptyset\})$ because it implies that conditional beliefs are confirmed by observed conditional frequencies at every history that is reached with positive probability in equilibrium (see the proof of Lemma 4). By inspection of Definitions 5 and 9 it is also clear that every fully unimprovable SCE is also an SCE. The following proposition implies that SCE and fully unimprovable SCE are realization equivalent.

Proposition 9 For every SCE there is a corresponding fully unimprovable SCE that yields the same probability distribution on terminal nodes.

Intuitively, since SCE does not rely on complete information and does not model strategic thinking, requiring rational reactions to unexpected moves adds little to the analysis: Agents are not assumed to know the preferences of others, hence they are not assumed to rule out irrational reactions to deviations.

8 Knowledge of the game and strategic reasoning

What if (Γ, f, ϕ) (or a part of it) is common knowledge? Then it makes sense to explore a notion of “rationalizable SCE” according to which agents reason strategically about their opponents taking into account their preferences and feedback functions (cf. Rubinstein and Wolinsky 1994, Battigalli 1999, Dekel *et al.* 1999, Esponda 2013, Fudenberg and Kamada 2015). We illustrate this with two simple examples.

Example 8 (Cf. Fudenberg and Levine, 1993). Figure 2 depicts the so called “Entry game” often used to illustrate the shortcomings of the Nash equilibrium concept. To complete the specification of (Γ, f, ϕ) assume observable payoffs (hence, in this case, perfect feedback) and let ϕ be an arbitrary pair of strictly increasing functions.³³ It can be checked that the set of SCE’s of (Γ, f, ϕ) is

$$\{\sigma^* \in \Sigma : \sigma_1^*(A) = 1, \text{ or } (\sigma_1^*(A) < 1, \sigma_2^*(L) = 1)\}.$$

First-movers who choose A (go Across) get no feedback and can thus hold trivially confirmed beliefs that make them play A. As for second-movers, their plan becomes relevant only if a

³³One can show that ambiguity aversion does not matter when players have only two strategies.




Figure 2: Entry game

positive fraction of first-movers choose D (go Down). In that case, the only rational choice is L (Left). In equilibrium, first-movers going Down correctly predict Left, which makes Down a strict best response. How can a first mover expect that the second mover goes Right with high probability (i.e., that a large fraction of second movers play Right)? Informally, this is possible if either the first mover gives a high probability to the second mover being irrational, or—more reasonably—if the first mover does not know v_2 . If instead the first mover knows v_2 and believes in the rationality of the second mover, then he predicts that Down would be followed by Left, and he would go Down. Intuitively, only one SCE is consistent with belief in rationality and knowledge of the game: (D,L). However, this is not formally captured by the confirmed beliefs condition of Definition 5. According to such definition, if an agent in population 1 does not play Down, he can keep the belief that the second mover would play Right after Down: If v_2 is unknown to first movers, they have no way to understand that only Left is rational. Relatedly, the set of equilibrium distributions on terminal nodes is unchanged if we adopt the stronger definition of unimprovability: The latter implies $\sigma_2^*(L) = 1$, but the first movers who go Across because they have wrong beliefs have no way to find out they are wrong; hence, the equilibrium distributions on terminal nodes are the same. ▲

Example 9 (Rubinstein and Wolinsky 1994, Battigalli 1999). Two players must simultaneously choose a location among the equally spaced points 0, 1, 2, 3 on the real line. Their payoff is the negative of the distance between their chosen locations and each player, that is, $U_i(s_1, s_2) = -|s_1 - s_2|$; furthermore, each player observes this distance and, of course, remembers his action: $F_i(s_1, s_2) = (s_i, |s_1 - s_2|)$. Again, let us fix an arbitrary concave ϕ and—for simplicity—let us focus on the symmetric equilibria. The set of symmetric SCE's is

$$\text{symSCE} = \{(0, 0), (1, 1), (2, 2), (3, 3), (1, 2), (2, 1)\}.$$

Of course, (1, 2) and (2, 1) are not Nash equilibria. They are symmetric SCE's supported by confirmed beliefs according to which each agent assigns the same probability to the co-player being either on his right, or on his left. However, if the game is common knowledge, an agent with such beliefs infers that with probability .5 his co-player is “cornered” at an extreme point and cannot be best responding to confirmed beliefs. This is unlikely to be a stable situation under complete information if players reason strategically, because no player has reasons to believe that the co-player would keep playing in the same way. ▲

Symmetrically rationalizable SCE We capture with an inductive definition the behavioral consequences of the following assumptions on rationality and interactive beliefs: (1) agents are rational (in the sense of full unimprovability) and have confirmed beliefs, and (2) there is common belief at the beginning of the game that (1) holds. This is easier to do in a particular case, i.e., when all agents in the same population follow the same plan (symmetric SCE) and there is common full belief of this.³⁴ In this case, the CPS $\mu_i(\cdot|\cdot)$ on (Σ_{-i}, \bar{H}_i) is supported by Dirac models, and therefore it is isomorphic to a CPS p_{μ_i} on (S_{-i}, \bar{H}_i) . A variation of the algorithm defined by Battigalli (1999) (see also Esponda 2013) precisely captures the foregoing epistemic assumptions. As in Section 7, also in this section we assume for simplicity that there are no chance moves.

Given a CPS $\mu_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$, we let

$$r_i(\mu_i(\cdot|\cdot), \phi_i) = \left\{ s_i \in S_i : \forall h \in H_i, s_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; s_i, \mu_i(\cdot|\cdot), \phi_i) \right\}$$

denote the set of fully (μ_i, ϕ_i) -unimprovable strategies of i . When $\mu_i(\cdot|\cdot)$ is isomorphic to a predictive CPS $p_i(\cdot|\cdot)$, as in the case we are considering right now, it makes sense to write $r_i(p_i(\cdot|\cdot), \phi_i)$.

Definition 10 For each $i \in I$, let $B_i^0 = S_i \times M$, and

$$B_i^{k+1} = \left\{ (\bar{s}_i, \bar{m}_i) \in B_i^k : \begin{array}{l} \exists p_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(S_{-i}), \bar{s}_i \in r_i(p_i(\cdot|\cdot), \phi_i), \\ p_i(F_{\bar{s}_i}^{-1}(\bar{m}_i) \cap \{s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k\} | \{\emptyset\}) = 1 \end{array} \right\}$$

for each $k \in \mathbb{N}_0$. A strategy profile \bar{s} is a **symmetrically rationalizable SCE** for game (Γ, f, ϕ) (without chance moves) if $(\bar{s}_i, F_i(\bar{s}))_{i \in I} \in \times_{i \in I} \bigcap_{k \in \mathbb{N}} B_i^k$. The set of symmetrically rationalizable self-confirming equilibria of (Γ, f, ϕ) is denoted by $\text{symRSCE}(\Gamma, f, \phi)$.

Intuitively, $(\bar{s}_i, \bar{m}_i) \in B_i^1$ if \bar{s}_i is justified by some CPS such that i is initially certain to get message \bar{m}_i if he plays \bar{s}_i . Thus, $(\bar{s}_i, F_i(\bar{s}))_{i \in I} \in \times_{i \in I} \bigcap_{k \in \mathbb{N}} B_i^1$ if \bar{s} is an unimprovable

³⁴Note, we did not say “belief at the beginning of the game,” because now we mean “probability 1 belief conditional on every information set.” This is called “full belief” in epistemic game theory (e.g., Battigalli et al. 2017).

SCE (with beliefs supported by Dirac models). Then, $(\bar{s}_i, \bar{m}_i) \in B_i^2$ if \bar{s}_i is justified by some CPS such that i is initially certain to get message \bar{m}_i if he plays \bar{s}_i , and furthermore he is initially certain that everybody else is playing strategies justified by confirmed beliefs. The iterations capture higher levels of (initial) mutual belief in rationality and in confirmation of beliefs.

Remark 5 A strategy profile \bar{s} is a **symmetrically rationalizable SCE** for (Γ, f, ϕ) if and only if there is a profile of finite subsets $(\bar{B}_i)_{i \in I} \in \times_{i \in I} 2^{S_i \times M}$ such that, for every $i \in I$, there is $\bar{m}_i \in M$ with $(\bar{s}_i, \bar{m}_i) \in \bar{B}_i$, and for every $(\hat{s}_i, \hat{m}_i) \in \bar{B}_i$, there is $p_i(\cdot | \cdot) \in \Delta^{\bar{H}_i}(S_{-i})$ with $\hat{s}_i \in r_i(p_i(\cdot | \cdot), \phi_i)$ and

$$p_i(F_{\bar{s}_i}^{-1}(\hat{m}_i) \cap \{s_{-i} : (s_j, F_j(\hat{s}_i, s_{-i}))_{j \neq i} \in \bar{B}_{-i}\} | \{\emptyset\}) = 1. \quad (10)$$

The Remark above shows that our inductive definition is an extensive-form version with ambiguity attitudes of Rubinstein's and Wolinsky's (1994) rationalizable conjectural equilibrium. Next we offer a characterization with sets of strategy profiles. Let $\text{symSCE}^0(\Gamma, f, \phi) = S$ and

$$\text{symSCE}^{k+1}(\Gamma, f, \phi) = \left\{ \bar{s} \in \text{symSCE}^k(\Gamma, f, \phi) : \begin{array}{l} \forall i \in I, \exists p_i(\cdot | \cdot) \in \Delta^{\bar{H}_i}(S_{-i}), \bar{s}_i \in r_i(p_i(\cdot | \cdot), \phi_i), \\ p_i(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} | \{\emptyset\}) = 1 \end{array} \right\},$$

where, for any subset $X \subseteq S$ and strategy $\bar{s}_i \in S_i$,

$$X_{\bar{s}_i} = \{s_{-i} \in S_{-i} : (\bar{s}_i, s_{-i}) \in X\}$$

is the section of X at \bar{s}_i (thus, $\text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i}$ is the section of $\text{symSCE}^k(\Gamma, f, \phi)$ at \bar{s}_i). Note that $\text{symSCE}^1(\Gamma, f, \phi)$ coincides with the set of fully unimprovable symmetric SCE's of (Γ, f, ϕ) justified by confirmed beliefs supported by Dirac models. Thus, $\bar{s} \in \text{symSCE}^2(\Gamma, f, \phi)$ if each \bar{s}_i is a best reply to a confirmed belief that assigns probability 1 to other players choosing best replies to confirmed beliefs supported by Dirac models, given \bar{s}_i and message $m_i = F_i(\bar{s})$.

The following result shows that symmetric RSCE is characterized by the iterated deletion of strategy profiles $(\text{symSCE}^k(\Gamma, f, \phi))_{k \in \mathbb{N}}$:

Lemma 2 For every $k \in \mathbb{N}$, and $\bar{s} \in S$, $\bar{s} \in \text{symSCE}^k(\Gamma, f, \phi)$ if and only if $(\bar{s}_i, F_i(\bar{s}))_{i \in I} \in B^k$. Therefore,

$$\bigcap_k \text{symSCE}^k(\Gamma, f, \phi) = \text{symRSCE}(\Gamma, f, \phi).$$

Example 10 One can easily check that the only rationalizable SCE for the game of Figure 2 is (D,L), independently of ϕ : Only L is fully unimprovable for player 2, thus $\text{symSCE}^1(\Gamma, f, \phi) = \{(A, L), (D, L)\}$ and $\text{symSCE}^1(\Gamma, f, \phi)_A = \{L\}$. Therefore,

$$\text{symSCE}^2(\Gamma, f, \phi) = \{(D, L)\}.$$

▲

Example 11 Now consider the rationalizable SCE's set for the distance game of Example 9:

$$\text{symSCE}^1(\Gamma, f, \phi) = \text{symSCE}(\Gamma, f, \phi) = \{(0, 0), (1, 1), (2, 2), (3, 3), (1, 2), (2, 1)\}.$$

Let $\bar{s} = (1, 2)$. We show that $\bar{s} \notin \text{symSCE}^2(\Gamma, f, \phi)$. Consider player 1. If we had $\bar{s} \in \text{symSCE}^2(\Gamma, f, \phi)$, then $\bar{s}_1 = 1$ would be a best reply to a belief p_1 such that $p_1(\{0, 2\} \cap \text{symSCE}^1(\Gamma, f, \phi)_{\bar{s}_2} | \{\emptyset\}) = 1$. But $\text{symSCE}^1(\Gamma, f, \phi)_{\bar{s}_2} = \{1, 2\}$. Thus

$$\{0, 2\} \cap \text{symSCE}^1(\Gamma, f, \phi)_{\bar{s}_2} = \{2\},$$

and the best reply to $s_2 = 2$ is $s_1 = 2$. A similar argument shows that $(2, 1) \notin \text{symSCE}^1(\Gamma, f, \phi)$. Hence only the pure Nash equilibria are symRSCE. \blacktriangle

We can prove for the symRSCE correspondence two monotonicity results analogous to those obtained for the symSCE correspondence (cf. Corollary 4 adapted to symSCE and Theorem 5):

Theorem 10 Fix two games with observable payoffs and no chance moves where no player moves more than once, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then,

$$\text{symRSCE}(\Gamma, f, \bar{\phi}) \subseteq \text{symRSCE}(\Gamma, f, \phi).$$

Theorem 11 Fix two games with ambiguity aversion, observable payoffs, and no chance moves, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then, for every $\bar{s} \in \text{symRSCE}(\Gamma, f, \bar{\phi})$ there is some $s \in \text{symRSCE}(\Gamma, f, \phi)$ such that, for each player i , s_i is realization equivalent to \bar{s}_i ; therefore,

$$\zeta(\text{symRSCE}(\Gamma, f, \bar{\phi})) \subseteq \zeta(\text{symRSCE}(\Gamma, f, \phi)).$$

Intuitively, these results rely on the following intermediate step:³⁵

Monotonicity of the justifiability correspondence: For every information set $h \in H_i$ and compact subset $\Sigma_{i|h} \subseteq \Delta(S_{-i}(h))$, if an action a_i is justified as a best reply to some conditional belief $\bar{p}_{i,h} \in \Delta(\Sigma_{i|h})$ given baseline ambiguity attitudes $\bar{\phi}_i$, then there is some $p_{i,h} \in \Delta(\Sigma_{i|h})$ that justifies a_i as a best reply given the more ambiguity averse attitudes ϕ_i . This holds, in particular, when $\Sigma_{i|h} = \{\delta_{s_{-i}} : s_{-i} \in S_{-i}(h)\}$ is the set of Dirac distributions in $\Delta(S_{-i}(h))$.

Then we can show that if a strategy is fully $(\bar{p}_i, \bar{\phi}_i)$ -unimprovable, then it is also fully (p_i, ϕ_i) -unimprovable for some suitably chosen belief system p_i that coincides with \bar{p}_i on the equilibrium path (by the usual argument that equilibrium actions are unambiguous, hence deviations become less appealing), and can be chosen off the equilibrium path invoking monotonicity of the justifiability correspondence.

³⁵See Lemma 10 and its corollary in the Appendix; cf. Battigalli et al. (2016a) and Weinstein (2016).

Rationalizable SCE We now define rationalizable SCE for population games, that is, we consider the general non-symmetric version of rationalizable SCE. By analogy with Lemma 2, we perform an iterated deletion of distributions of strategy profiles.

To ease notation, let

$$\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) = \left\{ \sigma_{-i} \in \Sigma_{-i} : \hat{F}_i(s_i, \sigma_{-i}) = \hat{F}_i(s_i, \bar{\sigma}_{-i}) \right\}$$

denote the partially identified set of co-players strategy distributions observationally equivalent for i to $\bar{\sigma}_{-i}$ given s_i . With this, let $SCE^0(\Gamma, f, \phi) = \times_{i \in I} \Delta(S_i)$, and

$$SCE^{k+1}(\Gamma, f, \phi) = \left\{ \bar{\sigma} \in SCE^k(\Gamma, f, \phi) : \begin{array}{l} \forall i \in I, \forall s_i \in \text{Supp} \bar{\sigma}_i, \exists \mu_{s_i}(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i}), s_i \in r_i(\mu_{s_i}(\cdot|\cdot), \phi_i), \\ \mu_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \phi) \mid \{\emptyset\} \right) = 1 \end{array} \right\}$$

for every $k \in \mathbb{N}_0$.³⁶ Note that $SCE^1(\Gamma, f, \phi)$ is the set of fully unimprovable SCE's of (Γ, f, ϕ) .

Definition 11 A profile of strategy distributions $\bar{\sigma} \in \Sigma$ is a rationalizable self-confirming equilibrium for (Γ, f, ϕ) if

$$\bar{\sigma} \in \bigcap_{k \in \mathbb{N}} SCE^k(\Gamma, f, \phi).$$

The set of rationalizable self-confirming equilibria of (Γ, f, ϕ) is denoted by $RSCE(\Gamma, f, \phi)$.

The most important difference with respect to the definition of rationalizable symmetric self-confirming equilibrium is best understood by looking at the second step, where agents check if it is possible that others are best responding to confirmed beliefs. If there is common belief of symmetry, an agent playing strategy s_i is certain that all the other agents in population i also use s_i , and this is taken into account when he checks whether the agents in co-players' populations are best responding to confirmed beliefs. If instead it is understood that different agents in population i may use different strategies, an agent playing s_i may think that only a negligible fraction of other agents in i is playing s_i , therefore, the fact that he is playing s_i does not enter this calculation. Essentially, we are assuming that each agent has a belief about the whole profile of distributions, that is, a belief over $\Sigma = \times_{j \in I} \Delta(S_j)$. For the purpose of computing best replies, only the marginal over Σ_{-i} matters. But in order to check whether everybody is best replying to confirmed

³⁶Note the following slight abuse of notation. If, for some $k \in \mathbb{N}_0$, $s_i \in S_i$, $\bar{\sigma}_{-i} \in \Sigma_{-i}$ the set $\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\phi)$ is not measurable, than the requirement

$$\mu_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\phi) \mid \{\emptyset\} \right) = 1$$

means that there is a measurable set $\Sigma' \subset \left(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\phi) \right)$ with $\mu_{s_i}(\Sigma' \mid \{\emptyset\}) = 1$.

beliefs, the distribution of strategies in population i is crucial (see Fudenberg and Kamada, 2015, 2017).

Say that Γ (hence also (Γ, f, ϕ)) has **observable deviators** if the factorization

$$S(h) = \times_{j \in I} S_j(h)$$

holds for each information set h of each player. Intuitively, this means that if an information set is reached unexpectedly the active player is able to understand who deviated from the expected path (cf. Fudenberg and Levine 1993, and Battigalli and Guaitoli 1998). We can prove the following monotonicity result for the RSCE correspondence:

Theorem 12 *Fix two games with observable payoffs and observable deviators where no player moves more than once, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then $RSC E(\Gamma, f, \bar{\phi}) \subseteq RSC E(\Gamma, f, \phi)$.*

We assume observable deviators because in the proof we use a version of the previous result about monotonicity of the justifiability correspondence where $S_{-i}(h)$ is replaced by the (compact) space of product distributions on $S_{-i}(h)$, which makes sense if $S_{-i}(h)$ is a product set.

9 Discussion

We discuss the relevance of some assumptions, some extensions, and possible modifications of solution concepts.

9.1 Pure equilibrium and absence of chance moves

We extended the monotonicity result of BCMM and proved monotonicity of the (rationalizable) SCE correspondence with respect to ambiguity aversion in games where players move at most once on any path, and in games without chance moves if attention is restricted to symmetric, i.e. “pure,” equilibria. Here we discuss the relevance of symmetric (rationalizable) SCE in games without chance moves.

We justified our assumption that players do not randomize arguing that if randomizations were feasible they could be represented in the extensive form by means of chance moves. Thus, when we exclude chance moves we implicitly assume that players cannot randomize at all. We point out that, in our analysis of full unimprovability and rationalizable SCE, this assumption is made for simplicity and that it is crucial for Theorem 11 only.

A methodological reason to be interested in pure equilibria is that we may want to analyze explicitly the population game as a grand game with many players who are randomly matched.³⁷ Assuming that individual agents do not randomize (see the previous comment),

³⁷That is, if each population i has n agents, the grand game has $n \times |I|$ players.

we should consider the pure equilibria of such grand game. While this argument justifies the focus on pure equilibria, it also implies that we should allow for chance moves.

We do not focus on pure equilibria of games without chance moves because they are, in general, plausible, but rather because such focus improves our understanding of the interaction between nonmonotonicity of the equilibrium correspondence and dynamic inconsistency of preferences.

9.2 Heterogeneous populations

Our analysis can be easily extended to the case of populations with heterogeneous personal traits affecting preferences, or feedback. Let Θ_i be a parameter space of possible personal traits of agents in population i . Parameter $\theta_i \in \Theta_i$ may affect tastes and risk attitudes (v_{θ_i}), ambiguity attitudes (ϕ_{θ_i}), and feedback (f_{θ_i}). Each agent's personal traits are privately known to him. Assume for simplicity that Θ_i is finite and let $\tau_i \in \Delta(\Theta_i)$ denote the exogenous distribution of personal traits in population i . For any $j \in I$ and $\theta_j \in \Theta_j$ we let σ_{θ_j} denote the strategy distribution in the subpopulation of agents with trait θ_j . For any profile of distributions $(\sigma_{\theta_j})_{\theta_j \in \Theta_j} \in \Delta(S_j)^{\Theta_j}$ and any $s_j \in S_j$ the aggregate fraction of agents playing s_j is $\sum_{\theta_j \in \Theta_j} \tau_j(\theta_j) \sigma_{\theta_j}(s_j)$. With this, we can extend SCE and the refinements analyzed in this paper by taking into account that unimprovability and belief confirmation depend on personal traits (cf. Ch. 7 in Battigalli 2017, and Dekel et al. 2004): A profile of distributions $\left((\sigma_{\theta_i})_{\theta_i \in \Theta_i} \right)_{i \in I}$ is a selfconfirming equilibrium if there is a profile of beliefs $\left((\mu_{\theta_i, s_i})_{\theta_i \in \Theta_i, s_i \in \text{Supp} \sigma_{\theta_i}} \right)_{i \in I}$ such that, for all $i \in I$, $\theta_i \in \Theta_i$, and $s_i \in \text{Supp} \sigma_{\theta_i}$

1. s_i is unimprovable given μ_{θ_i, s_i} , v_{θ_i} , and ϕ_{θ_i} ,
2. $\mu_{\theta_i, s_i} \left(\left\{ \sigma'_{-i} \in \Sigma_{-i} : \hat{F}_{\theta_i}(s_i, \sigma'_{-i}) = \hat{F}_{\theta_i} \left(s_i, \times_{j \neq i} \left(\sum_{\theta_j \in \Theta_j} \tau_j(\theta_j) \sigma_{\theta_j} \right) \right) \right\} \right) = 1$.

The definition of rationalizable SCE can be extended in a similar way, taking into account what is assumed to be commonly known about the exogenous distributions of traits. Our results about nonsymmetric SCE can be seamlessly extended to such notion of (rationalizable) SCE. As for the results about symmetric (rationalizable) SCE in games without chance moves, they only apply to the case in which *all* agents of the same population use the same strategy despite differences in personal traits. Therefore, these results are less relevant.

9.3 Incomplete information

A related extension concerns the case where some features of game are not known or commonly known. It is straightforward to relax the assumption that the probabilities of

chance moves are (commonly) known: Chance can be analyzed as an indifferent player;³⁸ then, one just applies the definitions we already have for games without chance moves, letting the chance player distribution be the one obtained from the objective probabilities of chance moves. Of course, there cannot be symmetric equilibria of such game, because the strategy distribution of the chance player is strictly positive. More generally, we can parametrize all the features of the game that are not commonly known with a parameter profile $\theta = (\theta_j)_{j \in I \cup \{0\}} \in \Theta$, where θ_0 parameterizes chance probabilities and other non personal aspects of the game such as some aspects of the consequence function, and each θ_i ($i \in I$) parameterizes privately known personal features. Then we can give definitions of SCE at a given θ , and of rationalizable SCE given that it is common knowledge that $\theta \in \Theta$ (cf. Ch. 7 in Battigalli 2017, Battigalli and Guaitoli 1988, and Esponda 2013).

9.4 Rationalizable selfconfirming equilibrium

We put forward a relatively simple definition of rationalizable SCE that can be formally derived from the epistemic assumptions of (a) subjective rationality, (b) belief confirmation, and (c) initial common belief of (a) and (b). These epistemic assumptions can be formally stated by means of extensive-form type structures (see, e.g., Battigalli et al. 2017). The derivation of symmetric rationalizable SCE from assumptions (a)-(c) is quite straightforward for the symmetric case, which is relevant for repeated games played by impatient players. The derivation from (a)-(c) of nonsymmetric rationalizable SCE in population games is more complex. We can think of one that goes through approximations *via* large finite populations. We provide below a critical discussion of rationalizable SCE based on our knowledge of dynamic epistemic game theory and our intuitions about learning dynamics. We emphasize that this discussion is not specifically related to the issue of ambiguity aversion.

While plain SCE can be justified as the set of rest points of learning dynamics (e.g., Battigalli et al. 1992, Battigalli et al. 2016c and references therein), we are not aware of any learning foundation for rationalizable SCE. As we think about one, the very definition of rationalizable SCE appears questionable. Let us consider the simplest possible scenario: A fixed set I of completely *impatient* players plays Γ infinitely many times, with imperfect monitoring about past periods outcomes (terminal nodes) given by the profile of feedback functions $f = (f_i)_{i \in I}$; if players have nonneutral attitudes toward ambiguity, they are represented by the profile of second-order utility functions $\phi = (\phi_i)_{i \in I}$; (Γ, f, ϕ) is common knowledge. Suppose that players are rational and there is *common belief in rationality at the beginning of the first period*. Also suppose for simplicity that players' initial beliefs about first-period strategies assign strictly positive probability to opponents' first-period strategies consistent with rationality and initial common belief in rationality,

³⁸More precisely, chance can be analyzed as an indifferent player no feedback. Ex post perfect recall applies only to real players.

called “**initially rationalizable**” strategies (Battigalli et al. 2017).³⁹ Then players never observe zero-probability events and have well defined updated beliefs that always satisfy one-period rationality and common initial belief in rationality; hence, they play initially rationalizable strategies of (Γ, f, ϕ) in every period. If one-period strategies and updated beliefs converge, then the limit is a pure SCE in initially rationalizable strategies of (Γ, f, ϕ) . This is a weaker equilibrium concept than rationalizable SCE, because it requires common initial belief in rationality, not common belief in confirmation. This (somewhat informal) argument shows that the conjunction of SCE and initial rationalizability is necessary condition of steady states of learning dynamics that satisfy the foregoing epistemic assumptions. The arguments offered by Rubinstein and Wolinsky (1994) suggest, however, that this condition may not be sufficient for stability (cf. Examples 9 and 11). We conjecture that rationalizable SCE is a sufficient condition for stability, but it is not clear to us why it should also be necessary.

Be as it may, why should we assume players hold strictly positive beliefs over the set of initially rationalizable strategies? After all, the epistemic assumptions that justify initial rationalizability do not require any kind of strict positivity of beliefs. If we drop this assumption about the support of initial beliefs, we open the door to surprises as the learning dynamic unfolds and we have no reason to suppose that surprised players keep believing in the strategic sophistication of the opponents, or even their subjective rationality, unless we strengthen the assumptions about players’ belief in the strategic sophistication of opponents. Indeed, the very definition of initial rationalizability allows surprised players to assign positive updated probabilities to never-best-replies of the opponents. Thus, we suggest that a more interesting approach is to assume “common strong belief” in rationality: players always assign probability 1 to the “highest level of strategic sophistication” of the opponents consistent with what they observe. Within a one-period play, these epistemic assumptions yield a solution concept called “**strong rationalizability**” (Battigalli et al. 2017). Suppose that these epistemic assumptions hold and the learning dynamic converges; then the limit must be a pure SCE in strongly rationalizable strategies, which is—essentially—the concept first put forward by Battigalli (1987) and Battigalli and Guaitoli (1988). Thus, SCE in strongly rationalizable strategies is a necessary condition for stability under the foregoing epistemic assumptions. Again, arguments *à la* Rubinstein and Wolinsky (1994) suggest that it may not be sufficient. We can give a definition of “**strongly rationalizable SCE**” in which “initial belief” is replaced by “strong belief.” We conjecture that strongly rationalizable SCE provides sufficient conditions for stability, but it is not clear to us that these conditions are also necessary.

Of course, the assumption that players are impatient is just a simplification. If a player is somewhat patient and plays a repeated game, he may sacrifice short-run expected payoffs to experiment, or to teach opponents to play in future periods in ways he finds advantageous

³⁹Initially rationalizable strategies can be obtained with an iterated elimination algorithm similar to the one used to define symSCE, but without the confirmed-belief requirement.

to him. This is why the literature on learning and convergence to equilibrium in games often focuses on population games: if populations are large, the incentive to teach is shut down, and the incentive to experiment disappears in the long run if updated beliefs converge. The interaction of epistemic assumptions and learning dynamics in recurrent play within a population game scenario is hard to analyze, and learning foundations for (nonsymmetric) rationalizable SCE are therefore harder to provide. An alternative route is to put forward a definition of rationalizable SCE for repeated games, similar to the definition of “subjective equilibrium” of Kalai and Lehrer (1993a,b, 1995), but with the additional requirement of common belief in rationality. Arguments similar to those offered above suggest that an equilibrium concept based on common strong belief in rationality and—possibly—in belief confirmation could capture interesting necessary and/or sufficient conditions for the stability of beliefs (about repeated games strategies) under learning.

To sum up, it is not surprising that more than two decades after Rubinstein and Wolinsky (1994) put forward the rationalizable SCE concept a learning foundation is still lacking. We suggest that, as one seriously tries to provide such a foundation, different versions of the rationalizable SCE will concepts turn out to be relevant.

10 Appendix

10.1 Proof of Proposition 3

Every finite game Γ has a sequential equilibrium in behavioral strategies $\bar{\beta} = (\bar{\beta}_i)_{i \in I}$. For each $i \in I$, let $\bar{\sigma}_i$ denote the mixed strategy associated with $\bar{\beta}_i$ according to Kuhn’s (1953) transformation:

$$\forall s_i \in S_i, \bar{\sigma}_i(s_i) = \prod_{h \in H_i} \bar{\beta}_i(s_{i,h}|h).$$

Let $\mu_{\bar{s}_i} = \delta_{\bar{\sigma}_i}$ for every i . By construction, these beliefs are correct, hence confirmed. To show that $\bar{\sigma}$ is an SCE justified by these confirmed beliefs, we must prove that, for each i , each $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$ is $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable.

It is well known that every pure strategy in the support of a sequential equilibrium is a sequential best reply to the equilibrium beliefs. Therefore, for every $i \in I$, $h \in H_i$ and \bar{s}_i such that $\prod_{h' \in H_i} \bar{\beta}_i(\bar{s}_{i,h'}|h') > 0$, that is, for every $\bar{s}_i \in \text{Supp} \bar{\sigma}_i$,

$$\bar{s}_{i,h} \in \arg \max_{a_i \in A_i(h)} \sum_{x \in h} \mathbb{P}_{(\bar{s}_{i,h}, a_i), \bar{\beta}_{-i}}(x|h) \sum_{z \in Z} \mathbb{P}_{(\bar{s}_{i,h}, a_i), \bar{\beta}_{-i}}(z|x) u_i(z),$$

where $\mathbb{P}_{s_i, \bar{\beta}_{-i}}(\cdot|\cdot)$ denotes the probability of a node conditional on an information set, or an earlier node, determined by the behavioral strategy profile $(s_i, \bar{\beta}_{-i})$ and the probabilities of chance moves. Since $\bar{\sigma}_{-i}$ is by construction realization-equivalent to $\bar{\beta}_{-i}$, $\mathbb{P}_{s_i, \bar{\beta}_{-i}}(\cdot|\cdot) =$

$\mathbb{P}_{s_i, \bar{\sigma}_{-i}}(\cdot|\cdot)$; hence,

$$\forall s_i \in S_i(h), U_i(s_i, \bar{\sigma}_{-i}|h) = \sum_{x \in h} \mathbb{P}_{s_i, \bar{\beta}_{-i}}(x|h) \sum_{z \in Z} \mathbb{P}_{s_i, \bar{\beta}_{-i}}(z|x) u_i(z).$$

Since $\mu_{\bar{s}_i} = \delta_{\bar{\sigma}_{-i}}$,

$$\forall a_i \in A_i(h), V_i(a_i|h; \mu_{\bar{s}_i}, \phi_i) = \phi_i^{-1}(\phi_i(U_i((\bar{s}_i|h, a_i), \bar{\sigma}_{-i}|h))) = U_i((\bar{s}_i|h, a_i), \bar{\sigma}_{-i}|h).$$

Therefore \bar{s}_i is $(\mu_{\bar{s}_i}, \phi_i)$ -unimprovable. \blacksquare

10.2 Example 6

We prove that in Example 6, *there exists a concave and strictly increasing transformation $\phi_1 = \varphi_1 \circ \bar{\phi}_1$ such that for every belief μ_1 , Out is not prescribed by any (ϕ_1, μ_1) -unimprovable strategy of player 1.* As anticipated in Example 5, we look for a ϕ_1 such that

$$\begin{aligned} \phi_1(u) &= \bar{\phi}_1(u) = \sqrt{u} \quad \text{if } u \in [0, 9], \\ \bar{\phi}_1(36) &= 6 > k = \phi_1(36) > \phi_1(9) = 3, \end{aligned}$$

with $k = \phi_1(36)$ close to 3.

For every σ_2 in $\Delta(A_2(\{\text{In}, \text{G}\})) = \Delta(\{\text{L}, \text{R}\})$, let

$$e_M(\sigma_2) = 1 \cdot \sigma_2(\text{L}) + 36 \cdot \sigma_2(\text{R})$$

denote the objective expected payoff of M in the subgame. For every probability measure ν on $\Delta(\{\text{L}, \text{R}\})$, define $G_\nu : \Delta(\{\text{L}, \text{R}\}) \rightarrow [0, 1]$ as

$$G_\nu(\sigma_2) = \nu(\{\sigma'_2 \in \Delta(\{\text{L}, \text{R}\}) : e_M(\sigma'_2) \leq e_M(\sigma_2)\}).$$

Player 1 prefers M to T only if

$$V_1(\text{M} | \{\{\text{In}, \text{G}\}\}; s_1, \mu_1, \phi_1) \geq 9 = V_1(\text{T} | \{\{\text{In}, \text{G}\}\}; s_1, \mu_1, \phi_1).$$

Let $\nu = \mu_1(\cdot | \{\{\text{In}, \text{G}\}\})$. Then, for every $\sigma_2 \in \Delta(\{\text{L}, \text{R}\})$,⁴⁰

$$\phi_1^{-1}(G_\nu(\sigma_2) \cdot \phi_1(e_M(\sigma_2)) + (1 - G_\nu(\sigma_2)) \cdot \phi_1(36)) \geq V_1(\text{M} | \{\{\text{In}, \text{G}\}\}; s_1, \mu_1, \phi_1),$$

that is,

$$G_\nu(\sigma_2) \cdot \phi_1(e_M(\sigma_2)) + (1 - G_\nu(\sigma_2)) \cdot \phi_1(36) \geq \phi_1(V_1(\text{M} | \{\{\text{In}, \text{G}\}\}; s_1, \mu_1, \phi_1)).$$

⁴⁰ Given the belief ν and a model σ_2 , $G_\nu(\sigma_2)$ is the probability assigned by ν to the models that, paired with action M, yield an (objective) expected utility lower or equal to the one obtained under model σ_2 . Therefore, the value of action M under ν cannot exceed the value of M under the following belief: the models that, paired with action M, yield an (objective) expected utility equal to the one obtained under model σ_2 have probability $G_\nu(\sigma_2)$, whereas δ_R has probability $(1 - G_\nu(\sigma_2))$.

Therefore, a necessary condition for 1 to choose M over T is that, for every $\sigma_2 \in \Delta(\{L, R\})$,

$$G_\nu(\sigma_2) \cdot \phi_1(e_M(\sigma_2)) + (1 - G_\nu(\sigma_2)) \cdot k \geq 3.$$

Solving for $\hat{\sigma}_2$ such that $e_M(\hat{\sigma}_2) = 17/2$, we get

$$G_\nu(\hat{\sigma}_2) \leq \frac{k-3}{k-\sqrt{17/2}}. \quad (11)$$

The interpretation is that the probability assigned to the models under which action M has an (objective) value lower or equal than $17/2$ must be sufficiently small if player 1 is to choose M at (In, G).

Now, we return to the choice of player 1 at the root of the game. Note that

$$\mu_1 = \mu_1(\cdot | \{\emptyset\}) = \mu_1(\cdot | \{(In, G)\})$$

because the belief about the co-player cannot change upon observing one's own move, or a chance move. The DM is sophisticated, therefore, before choosing between In and Out, he predicts his behavior at (In, G). There are three cases:

1. Player 1 understands that his subjective belief μ_1 makes him play T at (In, G). Then, since In.T is preferred to Out for every belief, he will play In at $\{\emptyset\}$.
2. Player 1 understands that his subjective belief μ_1 makes him play M at (In, G). But then, it must be the case that (11) holds. Therefore, the probability assigned to the models that (given In.G) yield an objective expected utility of choosing action M larger than or equal to $\frac{17}{2}$ is at least $1 - \frac{k-3}{k-\sqrt{17/2}}$. In turn, this implies that the probability assigned to the models that, ex-ante, yields an objective expected utility of strategy In.M larger than or equal to $\frac{17}{4}$ is at least $1 - \frac{k-3}{k-\sqrt{17/2}}$. But then, the evaluation of strategy In.M under belief μ_1 satisfies the following conditions:⁴¹

$$\begin{aligned} \phi_1(V_1(In | \{\emptyset\}; In.M, \mu_1, \phi_1)) &\geq \phi_1(1/2) \cdot \frac{k-3}{k-\sqrt{17/2}} + \phi_1(17/4) \cdot \left(1 - \frac{k-3}{k-\sqrt{17/2}}\right) \\ &= \frac{\sqrt{17}}{2} - \frac{k-3}{k-\sqrt{17/2}} \cdot \left(\frac{\sqrt{17}}{2} + 1\right). \end{aligned}$$

For $k = \phi_1$ (36) sufficiently close to 3, $\phi_1(V_1(In | \{\emptyset\}; In.M, \mu_1, \phi_1))$ is higher than 2, thus player 1 chooses In over Out.

⁴¹The first inequality is due to the following fact. The probability assigned by μ_1 to the models that, ex-ante, given strategy In.M, yield an (objective) expected utility larger or equal than $\frac{17}{4}$ must be larger or equal than $1 - \frac{k-3}{k-\sqrt{17/2}}$. Therefore, the value of action In given strategy In.M under μ_1 cannot be lower than its value under the following belief: the set of models that, given strategy In.M, yield an (objective) expected utility equal to $\frac{17}{4}$ has subjective probability $1 - \frac{k-3}{k-\sqrt{17/2}}$, whereas δ_L has subjective probability $\frac{k-3}{k-\sqrt{17/2}}$.

3. Player 1 understands that his subjective belief μ_1 makes him play B at (In, G). A similar argument as for the previous case shows that for k sufficiently close to 3 also in this case player 1 chooses In over Out.

Summing up, there is a concave and strictly increasing transformation $\phi_1 = \varphi_1 \circ \bar{\phi}_1$ such that Out is not prescribed by any (μ_1, ϕ_1) -unimprovable strategy for all beliefs μ_1 .

10.3 Monotonicity of the SCE correspondence

Define the set of messages for i consistent with an information set $h \in H_i$ as follows:

$$M_i(h) = \{m : \exists (x, z) \in h \times Z, (x \prec z) \wedge (m = f_i(z))\}.$$

Lemma 3 *Ex post perfect recall implies that*

$$S_{0,-i}(h) = \bigcup_{m \in M_i(h)} F_{i,s_i}^{-1}(m)$$

for all $i \in I$, $h \in H_i$, and $s_i \in S_i(h)$.

In words, $S_{0,-i}(h)$ is the union of the sets of preimages of messages consistent with h , because these messages “record” that h has been reached.

Proof. Fix i , $h \in H_i$, and $s_i \in S_i(h)$ arbitrarily. First note that perfect recall implies

$$S_{0,I}(h) = S_i(h) \times S_{0,-i}(h).$$

We first prove that

$$S_{0,-i}(h) \subseteq \bigcup_{m \in M_i(h)} F_{i,s_i}^{-1}(m).$$

Fix any $s_{0,-i} \in S_{0,-i}(h)$; since $s_i \in S_i(h)$ and $S_{0,I}(h) = S_i(h) \times S_{0,-i}(h)$, then $(s_i, s_{0,-i}) \in S_{0,I}(h)$, that is, $x \prec \zeta(s_i, s_{0,-i})$ for some $x \in h$. Thus, by definition of $M_i(h)$, $f_i(\zeta(s_i, s_{0,-i})) \in M_i(h)$. Hence, $s_{0,-i} \in F_{i,s_i}^{-1}(m)$ for some $m \in M_i(h)$.

Next we prove by contraposition that the converse

$$S_{0,-i}(h) \supseteq \bigcup_{m \in M_i(h)} F_{i,s_i}^{-1}(m)$$

is implied by *ex post* perfect recall. Suppose that we can find some $m \in M_i(h)$ and $s'_{0,-i} \in F_{i,s_i}^{-1}(m) \setminus S_{0,-i}(h)$. We show that this implies a violation of *ex post* perfect recall. Since $m \in M_i(h)$, there is a pair $(x, z) \in h \times Z$ such that $x \prec z$ and $f_i(z) = m$. Fix any $s_{0,-i} \in \text{proj}_{S_{0,-i}} \zeta^{-1}(z)$, so that $(s_i, s_{0,-i}) \in S_i(h) \times S_{0,-i}(h) = S_{0,I}(h)$ for some $s_i \in S_i(h)$. Let $z = \zeta(s_i, s_{0,-i})$ and $z' = \zeta(s_i, s'_{0,-i})$. Then, by choice of $s_{0,-i}$ and $s'_{0,-i}$, $f_i(z) =$

$m = f_i(z')$, z is preceded by a node of h and z' is not preceded by any node of h . Hence, there are $z, z' \in Z$ such that z has a predecessor in h , z' has no predecessor in h , and yet $f_i(z) = m = f_i(z')$, thus violating ex post perfect recall. \blacksquare

The following result says that the value of equilibrium actions is unambiguous, hence independent of ambiguity attitudes:

Lemma 4 *Let $\bar{\sigma}$ be an SCE of the game with observable payoffs (Γ, f, ϕ) justified by the confirmed beliefs $(\mu_{s_i})_{i \in I, s_i \in \text{Supp} \bar{\sigma}_i}$. For every $i \in I$ and $s_i \in \text{Supp} \bar{\sigma}_i$ and $h \in H_i(\mu_{s_i}) \cap H_i(s_i)$, action $s_{i,h}$ is $\mu_i(\cdot|h)$ -unambiguous, and its value is the conditional objective expected payoff, that is,*

$$V_i(s_{i,h}|h; s_i, \mu_{s_i}, \phi_i) = U_i(s_i, \bar{\sigma}_{-i}|h).$$

Proof. By ex post perfect recall and Lemma 3,

$$S_{0,-i}(h) = \bigcup_{m \in M_i(h)} F_{i,s_i}^{-1}(m) \quad (12)$$

for each $h \in H_i$. Also, recall that

$$\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) = \left\{ \sigma_{-i} \in \Sigma_{-i} : \forall m, (\sigma_0 \times \sigma_{-i})(F_{i,s_i}^{-1}(m)) = (\sigma_0 \times \bar{\sigma}_{-i})(F_{i,s_i}^{-1}(m)) \right\}$$

is the partially identified set of co-players distributions of strategies observationally equivalent (obs.eq) for i to $\bar{\sigma}_{-i}$ given s_i .

Fix any $h \in H_i(\mu_{s_i}) \cap H_i(s_i)$. Then,

$$\begin{aligned} & \forall \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}), \\ & (\sigma_0 \times \sigma_{-i})(S_{0,-i}(h)) \stackrel{(12, \text{obs.eq})}{=} (\sigma_0 \times \bar{\sigma}_{-i})(S_{0,-i}(h)) \stackrel{(12, \text{conf})}{=} (\sigma_0 \times p_{\mu_{s_i}})(S_{0,-i}(h)) > 0, \end{aligned} \quad (13)$$

where the first equality follows from eq. (12) and the fact that σ_{-i} is observationally equivalent to $\bar{\sigma}_{-i}$, the second equality follows from eq. (12) and belief confirmation (conf), that is $\mu_{s_i}(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i})) = 1$, and the inequality follows from $h \in H_i(\mu_{s_i})$.

Fix any $m \in M_i(h)$; by observable payoffs (obs.p), there is $u^m \in \mathbb{R}$ such that $u_i(\zeta(s_0, s_i, s_{-i})) = u^m$ for all $(s_0, s_{-i}) \in S_0 \times S_{-i}$ with $f_i(\zeta(s_0, s_i, s_{-i})) = m$. Then, observable payoffs and eq.s (12)-(13) imply that

$$\forall \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}), U_i(s_i, \sigma_{-i}|h) = U_i(s_i, \bar{\sigma}_{-i}|h). \quad (14)$$

Indeed, we have

$$U_i(s_i, \sigma_{-i}|h) \stackrel{(\text{def})}{=} \sum_{(s_0, s_{-i}) \in S_{0,-i}(h)} \frac{\sigma_0(s_0) \cdot \sigma_{-i}(s_{-i})}{(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))} u_i(\zeta(s_0, s_i, s_{-i}))$$

$$\begin{aligned}
&\stackrel{(12,\text{obs.p})}{=} \sum_{m \in M_i(h)} \frac{(\sigma_0 \times \sigma_{-i})(F_{i,s_i}^{-1}(m))v^m}{(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))} \stackrel{(\text{obs.eq})}{=} \sum_{m \in M_i(h)} \frac{(\sigma_0 \times \bar{\sigma}_{-i})(F_{i,s_i}^{-1}(m))v^m}{(\sigma_0 \times \bar{\sigma}_{-i})(S_{0,-i}(h))} \\
&\stackrel{(12,\text{obs.p})}{=} \sum_{(s_0, s_{-i}) \in S_{0,-i}(h)} \frac{\sigma_0(s_0) \cdot \bar{\sigma}_{-i}(s_{-i})}{(\sigma_0 \times \bar{\sigma}_{-i})(S_{0,-i}(h))} u_i(\zeta(s_0, s_i, s_{-i})) \stackrel{(\text{def})}{=} U_i(s_i, \bar{\sigma}_{-i}|h).
\end{aligned}$$

Confirmed beliefs $(\mu_{s_i}(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i})) = 1)$, the updating formula (5), and eq.s (13)-(14) yield

$$\begin{aligned}
V_i(s_{i,h}|h; s_i, \mu_{s_i}, \phi_i) &= \phi_i^{-1} \left(\int_{\Sigma_{-i}(h)} \phi_i(U_i(s_i, \sigma_{-i}|h)) \mu_{s_i}(d\sigma_{-i}|h) \right) \\
&\stackrel{(5)}{=} \phi_i^{-1} \left(\int \phi_i(U_i(s_i, \sigma_{-i}|h)) \frac{(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))}{(\sigma_0 \times p_{\mu_{s_i}})(S_{0,-i}(h))} \mu_{s_i}(d\sigma_{-i}) \right) \\
&\stackrel{(14,\text{conf.})}{=} \phi_i^{-1} \left(\phi_i(U_i(s_i, \sigma_{-i}|h)) \int \frac{(\sigma_0 \times \sigma_{-i})(S_{0,-i}(h))}{(\sigma_0 \times p_{\mu_{s_i}})(S_{0,-i}(h))} \mu_{s_i}(d\sigma_{-i}) \right) \\
&\stackrel{(13)}{=} \phi_i^{-1}(\phi_i(U_i(s_i, \bar{\sigma}_{-i}|h))) = U_i(s_i, \bar{\sigma}_{-i}|h).
\end{aligned}$$

■

An increase in ambiguity aversion (weakly) decreases the values of actions:

Lemma 5 *If (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then,*

$$V_i(a_i|h; s_i, \mu_i, \bar{\phi}_i) \geq V_i(a_i|h; s_i, \mu_i, \phi_i)$$

for all $i \in I$, $\mu_i \in \Delta(\Sigma_{-i})$, $s_i \in S_i$, $h \in H_i(\mu_i)$, and $a_i \in A_i(h)$.

Proof. Fix i and s_i arbitrarily. For every $h \in H_i$ and $a_i \in A_i(h)$, define the following auxiliary function:

$$\begin{aligned}
U_{s_i, h, a_i} : \Sigma_{-i}(h) &\rightarrow \mathbb{R}, \\
\sigma_{-i} &\mapsto U_i((s_i|_h, a_i), \sigma_{-i}|h).
\end{aligned}$$

With this, for every $\mu_i \in \Delta(\Sigma_{-i})$ and $h \in H_i(\mu_i)$, the conditional belief $\mu_i(\cdot|h)$ is well defined and

$$\begin{aligned}
V_i(a_i|h; s_i, \mu_i, \bar{\phi}_i) &= \bar{\phi}_i^{-1}(\mathbb{E}_{\mu_i(\cdot|h)}[\bar{\phi}_i \circ U_{s_i, h, a_i}]), \\
V_i(a_i|h; s_i, \mu_i, \phi_i) &= \phi_i^{-1}(\mathbb{E}_{\mu_i(\cdot|h)}[\phi_i \circ U_{s_i, h, a_i}]).
\end{aligned} \tag{15}$$

Since $\phi_i = \varphi \circ \bar{\phi}_i$ for some concave and strictly increasing $\varphi : \bar{\phi}_i(\mathbb{V}_i) \rightarrow \mathbb{R}$, Jensen's inequality implies

$$\varphi \left(\mathbb{E}_{\mu_i(\cdot|h)} [\bar{\phi}_i \circ U_{s_i, h, a_i}] \right) \geq \mathbb{E}_{\mu_i(\cdot|h)} [\varphi \circ \bar{\phi}_i \circ U_{s_i, h, a_i}] = \mathbb{E}_{\mu_i(\cdot|h)} [\phi_i \circ U_{s_i, h, a_i}].$$

By monotonicity of φ and $\bar{\phi}_i$, and recalling that $\phi_i^{-1} = \bar{\phi}_i^{-1} \circ \varphi^{-1}$, we obtain

$$\mathbb{E}_{\mu_i(\cdot|h)} [\bar{\phi}_i \circ U_{s_i, h, a_i}] \geq \varphi^{-1} \left(\mathbb{E}_{\mu_i(\cdot|h)} [\phi_i \circ U_{s_i, h, a_i}] \right),$$

$$\bar{\phi}_i^{-1} \left(\mathbb{E}_{\mu_i(\cdot|h)} [\bar{\phi}_i \circ U_{s_i, h, a_i}] \right) \geq \left(\bar{\phi}_i^{-1} \circ \varphi^{-1} \right) \left(\mathbb{E}_{\mu_i(\cdot|h)} [\phi_i \circ U_{s_i, h, a_i}] \right) = \phi_i^{-1} \left(\mathbb{E}_{\mu_i(\cdot|h)} [\phi_i \circ U_{s_i, h, a_i}] \right).$$

By eq. (15), $V_i(a_i|h; s_i, \mu_i, \bar{\phi}_i) \geq V_i(a_i|h; s_i, \mu_i, \phi_i)$. ■

Lemma 6 *Fix two games with observable payoffs, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, such that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Fix an SCE $\bar{\sigma}$ of $(\Gamma, f, \bar{\phi})$ and justifying beliefs $(\mu_{s_i})_{i \in I, s_i \in \text{Supp} \bar{\sigma}_i}$. Suppose that for each $i \in I$, every $s_i \in \text{Supp} \bar{\sigma}_i$ is (μ_{s_i}, ϕ_i) -sequentially optimal. Then, there exist maps $\bar{s}_i : \text{Supp} \bar{\sigma}_i \rightarrow S_i$, $i \in I$, such that*

- (i) *for every $i \in I$ and $h \in H_i(s_i)$, $\bar{s}_i(s_i)_h = s_{i,h}$, and*
- (ii) *$\sigma = (\bar{\sigma}_i \circ \bar{s}_i^{-1})_{i \in I}$ is an SCE of (Γ, f, ϕ) where for every $i \in I$ and $s_i \in \text{Supp} \bar{\sigma}_i$, μ_{s_i} justifies $\bar{s}_i(s_i)$.*

(So, σ and $\bar{\sigma}$ are realization equivalent and are justified by the same beliefs.)

Proof For every $i \in I$ and $s_i \in \text{Supp} \bar{\sigma}_i$, we construct a (μ_{s_i}, ϕ_i) -unimprovable strategy $\bar{s}_i(s_i)$. Map $\bar{s}_i : \text{Supp} \bar{\sigma}_i \rightarrow S_i$ is such that (1) $\bar{s}_i(s_i)_h = s_{i,h}$ for all $h \in H_i(s_i)$, and (2) $\bar{s}_i(s_i)$ is derived by folding back on $H_i(\mu_{s_i}) \setminus H_i(s_i)$ given μ_{s_i} and ϕ_i .⁴² Therefore, by construction, for every $h \in H_i(\mu_{s_i}) \setminus H_i(s_i)$,

$$\bar{s}_i(s_i)_h \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; \bar{s}_i(s_i), \mu_{s_i}, \phi_i).$$

Now, let $h \in H_i(\mu_{s_i}) \cap H_i(s_i)$. For every $a_i \in A_i(h)$, we have

$$\begin{aligned} V_i(\bar{s}_i(s_i)_h|h; \bar{s}_i(s_i), \mu_{s_i}, \phi_i) &= V_i(s_{i,h}|h; s_i, \mu_{s_i}, \phi_i) \\ &\stackrel{\text{(L4)}}{=} U_i(s_i, \bar{\sigma}_{-i}|h) \\ &\stackrel{\text{(L4)}}{=} V_i(s_{i,h}|h; s_i, \mu_{s_i}, \bar{\phi}_i) \\ &\stackrel{\text{(s.opt)}}{\geq} V_i(a_i|h; \bar{s}_i(s_i), \mu_{s_i}, \bar{\phi}_i) \\ &\stackrel{\text{(L5)}}{\geq} V_i(a_i|h; \bar{s}_i(s_i), \mu_{s_i}, \phi_i), \end{aligned}$$

⁴²Note that if $h \notin H_i(s_i)$, and $h' \succ h$, $h' \notin H_i(s_i)$, therefore the folding back construction is well defined. See Section 4 for the description of folding back optimality.

where the first equality follows from construction of $\bar{s}_i(s_i)$, the second and the third equalities from Lemma 4, the first inequality from sequential optimality (s.opt), and the second inequality from Lemma 5. This shows that $\bar{s}_i(s_i)$ is (μ_{s_i}, ϕ_i) -unimprovable.

To conclude, note that the profile $\left(\bar{\sigma}_i \circ \bar{s}_i^{-1}, (\mu_{\hat{s}_i})_{\hat{s}_i \in \text{Supp}(\bar{\sigma}_i \circ \bar{s}_i^{-1})}\right)_{i \in I}$, where $\mu_{\hat{s}_i}(\cdot) = \mu_{s_i}(\cdot)$ for some s_i with $\hat{s}_i = \bar{s}_i(s_i)$, also satisfies the confirmed beliefs condition (because $\left(\bar{s}_i(s_i), (\bar{\sigma}_j \circ \bar{s}_j^{-1})_{j \neq i}\right)$ and $(s_i, \bar{\sigma}_{-i})$, given σ_0 , induce the same distribution over terminal nodes). Therefore $(\bar{\sigma}_i \circ \bar{s}_i^{-1})_{i \in I}$ is an SCE of (Γ, f, ϕ) . \blacksquare

10.4 Symmetric equilibria

In this subsection and the following ones, we consider games *without chance moves*. Therefore, the outcome function and the strategic-form feedback and payoff functions of player i are, respectively,

$$\zeta : S \rightarrow Z,$$

$$F_i = f_i \circ \zeta : S \rightarrow M,$$

and

$$U_i = u_i \circ \zeta : S \rightarrow \mathbb{R}.$$

The proof of Theorem 5 requires some preliminary results. For every $i \in I$, consider the augmented collection

$$\hat{H}_i = H_i \cup \{f_i^{-1}(m) : m \in f_i(Z)\}$$

that includes also i 's terminal information sets. By ex post perfect recall of \hat{H}_i , we can derive from the game tree a transitive and antisymmetric precedence relation \preceq on \hat{H}_i that makes it a directed forest (collection of directed trees). Furthermore, we can extend to \hat{H}_i the definition of the sets $S(h)$, $S_i(h)$ and $S_{-i}(h)$ so that $S(h) = S_i(h) \times S_{-i}(h)$; in particular, for each $m \in f_i(Z)$,

$$\begin{aligned} S(f_i^{-1}(m)) &= F_i^{-1}(m) = \text{proj}_{S_i} F_i^{-1}(m) \times \text{proj}_{S_{-i}} F_i^{-1}(m) \\ &= S_i(f_i^{-1}(m)) \times S_{-i}(f_i^{-1}(m)). \end{aligned}$$

Note that, for each μ_i , we can define the collection of information sets in \hat{H}_i that are possible under μ_i :

$$\hat{H}_i(\mu_i) = \left\{ h \in \hat{H}_i : p_{\mu_i}(S_{-i}(h)) > 0 \right\}.$$

Note that, $H_i(\mu_i) \subset \hat{H}_i(\mu_i)$. Also, for each $h \in H_i$ and $a_i \in A_i(h)$, we define the collection of information sets in \hat{H}_i that “immediately” follow h and action a_i :

$$\hat{H}_i(h, a_i) = \left\{ h' \in \hat{H}_i : (h \prec h') \wedge (\alpha_i(h, h') = a_i) \wedge (\nexists h'' \in H_i, h \prec h'' \prec h') \right\}.$$

Finally, it is convenient to extend the formula for the values of actions (given strategy and belief) from $H_i(\mu_i)$ to $\hat{H}_i(\mu_i)$. Thus, we stipulate that if $h = f_i^{-1}(m)$ then $A_i(h) = \{a_i^m\}$ is the singleton that contains only the pseudo-action “ i observes m .” By observability of payoff, if $h = f_i^{-1}(m)$, then $V_i(a_i^m|h; s_i, \mu_i, \phi_i) = u_i(z)$ for every $z \in h$, independently of (s_i, μ_i, ϕ_i) .

Remark 6 *By ex post perfect recall, for every $i \in I$, $h \in H_i$ and $a_i \in A_i(h)$, $\{S_{-i}(h')\}_{h' \in \hat{H}_i(h, a_i)}$ is a partition of $S_{-i}(h)$.*

Proof We first prove that distinct elements of the collection $\{S_{-i}(h')\}_{h' \in \hat{H}_i(h, a_i)}$ are disjoint. By ex post perfect recall $S(h') = S_i(h') \times S_{-i}(h')$ for each $h' \in \hat{H}_i$. Let

$$S_i(h, a_i) = \{s_i \in S_i(h) : s_{i,h} = a_i\}$$

denote the strategies of i allowing for h and choosing a_i at h . Then, by definition of $\hat{H}_i(h, a_i)$, $S_i(h') = S_i(h, a_i)$ for each $h' \in \hat{H}_i(h, a_i)$, because i does not move again after choosing a_i and before any such h' . Therefore,

$$\forall h' \in \hat{H}_i(h, a_i), S(h') = S_i(h, a_i) \times S_{-i}(h').$$

Take any $h', h'' \in \hat{H}_i(h, a_i)$ with $h' \neq h''$ (if such distinct information sets exist). By definition of $\hat{H}_i(h, a_i)$, $h' \not\prec h''$ and $h'' \not\prec h'$. By perfect recall, this implies that

$$\{z' \in Z : \exists x' \in h', x' \preceq z'\} \cap \{z'' \in Z : \exists x'' \in h'', x'' \preceq z''\} = \emptyset,$$

which in turn implies that $S(h') \cap S(h'') = \emptyset$. Since $S_i(h') = S_i(h, a_i) = S_i(h'')$, then $S_{-i}(h') \cap S_{-i}(h'') = \emptyset$.

Note that

$$\bigcup_{h' \in \hat{H}_i(h, a_i)} S_{-i}(h') \subseteq S_{-i}(h)$$

because $h \prec h'$, hence $S_{-i}(h') \subseteq S_{-i}(h)$, for each $h' \in \hat{H}_i(h, a_i)$.

To prove the converse, pick any $(s_i, s_{-i}) \in S_i(h, a_i) \times S_{-i}(h)$ and let h' be the first information set in \hat{H}_i after h (possibly a terminal information set) reached by the path with terminal node $\zeta(s_i, s_{-i})$. Then $h' \in \hat{H}_i(h, a_i)$ and $s_{-i} \in S_{-i}(h')$. Therefore,

$$S_{-i}(h) \subseteq \bigcup_{h' \in \hat{H}_i(h, a_i)} S_{-i}(h').$$

We conclude that $\{S_{-i}(h')\}_{h' \in \hat{H}_i(h, a_i)}$ is a partition. ■

We say that a belief $\mu_i \in \Delta(\Sigma_{-i})$ is **supported by Dirac models** if

$$\mu_i(\{\delta_{s_{-i}} : s_{-i} \in S_{-i}\}) = 1.$$

Then

$$\forall s_{-i} \in S_{-i}, \mu_i(\delta_{s_{-i}}) = p_{\mu_i}(s_{-i})$$

and

$$\forall h \in \hat{H}_i(\mu_i), \forall s_{-i} \in S_{-i}(h), \mu_i(\delta_{s_{-i}|h}) = p_{\mu_i}(s_{-i}|h).$$

If μ_i is supported by Dirac models,

$$V_i(a_i|h; s_i, \mu_i, \phi_i) = \phi_i^{-1} \left(\sum_{s_{-i} \in S_{-i}(h)} p_{\mu_i}(s_{-i}|h) \phi_i(U_i((s_i|h, a_i), s_{-i})) \right).$$

Since ϕ_i is strictly increasing, maximizing $V_i(a_i|h; s_i, \mu_i, \phi_i)$ is the same as maximizing

$$\phi_i(V_i(a_i|h; s_i, \mu_i, \phi_i)) = \sum_{s_{-i} \in S_{-i}(h)} p_{\mu_i}(s_{-i}|h) \phi_i(U_i((s_i|h, a_i), s_{-i}))$$

and we can focus on the latter expected value. Under a belief supported by Dirac models, $\phi_i(V_i(s_i, h|h; s_i, \mu_i, \phi_i))$ is a weighted sum of the expected values given by s_i at the next information sets (including the terminal ones), where the weights are the predictive probabilities of reaching them.

Lemma 7 Fix any $i \in I$ and a belief supported by Dirac models $\hat{\mu}_i$. Then, for every $s_i \in S_i$ and $h \in H_i(\hat{\mu}_i)$,

$$\phi_i(V_i(s_i, h|h; s_i, \hat{\mu}_i, \phi_i)) = \sum_{h' \in \hat{H}_i(h, s_i, h)} p_{\hat{\mu}_i}(S_{-i}(h')|h) \phi_i(V_i(s_i|h', h'; s_i, \hat{\mu}_i, \phi_i)).$$

Proof. Fix $s_i \in S_i$ and $h \in H_i(\hat{\mu}_i)$ arbitrarily. By Remark 6 and taking into account that $\hat{\mu}_i$ is supported by Dirac models, we can write

$$\begin{aligned} \phi_i(V_i(s_i, h|h; s_i, \hat{\mu}_i, \phi_i)) &= \sum_{s_{-i} \in S_{-i}(h)} p_{\hat{\mu}_i}(s_{-i}|h) \phi_i(U_i(s_i|h, s_{-i})) \\ &= \sum_{h' \in \hat{H}_i(h, s_i, h)} \sum_{s_{-i} \in S_{-i}(h')} p_{\hat{\mu}_i}(s_{-i}|h) \phi_i(U_i(s_i|h, s_{-i})) \\ &= \sum_{h' \in \hat{H}_i(h, s_i, h): p_{\hat{\mu}_i}(S_{-i}(h')|h) > 0} p_{\hat{\mu}_i}(S_{-i}(h')|h) \sum_{s_{-i} \in S_{-i}(h')} \frac{p_{\hat{\mu}_i}(s_{-i}|h)}{p_{\hat{\mu}_i}(S_{-i}(h')|h)} \phi_i(U_i(s_i|h, s_{-i})). \end{aligned}$$

By the chain rule, for every $h' \in \hat{H}_i(h, s_i)$ with $p_{\hat{\mu}_i}(S_{-i}(h')|h) > 0$ and for every $s_{-i} \in S_{-i}(h')$,

$$\frac{p_{\hat{\mu}_i}(s_{-i}|h)}{p_{\hat{\mu}_i}(S_{-i}(h')|h)} = p_{\hat{\mu}_i}(s_{-i}|h').$$

Furthermore, for every $h' \in \hat{H}_i(h, s_{i,h})$ and $s_{-i} \in S_{-i}(h')$, $U_i(s_{i|h}, s_{-i}) = U_i(s_{i|h'}, s_{-i})$, because both $s_{i|h}$ and $s_{i|h'}$ make h' reachable and make the same choices before h' (by perfect recall), at h' , and after h' (by definition of replacement); thus,

$$\begin{aligned} & \sum_{s_{-i} \in S_{-i}(h')} \frac{p_{\hat{\mu}_i}(s_{-i}|h)}{p_{\hat{\mu}_i}(S_{-i}(h')|h)} \phi_i(U_i(s_{i|h}, s_{-i})) \\ = & \sum_{s_{-i} \in S_{-i}(h')} p_{\hat{\mu}_i}(s_{-i}|h') \phi_i(U_i(s_{i|h'}, s_{-i})) = \phi_i(V_i(s_{i|h'}|h'; s_i, \hat{\mu}_i, \phi_i)) \end{aligned}$$

and we obtain the desired result. \blacksquare

Iterating the previous formula and using the chain rule, one can prove the following:

Corollary 13 *Fix $i \in I$ and a belief $\hat{\mu}_i$ supported by Dirac models. Then, for every $s_i \in S_i$ and $h \in H_i(\hat{\mu}_i)$,*

$$\phi_i(V_i(s_{i,h}|h; s_i, \hat{\mu}_i, \phi_i)) = \sum_{s_{-i} \in S_{-i}} p_{\hat{\mu}_i}(s_{-i}|h) \phi_i(U_i(s_{i|h}, s_{-i})).$$

Thus, maximizing the conditional value of a strategy under a belief $\hat{\mu}_i$ supported by Dirac models is the same as maximizing the subjective expectation of the transformed utility $\hat{u}_i = \phi_i \circ U_i$. Applying Proposition 2 to $\hat{\mu}_i$ and \hat{u}_i we obtain the following:

Lemma 8 *For every $i \in I$ and $\hat{\mu}_i \in \Delta(\Sigma_{-i})$, if $\hat{\mu}_i$ is a belief supported by Dirac models, then the set of $(\hat{\mu}_i, \phi_i)$ -unimprovable and $(\hat{\mu}_i, \phi_i)$ -sequentially optimal strategies of i coincide.*

Note that, for every belief $\mu_i \in \Delta(\Sigma_{-i})$ there is a unique belief $\hat{\mu}_i$ supported by Dirac models with the same predictives as μ_i :

$$\forall s_{-i} \in S_{-i}, \hat{\mu}_i(\delta_{s_{-i}}) = p_{\mu_i}(s_{-i}).$$

This implies that

$$\forall h \in H_i(\mu_i), \forall s_{-i} \in S_{-i}(h), \hat{\mu}_i(\delta_{s_{-i}}|h) = p_{\mu_i}(s_{-i}|h).$$

We show that, for an ambiguity averse agent, $\hat{\mu}_i$ is the “most pessimistic” belief with the same predictive as μ_i in the following sense: for every $h \in H_i(\mu_i) = H_i(\hat{\mu}_i)$, the value at $h \in H_i(\mu_i)$ of a (μ_i, ϕ_i) -unimprovable strategy is at least as high as the value at h of a $(\hat{\mu}_i, \phi_i)$ -unimprovable strategy. Formally:

Lemma 9 *Fix $i \in I$, $\mu_i \in \Delta(\Sigma_{-i})$, and a concave second-order utility function ϕ_i . Let $\hat{\mu}_i \in \Delta(\Sigma_{-i})$ denote the belief supported by Dirac models with the same predictive as μ_i , and let s_i^* be (μ_i, ϕ_i) -unimprovable; then,*

$$\forall s_i \in S_i, \forall h \in \hat{H}_i(\mu_i), V_i(s_{i,h}^*|h; s_i^*, \mu_i, \phi_i) \geq V_i(s_{i,h}|h; s_i, \hat{\mu}_i, \phi_i).$$

Proof. First note that $\hat{H}_i(\mu_i)$ —with the obvious precedence relation—is a directed forest and its leaves are also leaves in \hat{H}_i as well, that is, if h is a leaf in $\hat{H}_i(\mu_i)$ then h is a set of terminal nodes ($h \in \{f_i^{-1}(m) : m \in f_i(Z)\}$). Furthermore, $\hat{H}_i(\mu_i) = \hat{H}_i(\hat{\mu}_i)$ because $p_{\mu_i}(S_{-i}(h)) = p_{\hat{\mu}_i}(S_{-i}(h))$ for every $h \in \hat{H}_i$. We prove the result by induction on the depth of h within $\hat{H}_i(\mu_i)$.

Basis step. Let $h = f_i^{-1}(m)$ be a terminal information set in $\hat{H}_i(\mu_i)$. Then the weak inequality holds trivially as an equality: in particular, by observable payoffs

$$\forall z \in h, V_i(s_{i,h}^*|h; s_i^*, \mu_i, \phi_i) = u_i(z) = V_i(s_{i,h}|h; s_i, \mu_i, \phi_i).$$

Inductive step. Let h have at least one follower in $\hat{H}_i(\mu_i)$ and suppose, by way of induction, that for every $h' \in \hat{H}_i(\mu_i)$ with $h \prec h'$,

$$V_i(s_{i,h'}^*|h'; s_i^*, \mu_i, \phi_i) \geq V_i(s_{i,h'}|h'; s_i, \hat{\mu}_i, \phi_i). \quad (\text{I.H.})$$

Consider the replacement plan $s_i^h = (s_{i|h}^*, s_{i,h})$, that is,

$$\forall h' \in H_i, \left(s_i^h\right)_{h'} = \left(s_{i|h}^*, s_{i,h}\right)_{h'} = \begin{cases} s_{i,h} & \text{if } h' = h, \\ \alpha_i(h', h) & \text{if } h' \prec h, \\ s_{i,h'}^* & \text{if } h' \not\prec h. \end{cases}$$

Since s_i^* is (μ_i, ϕ_i) -unimprovable

$$V_i(s_{i,h}^*|h; s_i^*, \mu_i, \phi_i) \geq V_i(s_{i,h}|h; s_i^*, \mu_i, \phi_i).$$

Next we prove that

$$V_i(s_{i,h}|h; s_i^*, \mu_i, \phi_i) \geq V_i(s_{i,h}|h; s_i, \hat{\mu}_i, \phi_i),$$

which implies the thesis.

To ease notation, write

$$\hat{H}_i(h, a_i, \mu_i) = \hat{H}_i(h, a_i) \cap \hat{H}_i(\mu_i)$$

for the collection of information sets that immediately follow h after a_i and can be reached with positive probability under μ_i ; $\hat{H}_i(h, a_i, \sigma_{-i})$ is similarly defined (thus, $\hat{H}_i(h, a_i, \sigma_{-i}) = \hat{H}_i(h, a_i, \delta_{\sigma_{-i}})$). By Remark 6, the definitions of V_i and s_i^h , Jensen's inequality, the chain rule for $\sigma_{-i}(\cdot|\cdot)$, Bayes rule for $\mu_i(\cdot|\cdot)$, Lemma 7, and the Inductive Hypothesis,

$$\phi_i(V_i(s_{i,h}|h; s_i^*, \mu_i, \phi_i)) \stackrel{(\text{R6,def})}{=} \int_{\Sigma_{-i}(h)} \phi_i \left(\sum_{h' \in \hat{H}_i(h, s_{i,h})} \sum_{s_{-i} \in S_{-i}(h')} \sigma_{-i}(s_{-i}|h) U_i(s_i^h, s_{-i}) \right) \mu_i(d\sigma_{-i}|h)$$

$$\begin{aligned}
&= \int_{\Sigma_{-i}(h)} \phi_i \left(\sum_{h' \in \hat{H}_i(h, s_{i,h}, \sigma_{-i})} \sigma_{-i}(S_{-i}(h') | h) \sum_{s_{-i} \in S_{-i}(h')} \frac{\sigma_{-i}(s_{-i}|h)}{\sigma_{-i}(S_{-i}(h') | h)} U_i(s_i^h, s_{-i}) \right) \mu_i(d\sigma_{-i}|h) \\
&\stackrel{(\text{Jen, ch})}{\geq} \int_{\Sigma_{-i}(h)} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \sigma_{-i})} \sigma_{-i}(S_{-i}(h') | h) \phi_i \left(\sum_{s_{-i} \in S_{-i}(h')} \sigma_{-i}(s_{-i}|h') U_i(s_i^h, s_{-i}) \right) \mu_i(d\sigma_{-i}|h) \\
&\stackrel{(\text{def})}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} \int_{\Sigma_{-i}(h')} \sigma_{-i}(S_{-i}(h') | h) \phi_i \left(U_i(s_{i|h'}^*, \sigma_{-i}|h') \right) \mu_i(d\sigma_{-i}|h) \\
&= \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}(S_{-i}(h') | h) \int_{\Sigma_{-i}(h')} \frac{\sigma_{-i}(S_{-i}(h') | h)}{p_{\mu_i}(S_{-i}(h') | h)} \phi_i \left(U_i(s_{i|h'}^*, \sigma_{-i}|h') \right) \mu_i(d\sigma_{-i}|h) \\
&\stackrel{(\text{Bayes})}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}(S_{-i}(h') | h) \int_{\Sigma_{-i}(h')} \phi_i \left(U_i(s_{i|h'}^*, \sigma_{-i}|h') \right) \mu_i(d\sigma_{-i}|h') \\
&\stackrel{(\text{def})}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}(S_{-i}(h') | h) \phi_i(V_i(s_{i,h'}^* | h'; s_i^*, \mu_i, \phi_i)) \\
&\stackrel{(\text{I.H.})}{\geq} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}(S_{-i}(h') | h) \phi_i(V_i(s_{i,h'} | h'; s_i, \hat{\mu}_i, \phi_i)) \stackrel{(\text{L7})}{=} \phi_i(V_i(s_{i,h} | h; s_i, \hat{\mu}_i, \phi_i)).
\end{aligned}$$

Since ϕ_i is strictly increasing, $V_i(s_{i,h} | h; s_i^*, \mu_i, \phi_i) \geq V_i(s_{i,h} | h; s_i, \hat{\mu}_i, \phi_i)$. \blacksquare

Proof of Theorem 5. Let $z^* \in \zeta(\text{symSCE}(\Gamma, f, \bar{\phi}))$. Then, there is a symmetric SCE $(s_i^*)_{i \in I}$ of $(\Gamma, f, \bar{\phi})$ such that $z^* = \zeta((s_i^*)_{i \in I})$. Let $(s_i^*)_{i \in I}$ be justified by the profile of confirmed beliefs $(\mu_i)_{i \in I}$. For each $i \in I$, let $m_i^* = f_i(z^*)$, and let $\hat{\mu}_i \in \Delta(\Sigma_{-i})$ the belief supported by Dirac models with the same predictive as μ_i , so that $H_i(\mu_i) = H_i(\hat{\mu}_i)$. Fix any $(\hat{\mu}_i, \bar{\phi}_i)$ -unimprovable strategy \hat{s}_i^* and let \hat{s}_i be the strategy that coincides with equilibrium strategy s_i^* on each path observationally equivalent to the equilibrium path, and with \hat{s}_i^* on the other paths, that is,

$$\forall h \in H_i, \hat{s}_{i,h} = \begin{cases} s_{i,h}^* & \text{if } h \prec f_i^{-1}(m_i^*), \\ \hat{s}_{i,h}^* & \text{if } h \not\prec f_i^{-1}(m_i^*). \end{cases}$$

By construction, $\zeta((\hat{s}_i)_{i \in I}) = \zeta((s_i^*)_{i \in I}) = z^*$. For each $i \in I$, since μ_i is a confirmed belief and $\hat{\mu}_i$ has the same predictive as μ_i ,

$$p_{\hat{\mu}_i}(S_{-i}(f^{-1}(m_i^*))) = p_{\mu_i}(S_{-i}(f^{-1}(m_i^*))) = 1.$$

Therefore, each belief in profile $(\hat{\mu}_i)_{i \in I}$ is confirmed under both $(s_i^*)_{i \in I}$ and $(\hat{s}_i)_{i \in I}$. We are going to prove that each \hat{s}_i is $(\hat{\mu}_i, \bar{\phi}_i)$ -sequentially optimal, which implies that $(\hat{s}_i)_{i \in I}$ is an SCE of $(\Gamma, f, \bar{\phi})$ and that Lemma 6 applies. Therefore, there is an equivalent symmetric SCE $(s_i)_{i \in I}$ of the more ambiguity averse game (Γ, f, ϕ) such that $\zeta((s_i)_{i \in I}) = z^*$, as desired.

Sequential optimality of \hat{s}_i under $(\bar{\phi}_i, \hat{\mu}_i)$. We show that \hat{s}_i is $(\bar{\phi}_i, \hat{\mu}_i)$ -unimprovable. Since $\hat{\mu}_i$ is supported by Dirac models, Lemma 8 implies that \hat{s}_i is also $(\hat{\mu}_i, \bar{\phi}_i)$ -sequentially optimal.

By belief confirmation and perfect recall, the collection of information sets over which s_i^* and \hat{s}_i coincide is given by those that are possible (indeed, certain) under $(\hat{s}_i, \hat{\mu}_i)$, that is,

$$H_i(\hat{\mu}_i) \cap H_i(\hat{s}_i) = \{h \in H_i : h \prec f_i^{-1}(m^*)\},$$

because i expects message m^* with certainty given that he plans to play \hat{s}_i .

By (ex post) perfect recall, for all $h, h' \in H_i(\hat{\mu}_i)$ such that $h \not\prec f_i^{-1}(m^*)$ (hence $h \notin H_i(\hat{s}_i)$), $h \prec h'$ implies $h' \not\prec f_i^{-1}(m^*)$. Therefore,

$$\forall h \in H_i(\hat{\mu}_i) \setminus H_i(\hat{s}_i), \hat{s}_{i,h} = \hat{s}_{i,h}^*$$

and

$$\forall h \in H_i(\hat{\mu}_i) \setminus H_i(\hat{s}_i), V_i(\cdot|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i) = V_i(\cdot|h; \hat{s}_i^*, \hat{\mu}_i, \bar{\phi}_i).$$

By $(\hat{\mu}_i, \bar{\phi}_i)$ -unimprovability of \hat{s}_i^* , this implies

$$\forall h \in H_i(\hat{\mu}_i) \setminus H_i(\hat{s}_i), \hat{s}_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i).$$

Next consider any $h \in H_i(\hat{\mu}_i) \cap H_i(\hat{s}_i)$, hence any $h \prec f_i^{-1}(m^*)$. Belief confirmation (for both μ_i and $\hat{\mu}_i$), payoffs observability, and the definition of \hat{s}_i imply

$$u_i(z^*) = V_i(s_{i,h}^*|h; s_i^*, \mu_i, \bar{\phi}_i).$$

By definition of \hat{s}_i ,

$$V_i(\hat{s}_{i,h}|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i) = u_i(z^*) = V_i(s_{i,h}^*|h; s_i^*, \mu_i, \bar{\phi}_i).$$

By Lemma 9,

$$V_i(s_{i,h}^*|h; s_i^*, \mu_i, \bar{\phi}_i) \geq V_i(\hat{s}_{i,h}^*|h; \hat{s}_i^*, \hat{\mu}_i, \bar{\phi}_i).$$

Hence:

$$V_i(\hat{s}_{i,h}|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i) \geq V_i(\hat{s}_{i,h}^*|h; \hat{s}_i^*, \hat{\mu}_i, \bar{\phi}_i).$$

By definition of \hat{s}_i and $(\hat{\mu}_i, \bar{\phi}_i)$ -unimprovability of \hat{s}_i^*

$$\forall a_i \in A_i(h) \setminus \{\hat{s}_{i,h}\}, V_i(\hat{s}_{i,h}^*|h; \hat{s}_i^*, \hat{\mu}_i, \bar{\phi}_i) \geq V_i(a_i|h; \hat{s}_i^*, \hat{\mu}_i, \bar{\phi}_i) = V_i(a_i|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i),$$

where the equality holds because $\hat{s}_{i,h} = \hat{s}_{i,h}^*$ and \hat{s}_i coincides with \hat{s}_i^* after deviations from an expected path. Collecting these equalities and inequalities, we obtain

$$\forall h \in H_i(\hat{\mu}_i) \cap H_i(\hat{s}_i), \hat{s}_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i).$$

Therefore \hat{s}_i is $(\hat{\mu}_i, \bar{\phi}_i)$ -unimprovable. ■

10.5 Full Unimprovability and Rationalizable SCE

For the reader's convenience we recall some notation and definitions of Section 7. $\bar{H}_i = H_i \cup \{\{\emptyset\}\}$ is the extended collection of information sets that includes the initial information set $\{\emptyset\}$ even if i is not a first mover ($\bar{H}_i = H_i$ if i is a first mover). With this, we let $\bar{H}_i(\mu_i) = H_i(\mu_i) \cup \{\{\emptyset\}\}$ be the subcollection of possible information sets given prior μ_i . A CPS $\bar{\mu}_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i}) \subseteq [\Delta(\Sigma_{-i})]^{\bar{H}_i}$ for player i specifies an initial belief $\bar{\mu}_i(\cdot|\{\emptyset\}) \in \Delta(\Sigma_{-i})$ and a conditional belief $\bar{\mu}_i(\cdot|h)$ for each $h \in H_i$. Full $(\bar{\mu}_i(\cdot|\cdot), \phi_i)$ -unimprovability requires value-maximization over actions at each $h \in H_i$, that is, at the information sets where i is active. The initial belief $\bar{\mu}_i(\cdot|\{\emptyset\})$ matters to compare full $(\bar{\mu}_i(\cdot|\cdot), \phi_i)$ -unimprovability with (μ_i, ϕ_i) -unimprovability, because in the latter μ_i has to be interpreted as an initial belief. Furthermore $\bar{\mu}_i(\cdot|\{\emptyset\})$ captures how i strategically analyzes the game before playing it. Finally, recall that here we assume that there are *no chance moves*.

10.5.1 Equivalence of SCE and fully unimprovable SCE

Proof of Proposition 9. First note that, for every prior μ_i on Σ_{-i} , one can find a CPS $\bar{\mu}_i(\cdot|\cdot)$ on (Σ_{-i}, \bar{H}_i) such that $\mu_i(\cdot|h) = \bar{\mu}_i(\cdot|h)$ for all information sets $h \in \bar{H}_i(\mu_i)$: Let $\bar{\mu}_i(\cdot|\{\emptyset\}) = \mu_i$ and derive $\bar{\mu}_i(\cdot|h)$ by conditioning for all $h \in H_i(\mu_i)$. Next, for every $h \in H_i \setminus H_i(\mu_i)$ whose immediate predecessor h in (\bar{H}_i, \preceq) belongs to $\bar{H}_i(\mu_i)$, fix some $\bar{\mu}_i(\cdot|h) \in \Delta(\Sigma_{-i}(h))$ such that $p_{\bar{\mu}_i(\cdot|h)}(S_{-i}(h')) > 0$ for all the information sets h' that weakly follow h (e.g., $\bar{\mu}_i(\cdot|h) = \frac{1}{|S_{-i}(h)|} \sum_{s_{-i} \in S_{-i}(h)} \delta_{\delta_{s_{-i}}}$) and derive $\bar{\mu}_i(\cdot|h')$ from $\bar{\mu}_i(\cdot|h)$ by conditioning. One can check that the constructed array $(\bar{\mu}_i(\cdot|h))_{h \in \bar{H}_i} \in [\Delta(\Sigma_{-i})]^{\bar{H}_i}$ is a CPS.

Fix an SCE σ justified by confirmed beliefs $(\mu_{s_i})_{i \in I, s_i \in \text{Supp}\sigma_i}$. For each $i \in I$ and $s_i \in \text{Supp}\sigma_i$, let $\bar{\mu}_{s_i}(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ be a CPS such that $\bar{\mu}_{s_i}(\cdot|h) = \mu_{s_i}(\cdot|h)$ for every $h \in H_i(\mu_{s_i})$ (see above). Now construct a new strategy $\hat{s}_i(s_i)$ so that (1) $\hat{s}_i(s_i)(h) = s_i(h)$ for all $h \in H_i(\mu_{s_i})$, and (2) $\hat{s}_i(s_i)$ is derived by folding back on $H_i \setminus H_i(\mu_{s_i})$ given $\bar{\mu}_{s_i}(\cdot|\cdot)$. Since s_i is (μ_{s_i}, ϕ_i) -unimprovable, $\hat{s}_i(s_i)$ must be *fully* $(\bar{\mu}_{s_i}(\cdot|\cdot), \phi_i)$ -unimprovable. By construction, (s_i, μ_i) and $(\hat{s}_i(s_i), \bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$ imply the same probabilities of terminal nodes, because, they reach the same information sets, $H_i(s_i) \cap H_i(\mu_{s_i}) = H_i(\hat{s}_i(s_i)) \cap H_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$, where s_i and $\hat{s}_i(s_i)$ prescribe the same actions.

By ex post perfect recall and the self-confirming conditions for σ , for every $i \in I$ and $s_i \in \text{Supp}\sigma_i$, (s_i, μ_{s_i}) and (s_i, σ_{-i}) yield the same probabilities of reaching information sets of i : $p_{\mu_{s_i}}(S_{-i}(h)) = \sigma_{-i}(S_{-i}(h))$ for every $h \in H_i(s_i)$. Therefore, by construction, $\times p_{\bar{\mu}_{s_i}}(S_{-i}(h) | \{\emptyset\}) = \sigma_{-i}(S_{-i}(h))$ for every $i \in I$, $s_i \in \text{Supp}\sigma_i$, and $h \in H_i(s_i)$. Now, for every $i \in I$, consider the pushforward measure $\hat{\sigma}_i = \sigma_i \circ \hat{s}_i^{-1}$, that is,

$$\forall s_i \in S_i, \hat{\sigma}_i(s_i) = \sum_{s'_i: \hat{s}_i(s'_i) = s_i} \sigma_i(s'_i).$$

By construction σ and $\hat{\sigma}$ yield the same distribution over terminal nodes, because they reach the same information sets and, for each $i \in I$, the pure strategies in the support of σ_i take the same actions as the associated pure strategies in the support of $\hat{\sigma}_i$ at all reachable information sets. Furthermore, the profile $\left(\hat{\sigma}_i, (\bar{\mu}_{\bar{s}_i}(\cdot|\cdot))_{\bar{s}_i \in \text{Supp}\hat{\sigma}_i}\right)_{i \in I}$, where $\bar{\mu}_{\bar{s}_i}(\cdot|\cdot) = \bar{\mu}_{s_i}(\cdot|\cdot)$ for some $s_i \in \hat{s}_i^{-1}(\bar{s}_i)$,⁴³ also satisfies the confirmed beliefs condition on top of the full unimprovability condition. Therefore $\hat{\sigma}$ is a fully unimprovable SCE. \blacksquare

10.5.2 Monotonicity of the symmetric RSCE correspondence

Recall that in the analysis of symmetric RSCE, we assume that there are no chance moves.

Proof of Remark 5 (Only if) Let $\bar{B}_i = \bigcap_{k \in \mathbb{N}} B_i^k$ for each $i \in I$; it can be checked that $(\bar{B}_i)_{i \in I}$ satisfies the required property.

(If) We show by induction that if $(\bar{B}_i)_{i \in I} \in \times_{i \in I} 2^{S_i \times M}$ has the required property, then $\bar{B}_i \subseteq B_i^k$ for every $i \in I$ and $k \in \mathbb{N}_0$. The claim is trivially true for $k = 0$. Suppose it is true for $k \in \mathbb{N}_0$. Fix $i \in I$ and $(\hat{s}_i, \hat{m}_i) \in \bar{B}_i$ arbitrarily. Then $(\hat{s}_i, \hat{m}_i) \in B_i^k$ (inductive hypothesis), and there is $p_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(S_{-i})$ with $\hat{s}_i \in r_i(p_i(\cdot|\cdot), \phi_i)$ such that eq. (10) holds. By the inductive hypothesis, $\bar{B}_{-i} \subseteq B_{-i}^k$; hence

$$\begin{aligned} & p_i \left(F_{\hat{s}_i}^{-1}(\hat{m}_i) \cap \left\{ s_{-i} : (s_j, F_j(\hat{s}_i, s_{-i}))_{j \neq i} \in B_{-i}^k \right\} | \{\emptyset\} \right) \\ & \geq p_i \left(F_{\hat{s}_i}^{-1}(\hat{m}_i) \cap \left\{ s_{-i} : (s_j, F_j(\hat{s}_i, s_{-i}))_{j \neq i} \in \bar{B}_{-i} \right\} | \{\emptyset\} \right), \end{aligned}$$

and (10) implies

$$p_i \left(F_{\hat{s}_i}^{-1}(\hat{m}_i) \cap \left\{ s_{-i} : (s_j, F_j(\hat{s}_i, s_{-i}))_{j \neq i} \in B_{-i}^k \right\} | \{\emptyset\} \right) = 1.$$

Therefore, $(\hat{s}_i, \hat{m}_i) \in B_i^{k+1}$. \blacksquare

⁴³It can be checked that it does not matter which $s_i \in \hat{s}_i^{-1}(\bar{s}_i)$ we pick.

Proof of Lemma 2 The statement is trivially true for $k = 0$. Suppose, by way of induction, that

$$\text{symSCE}^k(\Gamma, f, \phi) = \left\{ \bar{s} : (\bar{s}_i, F_i(\bar{s}))_{i \in I} \in B^k \right\}.$$

We first show that, for every fixed \bar{s} in the above set and $i \in I$

$$F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} = F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k \right\}. \quad (16)$$

(*Proof of \subseteq*) By definition of section and the inductive hypothesis,

$$\begin{aligned} \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} &= \left\{ s_{-i} : (\bar{s}_i, s_{-i}) \in \text{symSCE}^k(\Gamma, f, \phi) \right\} \\ &= \left\{ s_{-i} : ((\bar{s}_i, F_i(\bar{s}_i, s_{-i})), (s_j, F_j(\bar{s}_i, s_{-i})))_{j \in I \setminus \{i\}} \in B^k \right\} \\ &\subseteq \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \neq i} \in B_{-i}^k \right\}. \end{aligned}$$

Hence

$$F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} \subseteq F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k \right\}.$$

(*Proof of \supseteq*) Let $\hat{s}_{-i} \in F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k \right\}$. Since $\hat{s}_{-i} \in F_{\bar{s}_i}^{-1}(F_i(\bar{s}))$ and $(\bar{s}_i, F_i(\bar{s})) \in B_i^k$, then $F_i(\bar{s}_i, \hat{s}_{-i}) = F_i(\bar{s})$ and

$$((\bar{s}_i, F_i(\bar{s}_i, \hat{s}_{-i})), (\hat{s}_j, F_j(\bar{s}_i, \hat{s}_{-i})))_{j \in I \setminus \{i\}} \in B^k.$$

Hence, $\hat{s}_{-i} \in \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i}$. Thus

$$F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k \right\} \subseteq F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i}.$$

This completes the proof of (16). \square

Now, let $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \phi)$; then – by definition – $\bar{s} \in \text{symSCE}^k(\Gamma, f, \phi)$ and the inductive hypothesis implies that $((\bar{s}_i, F_i(\bar{s})))_{i \in I} \in B^k$. Let $(p_i(\cdot|\cdot))_{i \in I}$ be as in the definition of $\text{symSCE}^{k+1}(\Gamma, f, \phi)$. Then, by eq. (16),

$$p_i \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \left\{ s_{-i} : (s_j, F_j(\bar{s}_i, s_{-i}))_{j \in I \setminus \{i\}} \in B_{-i}^k \right\} \mid \{\emptyset\} \right) = 1$$

for every $i \in I$, and $(\bar{s}_i, F_i(\bar{s}))_{i \in I} \in B^{k+1}$. Similarly, let $(\bar{s}_i, F_i(\bar{s}))_{i \in I} \in B^{k+1}$; then, by eq. (16),

$$p_i \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} \mid \{\emptyset\} \right) = 1$$

for every $i \in I$, and $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \phi)$. \blacksquare

Our results about monotonicity of the RSCE correspondence rely on the following result of *monotonicity of the justifiability correspondence*: Consider a decision problem under uncertainty $(\hat{A}, \hat{S}, \hat{u})$ where \hat{A} and \hat{S} are finite sets of actions and states, and $\hat{u} : \hat{A} \times \hat{S} \rightarrow \mathbb{R}$ is a vNM utility function. Fix a non-empty and *compact* set of distributions $\hat{\Sigma} \subseteq \Delta(\hat{S})$. Let $\mathcal{IC} \subseteq \mathbb{R}^{\mathbb{R}}$ denote the set of *strictly increasing* and *continuous* functions. For each $\phi \in \mathcal{IC}$ and $\mu \in \Delta(\hat{\Sigma})$ define the set of (μ, ϕ) -best replies:

$$\hat{r}(\mu, \phi) = \arg \max_{a \in \hat{A}} \int_{\hat{\Sigma}} \phi(\mathbb{E}_{\sigma} [\hat{u}(a, \cdot)]) \mu(d\sigma).$$

Similarly, for each $p \in \Delta(\hat{S})$ we write

$$\hat{r}(p, \phi) = \arg \max_{a \in \hat{A}} \sum_{s \in \hat{S}} \phi(\hat{u}(a, s)) p(s),$$

which is the special case when μ is supported by Dirac beliefs and p is the corresponding predictive belief. If an action is a (μ, ϕ) -best reply we say that it is ϕ -**justified** by μ ; we say that it is ϕ -**justifiable** if it is ϕ -justified by some μ . Battigalli et al. (2016a) proved that the ϕ -justifiability correspondence is monotone with respect to concave and strictly increasing transformations (see also Weinstein, 2016):

Lemma 10 *For all $\bar{\phi}, \tau \in \mathcal{IC}$, if $\tau \in \mathcal{IC}$ is a concave transformation then*

$$\bigcup_{\mu \in \Delta(\hat{\Sigma})} \hat{r}(\mu, \tau \circ \bar{\phi}) \supseteq \bigcup_{\bar{\mu} \in \Delta(\hat{\Sigma})} \hat{r}(\bar{\mu}, \bar{\phi}).$$

Since $\hat{S} \cong \{\delta_s : s \in \hat{S}\} \subseteq \Delta(\hat{S})$, letting $\hat{\Sigma} = \{\delta_s : s \in \hat{S}\}$ in Lemma 10 we obtain the following:

Corollary 14 *For all $\bar{\phi}, \tau \in \mathcal{IC}$, if $\tau \in \mathcal{IC}$ is a concave transformation then*

$$\bigcup_{p \in \Delta(\hat{S})} \hat{r}(p, \tau \circ \bar{\phi}) \supseteq \bigcup_{\bar{p} \in \Delta(\hat{S})} \hat{r}(\bar{p}, \bar{\phi}).$$

From now on, whenever we refer to games in which every player moves at most once along every path, strategies will be omitted from the value formulas. In this class of games, the value of an action at an information set depends only on the action itself and not on the overall strategy of the agent.

Proof of Theorem 10 To prove the result, we will show that

$$\forall k \in \mathbb{N}, \text{symSCE}^k(\Gamma, f, \bar{\phi}) \subseteq \text{symSCE}^k(\Gamma, f, \phi).$$

Then the claim follows from Lemma 2. The statement is trivially true for $k = 0$. Suppose, by way of induction, that

$$\text{symSCE}^k(\Gamma, f, \bar{\phi}) \subseteq \text{symSCE}^k(\Gamma, f, \phi). \quad (\text{I.H.})$$

Fix $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \bar{\phi})$ and $i \in I$ arbitrarily. By definition of symSCE^{k+1} and the inductive hypothesis (I.H.), there is $\bar{p}_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(S_{-i})$ such that \bar{s}_i is fully $(\bar{p}_i(\cdot|\cdot), \bar{\phi}_i)$ -unimprovable and

$$\text{Supp}\bar{p}_i(\cdot|\{\emptyset\}) \subseteq \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \bar{\phi})_{\bar{s}_i} \right) \stackrel{(\text{I.H.})}{\subseteq} \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} \right).$$

We construct $p_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ such that $p_i(\cdot|h) = \bar{p}_i(\cdot|h)$ for each $h \in \bar{H}_i(\bar{p}_i(\cdot|\{\emptyset\}))$ and \bar{s}_i is fully $(p_i(\cdot|\cdot), \phi_i)$ -unimprovable. Since

$$\text{Supp}p_i(\cdot|\{\emptyset\}) = \text{Supp}\bar{p}_i(\cdot|\{\emptyset\}) \subseteq \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \phi)_{\bar{s}_i} \right)$$

and the construction holds for each i , this implies that $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \phi)$.

(*Construction of $p_i(\cdot|\cdot)$*) Since i moves at most once on every path, $H_i(s_i) = H_i$; furthermore, the value of any action $a_i \in A_i(h)$ ($h \in H_i$) is independent of i 's strategy. To construct $p_i(\cdot|\cdot)$, keep the same initial belief as \bar{p}_i : $p_i(\cdot|\{\emptyset\}) = \bar{p}_i(\cdot|\{\emptyset\})$. By symmetry (pure equilibrium), no chance moves, confirmed beliefs and ex-post perfect recall, there is a unique $h \in H_i(\bar{p}_i(\cdot|\{\emptyset\})) = H_i(p_i(\cdot|\{\emptyset\}))$ that i expects to reach with probability 1 under \bar{p}_i . Thus, let $p_i(\cdot|h) = \bar{p}_i(\cdot|\{\emptyset\}) = \bar{p}_i(\cdot|h)$ for $h \in H_i(\bar{p}_i(\cdot|\{\emptyset\}))$ even if $h \neq \{\emptyset\}$. It follows that for the unique $h \in H_i(p_i(\cdot|\{\emptyset\}))$ and for every $a_i \in A_i(h)$,

$$\begin{aligned} V_i(\bar{s}_{i,h}|h; p_i, \phi_i) &\stackrel{(\text{L4})}{=} U_i(\bar{s}) \\ &\stackrel{(\text{L4})}{=} V_i(\bar{s}_{i,h}|h; p_i, \bar{\phi}_i) \\ &\stackrel{(\text{uprv.})}{\geq} V_i(a_i|h; p_i, \bar{\phi}_i) \\ &\stackrel{(\text{L5})}{\geq} V_i(a_i|h; p_i, \phi_i), \end{aligned}$$

where the equalities follow from $h \in H_i(\bar{p}_i(\cdot|\{\emptyset\})) = H_i(p_i(\cdot|\{\emptyset\}))$ and (given the observability of payoffs) Lemma 4, the first inequality follows from $(p_i(\cdot|\{\emptyset\}), \bar{\phi}_i)$ -unimprovability of \bar{s}_i (uprv.) and the second one from Lemma 5.

Now, consider any $h \in H_i \setminus H_i(\bar{p}_i(\cdot|\{\emptyset\}))$. Since \bar{s}_i is fully $(\bar{p}_i(\cdot|\cdot), \bar{\phi}_i)$ -unimprovable, we have

$$\bar{s}_{i,h} \in \arg \max_{a_i \in A_i(h)} V_i(a_i|h; \bar{p}_i, \bar{\phi}_i) = \arg \max_{a_i \in A_i(h)} \mathbb{E}_{\bar{p}_i(\cdot|h)} [\bar{\phi}_i \circ U_{h,a_i}].$$

By Corollary 14 there exists some $p_{i,h} \in \Delta(S_{-i}(h))$ such that

$$\bar{s}_{i,h} \in \arg \max_{a_i \in A_i(h)} \mathbb{E}_{p_{i,h}} [\varphi \circ \bar{\phi}_i \circ U_{h,a_i}] = \arg \max_{a_i \in A_i(h)} \mathbb{E}_{p_{i,h}} [\phi_i \circ U_{h,a_i}].$$

Let $p_i(\cdot|h) = p_{i,h}$. By the one-move assumption, h has no strict predecessors or followers in H_i . Then, the array $(p_i(\cdot|h))_{h \in \bar{H}_i}$ is a CPS. By construction \bar{s}_i is fully $(p_i(\cdot|\cdot), \phi_i)$ -unimprovable. \square

We conclude that $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \phi)$. \blacksquare

Proof of Theorem 11 Throughout the proof, for each $s = (s_j)_{j \in I}$, let $[s] = \times_{j \in I} [s_j]$ and $[s_{-i}] = \times_{j \neq i} [s_j]$. To prove the result, we will show that for every $k \in \mathbb{N}$ and $\bar{s} \in \text{symSCE}^k(\Gamma, f, \bar{\phi})$, there exists $s \in \text{symSCE}^k(\Gamma, f, \phi)$ with $s \in [\bar{s}]$. Then the claim follows from Lemma 2. The statement is trivially true for $k = 0$.

Suppose, by way of induction, that the statement is true for k . Let

$$\bar{s} = (\bar{s}_i)_{i \in I} \in \text{symSCE}^{k+1}(\Gamma, f, \bar{\phi}) \subseteq \text{symSCE}^k(\Gamma, f, \bar{\phi})$$

(the inclusion holds by definition). By the inductive hypothesis, for every $i \in I$, there exists some $s^* = (s_i^*)_{i \in I} \in \text{symSCE}^k(\Gamma, f, \phi)$ such that $s^* \in [\bar{s}]$. Since $\bar{s} \in \text{symSCE}^{k+1}(\Gamma, f, \bar{\phi})$, for every $i \in I$ there is some CPS $\bar{p}_i(\cdot|\cdot) \in \Delta^{\bar{H}_i}(S_{-i})$ such that \bar{s}_i is fully $(\bar{p}_i(\cdot|\cdot), \bar{\phi}_i)$ -unimprovable and

$$\text{Supp} \bar{p}_i(\cdot|\{\emptyset\}) \subseteq \left(F_{\bar{s}_i}^{-1}(F_i(\bar{s})) \cap \text{symSCE}^k(\Gamma, f, \bar{\phi})_{\bar{s}_i} \right).$$

It can be checked that, for any $s = (s_i, s_{-i}) \in \text{symSCE}^k(\Gamma, f, \phi)$, if $s_i \in [s_i^*]$, then $(s_i^*, s_{-i}) \in \text{symSCE}^k(\Gamma, f, \phi)$ as well. Then, by the inductive hypothesis, for every $s_{-i} \in \text{symSCE}^k(\Gamma, f, \bar{\phi})_{\bar{s}_i}$, there exists $s'_{-i} \in \text{symSCE}^k(\Gamma, f, \phi)_{s_i^*}$ such that $s'_{-i} \in [s_{-i}]$. Moreover, since strategic feedback depends only on the realization equivalence classes of strategies, for each $s_{-i} \in F_{\bar{s}_i}^{-1}(F_i(\bar{s}))$ and $s'_{-i} \in [s_{-i}]$, we have $s'_{-i} \in F_{s_i^*}^{-1}(F_i(s^*))$. So, we can construct a belief $p_i^* \in \Delta(S_{-i})$ with

$$\text{Supp} p_i^* \subseteq \left(F_{s_i^*}^{-1}(F_i(s^*)) \cap \text{symSCE}^k(\Gamma, f, \phi)_{s_i^*} \right)$$

such that for each $\hat{s}_{-i} \in \text{Supp} p_i^*$,

$$p_i^*([\hat{s}_{-i}]) = \bar{p}_i([\hat{s}_{-i}]|\{\emptyset\}). \quad (17)$$

Since $\bar{p}_i(\cdot|\{\emptyset\})$ is the predictive probability isomorphic to a belief over Dirac models, by Corollary 13 player i 's preferences are dynamically consistent. Thus,⁴⁴

$$\bar{s}_i \in \arg \max_{s_i \in S_i} \sum_{s_{-i}} \bar{\phi}_i(U_i(s_i, s_{-i}) \bar{p}_i(s_{-i}|\{\emptyset\})).$$

⁴⁴Recall that, for every $p \in \Delta(S)$ and every strictly increasing function ϕ ,

$$\arg \max_{s_i \in S_i} \phi^{-1}(\mathbb{E}_p[\phi(U_{i,s_i})]) = \arg \max_{s_i \in S_i} \mathbb{E}[\phi(U_{i,s_i})].$$

But then, by eq. (17), we also have

$$\bar{s}_i \in \arg \max_{s_i \in \hat{S}_i} \sum_{s_{-i}} \bar{\phi}_i(U_i(s_i, s_{-i})p_i^*(s_{-i})). \quad (18)$$

Consider the decision problem $(\hat{A}, \hat{S}, \hat{u})$ where $\hat{A} = S_i$, $\hat{S} = F_{s_i^*}^{-1}(F_i(s^*)) \cap \text{symSCE}^k(\Gamma, f, \phi)_{s_i^*}$, and \hat{u} is the restriction of $\bar{\phi}_i \circ U_i$ on $\hat{A} \times \hat{S}$. By eq. (18), \bar{s}_i is $\bar{\phi}_i$ -justified by $p_i^* \in \Delta(\hat{S})$.

By Corollary 14, there exists $p_i \in \Delta(\hat{S})$ that ϕ_i -justifies \bar{s}_i . Let $p_i(\cdot | \{\emptyset\}) = p_i$. Then

$$p_i(\cdot | \{\emptyset\}) \in \Delta\left(F_{s_i^*}^{-1}(F_i(s^*)) \cap \text{symSCE}^k(\Gamma, f, \phi)_{s_i^*}\right)$$

and

$$\bar{s}_i \in \arg \max_{s_i \in \hat{S}_i} \sum_{s_{-i}} \phi_i(U_i(s_i, s_{-i})p_i(s_{-i} | \{\emptyset\})).$$

By dynamic consistency of player i 's preferences and standard arguments, this implies that for every $s_i \in [\bar{s}_i]$, for every $h' \in H_i(p_i(\cdot | \{\emptyset\})) \cap H_i(\bar{s}_i)$,

$$s_{i,h'} \in \arg \max_{a_i \in A_i(h')} \sum_{s_{-i}} \phi_i(U_i((s_i|_{h'}, a_i), s_{-i}|h')p_i(s_{-i}|h')),$$

where $p_i(s_{-i}|h')$ is derived by conditioning.

Next, fix $h \in H_i(\bar{s}_i) \setminus H_i(p_i(\cdot | \{\emptyset\}))$ with $h' \in H_i(p_i(\cdot | \{\emptyset\}))$ for all $h' \prec h$. By dynamic consistency, we can frame the continuation of the game as a (static) decision problem $(\hat{A}, \hat{S}, \hat{u})$, with $\hat{A} = S_i(h)$, $\hat{S} = S_{-i}(h)$ and \hat{u} is the restriction of $\bar{\phi}_i \circ U_i$ on $\hat{A} \times \hat{S} = S_i(h) \times S_{-i}(h)$. By the same argument as above, since \bar{s}_i is $\bar{\phi}_i$ -justified by $\bar{p}_i(\cdot | h)$, Corollary 14 implies that \bar{s}_i is ϕ_i -justified by some $p_i(\cdot | h) \in \Delta(S_{-i}(h))$. By dynamic consistency of player i 's preferences and standard arguments, this implies that for every $s_i \in [\bar{s}_i]$, for every $h' \in H_i(p_i(\cdot | h)) \cap H_i(\bar{s}_i)$ with $h \succeq h'$,

$$s_{i,h'} \in \arg \max_{a_i \in A_i(h')} \sum_{s_{-i}} \phi_i(U_i((s_i|_{h'}, a_i), s_{-i}|h')p_i(s_{-i}|h')),$$

where $p_i(s_{-i}|h')$ is derived from $p_i(s_{-i}|h)$ by conditioning.

Repeating iteratively the operation at all information sets that are not reached with positive probability under the probability measures already constructed, one can finally construct a CPS $\hat{p}_i(\cdot | \cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ such that $\hat{p}_i(\cdot | h) = p_i(\cdot | h)$ for each $h \in H_i(\bar{s}_i)$ with $p_i(S_{-i}(h)|h') = 0$ for all $h' \prec h$. Since every $s_i \in [\bar{s}_i]$ is $(\hat{p}_i(\cdot | \cdot), \phi_i)$ -unimprovable at every $h \in H_i(\bar{s}_i)$, there exists a fully $(\hat{p}_i(\cdot | \cdot), \phi_i)$ -unimprovable $s_i \in [\bar{s}_i] = [s_i^*]$. Let $s = (s_i)_{i \in I}$. By construction,

$$\hat{p}_i(\cdot | \{\emptyset\}) \in \Delta\left(F_{s_i^*}^{-1}(F_i(s^*)) \cap \text{symSCE}^k(\Gamma, f, \phi)_{s_i^*}\right).$$

By $s_i \in [s_i^*]$ and $s \in [s^*]$, $F_{s_i^*}^{-1}(F_i(s^*)) = F_{s_i}^{-1}(F_i(s))$. So, $s \in \text{symSCE}^1(\Gamma, f, \phi)$. If $k > 0$, for all $s'_{-i} \in S_{-i}$, $(s_i^*, s'_{-i}) \in \text{symSCE}^1(\Gamma, f, \phi)$ if and only if $(s_i, s'_{-i}) \in \text{symSCE}^1(\Gamma, f, \phi)$. So, $\text{symSCE}^1(\Gamma, f, \phi)_{s_i^*} = \text{symSCE}^1(\Gamma, f, \phi)_{s_i}$. But then, $\hat{p}_i(\cdot | \{\emptyset\}) \in \Delta(F_{s_i}^{-1}(F_i(s)) \cap \text{symSCE}^1(\Gamma, f, \phi)_{s_i})$, and so $s \in \text{symSCE}^2(\Gamma, f, \phi)$. Inductively,

$$\hat{p}_i(\cdot | \{\emptyset\}) \in \Delta\left(F_{s_i}^{-1}(F_i(s)) \cap \text{symSCE}^k(\Gamma, f, \phi)_{s_i}\right),$$

and so $s \in \text{symSCE}^{k+1}(\Gamma, f, \phi)$. ■

10.5.3 Monotonicity of the RSCE correspondence

We first use Lemma 10 to prove a preliminary monotonicity result for the set of fully unimprovable SCEs.

Lemma 11 *Fix two games with observable payoffs where no player moves more than once, (Γ, f, ϕ) and $(\Gamma, f, \bar{\phi})$, so that (Γ, f, ϕ) features more ambiguity aversion than $(\Gamma, f, \bar{\phi})$. Then $\text{SCE}^1(\Gamma, f, \bar{\phi}) \subseteq \text{SCE}^1(\Gamma, f, \phi)$.*

Proof of Lemma 11 Let $\bar{\sigma} \in \text{SCE}^1(\Gamma, f, \bar{\phi})$. Fix $i \in I$, $s_i \in \text{Supp}\bar{\sigma}_i$ and let $\bar{\mu}_{s_i}(\cdot | \cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ be such that s_i is fully $(\bar{\mu}_{s_i}(\cdot | \cdot), \bar{\phi}_i)$ -unimprovable and

$$\text{Supp}\bar{\mu}_{s_i}(\cdot | \{\emptyset\}) \subseteq \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}) \stackrel{(\text{def.})}{=} \left\{ \sigma_{-i} \in \Sigma_{-i} : \hat{F}_i(s_i, \sigma_{-i}) = \hat{F}_i(s_i, \bar{\sigma}_{-i}) \right\}.$$

We will construct $\mu_{s_i}(\cdot | \cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ such that $\mu_{s_i}(\cdot | h) = \bar{\mu}_{s_i}(\cdot | h)$ for each $h \in \bar{H}_i(\bar{\mu}_{s_i}(\cdot | \{\emptyset\}))$, and s_i is fully $(\mu_{s_i}(\cdot | \cdot), \phi_i)$ -unimprovable. Since

$$\text{Supp}\bar{\mu}_{s_i}(\cdot | \{\emptyset\}) = \text{Supp}\mu_{s_i}(\cdot | \{\emptyset\}) \subseteq \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}),$$

this implies that $\bar{\sigma} \in \text{SCE}^1(\Gamma, f, \phi)$.

Recall that $H_i(s_i) = H_i$ because i moves at most once on every path. To construct $\mu_{s_i}(\cdot | \cdot)$, keep the same prior belief: $\mu_{s_i}(\cdot | \{\emptyset\}) = \bar{\mu}_{s_i}(\cdot | \{\emptyset\})$. For each $h \in H_i(\bar{\mu}_{s_i}(\cdot | \{\emptyset\})) = H_i(\mu_{s_i}(\cdot | \{\emptyset\}))$ define $\mu_{s_i}(\cdot | h)$ by conditioning. Hence, for each $h \in \bar{H}_i(\bar{\mu}_{s_i}(\cdot | \{\emptyset\}))$, $\mu_{s_i}(\cdot | h) = \bar{\mu}_{s_i}(\cdot | h)$. Thus, for every $a_i \in A_i(h)$,

$$\begin{aligned} V_i(s_{i,h}|h; \mu_{s_i}, \phi_i) &\stackrel{(\text{L4})}{=} U_i(s_i, \bar{\sigma}_{-i}|h) \\ &\stackrel{(\text{L4})}{=} V_i(s_{i,h}|h; \mu_{s_i}, \bar{\phi}_i) \\ &\stackrel{(\text{uprv.})}{\geq} V_i(a_i|h; \mu_{s_i}, \bar{\phi}_i) \\ &\stackrel{(\text{L5})}{\geq} V_i(a_i|h; \mu_{s_i}, \phi_i), \end{aligned}$$

where the equalities follow from Lemma 4, the first inequality from full $(\bar{\mu}_{s_i}(\cdot|\cdot), \bar{\phi}_i)$ -unimprovability of s_i (uprv.), and the second one from Lemma 5.

Now, consider any $h \in H_i \setminus H_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$. Since s_i is fully $(\bar{\mu}_{s_i}(\cdot|\cdot), \bar{\phi}_i)$ -unimprovable and $\bar{\mu}_{s_i}(\Sigma_{-i}(h)|h) = 1$ by condition (1) of Definition 6, we have

$$\begin{aligned} s_{i,h} &\in \arg \max_{a_i \in A_i(h)} \int_{\Sigma_{-i}} \bar{\phi}_i(U_i((s_i|_h, a_i), \sigma_{-i}|h)) \bar{\mu}_{s_i}(d\sigma_{-i}|h) \\ &= \arg \max_{a_i \in A_i(h)} \int_{\Sigma_{-i}(h)} \bar{\phi}_i \left(\sum_{s_{-i} \in S_{-i}(h)} \sigma_{-i}(s_{-i}|h) u_i(\zeta(a_i, s_{-i})) \right) \bar{\mu}_{s_i}(d\sigma_{-i}|h), \end{aligned} \quad (19)$$

where we abuse notation and write $\zeta(a_i, s_{-i})$, because i does not move before h , nor after h , hence the terminal node reached depends only on $s_{-i} \in S_{-i}(h)$ and the action chosen by i at h . Next, we consider the problem of choice under uncertainty $(\hat{A}, \hat{S}, \hat{u})$ where $A = A_i(h)$, $\hat{S} = S_{-i}(h)$, $\hat{u}(a_i, s_{-i}) = u_i(\zeta(a_i, s_{-i}))$. Recall that we assumed *observable deviators*, hence

$$S_{-i}(h) = \times_{j \neq i} S_j(h)$$

and it makes sense to consider the following compact set of product distributions

$$\Sigma_{-i|h} = \{\sigma_{-i} \in \Sigma_{-i} : \sigma_{-i}(\times_{j \neq i} S_j(h)) = 1\} \subseteq \Sigma_{-i}(h).$$

The following map associates each $\sigma_{-i} \in \Sigma_{-i}(h)$ with the corresponding updated distribution $\sigma_{-i}(\cdot|h)$ defined in (1):

$$\begin{aligned} \varsigma_{h,-i} : \Sigma_{-i}(h) &\rightarrow \Sigma_{-i|h} \\ \sigma_{-i} &\mapsto \mathbf{1}_{S_{-i}(h)}(\cdot) \frac{\sigma_{-i}(\cdot)}{\sigma_{-i}(S_{-i}(h))} \end{aligned}$$

$[\mathbf{1}_{S_{-i}(h)}]$ denotes the indicator function of set $S_{-i}(h)$. Note that $\Sigma_{-i}(h)$ is a (relatively) open subset of the Polish space Σ_{-i} and $\varsigma_{h,-i}$ is continuous. Furthermore, the restriction $\varsigma_{h,-i}|_{\Sigma_{-i|h}}$ is the identity on $\Sigma_{-i|h}$, because $\sigma_{-i}(S_{-i}(h)) = \sigma_{-i}(\times_{j \neq i} S_j(h)) = 1$ implies

$$\sigma_{-i}(s_{-i}|h) = \frac{\sigma_{-i}(s_{-i})}{\sigma_{-i}(S_{-i}(h))} = \sigma_{-i}(s_{-i})$$

for each s_{-i} . Therefore $\varsigma_{h,-i}$ is also onto and $\varsigma_{h,-i}$ is a measurable surjection that yields the *onto* pushforward map

$$\begin{aligned} \hat{\varsigma}_{h,-i} : \Delta(\Sigma_{-i}(h)) &\rightarrow \Delta(\Sigma_{-i|h}), \\ \mu_i &\mapsto \mu_i \circ \varsigma_{h,-i}^{-1}. \end{aligned}$$

Let $\bar{\mu}_{s_i|h} = \bar{\mu}_{s_i}(\cdot|h) \circ \varsigma_{h,-i}^{-1} \in \Delta(\Sigma_{-i|h})$, that is,

$$\forall E \in \mathcal{B}(\Sigma_{-i|h}), \quad \bar{\mu}_{s_i|h}(E) = \bar{\mu}_{s_i}(\varsigma_{h,-i}^{-1}(E)|h).$$

For every $a_i \in A_i(h)$, we have:

$$\int_{\Sigma_{-i}(h)} \bar{\phi}_i \left(\sum_{s_{-i} \in S_{-i}(h)} \sigma_{-i}(s_{-i}|h) v_i(\gamma(\zeta(a_i, s_{-i}))) \right) \bar{\mu}_{s_i}(\mathrm{d}\sigma_{-i}|h) =$$

$$\int_{\Sigma_{-i}|h} \bar{\phi}_i \left(\sum_{s_{-i} \in S_{-i}(h)} \sigma_{-i}(s_{-i}|h) v_i(\gamma(\zeta(a_i, s_{-i}))) \right) \bar{\mu}_{s_i|h}(\mathrm{d}\sigma_{-i}(\cdot|h)).$$

Lemma 10, eq. (19), and the above equality imply that $s_{i,h}$ is ϕ_i -justified by some belief $\mu_{i|h} \in \Delta(\Sigma_{-i}|h)$, that is,

$$s_{i,h} \in \arg \max_{a_i \in A_i(h)} \int_{\Delta(\Sigma_{-i}|h)} \phi_i \left(\sum_{s_{-i} \in S_{-i}(h)} \sigma_{-i}(s_{-i}|h) v_i(\gamma(\zeta(a_i, s_{-i}))) \right) \mu_{i|h}(\mathrm{d}\sigma_{-i}(\cdot|h)).$$

Now go back to a belief on $\Sigma_{-i}(h)$: Since the pushforward map $\hat{\varsigma}_{h,-i}$ is onto we can find some belief $\mu_{s_i}(\cdot|h) \in \hat{\varsigma}_{h,-i}^{-1}(\mu_{i|h}) \subseteq \Delta(\Sigma_{-i}(h))$ such that

$$s_{i,h} \in \arg \max_{a_i \in A_i(h)} \int_{\Delta(\Sigma_{-i}(h))} \phi_i \left(\sum_{s_{-i} \in S_{-i}(h)} \sigma_{-i}(s_{-i}|h) v_i(\gamma(\zeta(a_i, s_{-i}))) \right) \mu_{s_i}(\mathrm{d}\hat{\sigma}_{-i}|h)$$

$$= \arg \max_{a_i \in A_i(h)} \int_{\Sigma_{-i}} \phi_i(U_i((s_i|h, a_i), \sigma_{-i}|h)) \mu_{s_i}(\mathrm{d}\sigma_{-i}|h).$$

We can do this for every off-path information set $h \in H_i \setminus H_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$ and obtain an array $(\mu_{s_i}(\cdot|h))_{h \in \bar{H}_i} \in \times_{h \in \bar{H}_i} \Delta(\Sigma_{-i}(h))$ (recall that μ_{s_i} coincides with $\bar{\mu}_{s_i}$ on $H_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$) that satisfies condition (1) of the definition of CPS by construction. Since every player moves at most once on every path, condition (2) trivially holds. Finally, by construction, s_i is fully $(\mu_i(\cdot|\cdot), \phi_i)$ -unimprovable. \blacksquare

Comment: What matters in the previous proof is that the range $\varsigma_{h,-i}(\Sigma_{-i}(h))$ of the continuous map $\varsigma_{h,-i}$ is a compact subset of $\Delta(S_{-i}(h))$ even though $\Sigma_{-i}(h)$ is not closed. We used the assumption of observable deviators to prove it. There are simple examples of games without observable deviators where $\varsigma_{h,-i}(\Sigma_{-i}(h))$ is not compact and therefore we cannot apply Lemma 10. However, we have not yet found examples of games without observable deviators where players move at most once and monotonicity of $\phi \mapsto SCE^1(\Gamma, f, \bar{\phi})$ (hence, monotonicity of RSCE) does not hold.

Proof of Theorem 12 Lemma 11 shows that the result holds for $k = 1$. Suppose by way of induction that $SCE^k(\Gamma, f, \bar{\phi}) \subseteq SCE^k(\Gamma, f, \phi)$. Fix $\sigma \in SCE^{k+1}(\Gamma, f, \bar{\phi})$, $i \in I$,

$s_i \in \text{Supp}\sigma_i$ arbitrarily and let $\bar{\mu}_{s_i}(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ be a confirmed CPS that $\bar{\phi}_i$ -justifies s_i , that is, s_i is fully $(\bar{\mu}_{s_i}(\cdot|\cdot), \bar{\phi})$ -unimprovable and

$$\bar{\mu}_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \bar{\phi}) | \{\emptyset\} \right) = 1.$$

Since $\sigma \in SCE^{k+1}(\Gamma, f, \bar{\phi}) \subseteq SCE^k(\Gamma, f, \bar{\phi})$, the inductive hypothesis implies $\sigma \in SCE^k(\Gamma, f, \phi)$, so there is $\hat{\mu}_{s_i}(\cdot|\cdot) \in \Delta^{\hat{H}_i}(\Sigma_{-i})$ such that s_i is fully $(\hat{\mu}_{s_i}(\cdot|\cdot), \phi_i)$ -unimprovable. Let $\mu_{s_i}(\cdot|\cdot) \in \Delta^{\bar{H}_i}(\Sigma_{-i})$ be such that $\mu_{s_i}(\cdot|h) = \bar{\mu}_{s_i}(\cdot|h)$ for each $h \in \bar{H}_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$ and $\mu_{s_i}(\cdot|h) = \hat{\mu}_{s_i}(\cdot|h)$ for each $h \in H_i \setminus \bar{H}_i(\bar{\mu}_{s_i}(\cdot|\{\emptyset\}))$. By same argument employed in the proof of Lemma 11, $\mu_{s_i}(\cdot|\cdot)$ is a CPS and s_i is $(\mu_{s_i}(\cdot|\cdot), \phi_i)$ -unimprovable. Moreover, by construction and the inductive hypothesis,

$$\begin{aligned} & \mu_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \phi) | \{\emptyset\} \right) \\ & \geq \mu_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \bar{\phi}) | \{\emptyset\} \right) \\ & = \bar{\mu}_{s_i} \left(\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) \cap \text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \bar{\phi}) | \{\emptyset\} \right) = 1. \end{aligned}$$

Therefore s_i is ϕ_i -justified by a confirmed CPS that initially believes $\text{proj}_{\Sigma_{-i}} SCE^k(\Gamma, f, \phi)$. This holds for every $i \in I$ and $\text{Supp}\sigma_i$, thus, $\sigma \in SCE^{k+1}(\Gamma, f, \phi)$. \blacksquare

References

- [1] BATTIGALLI, P. (1987): *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, unpublished thesis, Università Bocconi.
- [2] BATTIGALLI, P. (1999): “A Comment On Non-Nash equilibria,” typescript, European University Institute.
- [3] BATTIGALLI, P. (2003): “Rationalizability in Infinite, Dynamic Games of Incomplete Information,” *Research in Economics*, 57, 1-38.
- [4] BATTIGALLI, P. (2017): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University.
- [5] BATTIGALLI, P., AND G. BONANNO (1999): “Synchronic Information and Common Knowledge in Extensive Games,” *Research in Economics*, 53, 77-100..
- [6] BATTIGALLI, P. AND D. GUAITOLI (1988): “Conjectural Equilibria and Rationalizability in a Macroeconomic Game with Incomplete Information,” *Quaderni di Ricerca*, Università Bocconi, 1988 (published in *Decisions, Games and Markets*, Kluwer, Dordrecht, 97-124, 1997).

- [7] BATTIGALLI P., M. GILLI AND M.C. MOLINARI (1992): "Learning and Convergence to Equilibrium in Repeated Strategic Interaction," *Research in Economics*, 46, 335-378.
- [8] BATTIGALLI, P., M. SINISCALCHI, AND A. FRIEDENBERG (2017): *Strategic Uncertainty: An Epistemic Approach to Game Theory*, Book in Preparation for MIT Press.
- [9] BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2015): "Self-Confirming Equilibrium and Model Uncertainty," *American Economic Review*, 105, 646-677.
- [10] BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2016a): "A Note on Comparative Ambiguity Aversion and Justifiability," *Econometrica*, 84, 1903-1916.
- [11] BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2016b): "Analysis of Information Feedback and Selfconfirming Equilibrium," *Journal of Mathematical Economics*, 66, 40-51.
- [12] BATTIGALLI, P., A. FRANGETICH, G. LANZANI, AND M. MARINACCI (2016c): "Learning and Self-Confirming Long Run Biases," IGIER Working Paper 588, Bocconi University.
- [13] BLUME, L., AND D. EASLEY (2006): "If You're so Smart, why aren't You Rich? Belief Selection in Complete and Incomplete markets," *Econometrica*, 74, 929-966.
- [14] BRANDTS, J., AND G. CHARNESS (2011): "The Strategy versus the Direct-Response Method: A First Survey of Experimental Comparisons," *Experimental Economics*, 14, 375-398.
- [15] CERREIA-VIOGLIO S., F. MACCHERONI, M. MARINACCI, AND L. MONTRUCCHIO (2013): "Classical Subjective Expected Utility," *Proceedings of the National Academy of Sciences*, 110, 6754-6759.
- [16] DEKEL, E., D. FUDENBERG AND D.K. LEVINE (1999): "Payoff Information and Self-confirming Equilibrium," *Journal of Economic Theory*, 89, 165-185.
- [17] DEKEL, E., D. FUDENBERG, AND D. LEVINE (2004): "Learning to Play Bayesian Games," *Games and Economic Behavior*, 46, 282-303.
- [18] EPSTEIN, L.G., AND M. SCHNEIDER (2007): "Learning Under Ambiguity," *The Review of Economic Studies*, 74, 1275-1303.
- [19] ESPONDA, I. (2013): "Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions," *Theoretical Economics* 8, 467-501.

- [20] FUDENBERG, D. AND Y. KAMADA (2015): “Rationalizable Partition-Confirmed Equilibrium,” *Theoretical Economics* 10, 775–806.
- [21] FUDENBERG, D. AND Y. KAMADA (2017): “Rationalizable Partition-Confirmed Equilibrium with Heterogeneous Beliefs,” mimeo.
- [22] FUDENBERG, D. AND D.K. LEVINE (1993): “Self-Confirming Equilibrium,” *Econometrica*, 61, 523-545.
- [23] FUDENBERG, D. AND D.K. LEVINE (2006): “Superstition and Rational Learning,” *American Economic Review*, 96, 630-651.
- [24] HANANY, E., AND P. KLIBANOFF (2009): “Updating Ambiguity Averse Preferences,” *The B.E. Journal of Theoretical Economics*, 9.
- [25] HANANY, E., P. KLIBANOFF, AND S. MUKERJI (2017): “Incomplete Information Games with Ambiguity Averse Players,” mimeo.
- [26] KALAI, E., AND E. LEHRER (1993a): “Rational Learning Leads to Nash Equilibrium,” *Econometrica*, 61, 1019-1045.
- [27] KALAI, E., AND E. LEHRER (1993b): “Subjective Equilibrium in Repeated Games,” *Econometrica*, 61, 11231-1240.
- [28] KALAI, E., AND E. LEHRER (1995): “Subjective Games and Equilibria,” *Games and Economic Behavior*, 8, 123-163.
- [29] KLIBANOFF P., M. MARINACCI AND S. MUKERJI (2005): “A Smooth Model of Decision Making Under Ambiguity,” *Econometrica*, 73, 1849-1892.
- [30] KUHN H.W. (1953): “Extensive Games and the Problem of Information,” in: H.W. KUHN AND A.W. TUCKER (Eds.), *Contributions to the Theory of Games II*, Princeton University Press, Princeton NJ, 193-216.
- [31] MARINACCI, M. (2015): “Model Uncertainty,” *Journal of the European Economic Association*, 13, 998-1076.
- [32] OSBORNE M. AND A. RUBINSTEIN (1994): *A Course in Game Theory*, Cambridge MA: MIT Press.
- [33] RUBINSTEIN A. AND A. WOLINSKY (1994): “Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability,” *Games and Economic Behavior*, 6, 299-311.
- [34] SCHELLING, T.C. (1960): *The Strategy of Conflict*. Cambridge MA: Harvard University Press.

- [35] SINISCALCHI, M. (2011): “Dynamic Choice Under Ambiguity,” *Theoretical Economics*, 6, 379-421.
- [36] WEINSTEIN, J. (2016): “The Effect of Changes of Risk Attitudes on Strategic Behavior,” *Econometrica*, 84, 1881-1902.