Games and Economic Behavior 115 (2019) 1–29

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Ambiguity attitudes and self-confirming equilibrium in sequential games ☆

P. Battigalli [a,*], E. Catonini [b], G. Lanzani [c], M. Marinacci [a]

[a] *Department of Decision Sciences and IGIER – Università Bocconi, Italy*
[b] *ICEF – National Research University Higher School of Economics, Russian Federation*
[c] *Department of Economics, Massachusetts Institute of Technology, United States of America*

ARTICLE INFO

Article history:
Received 21 March 2018
Available online 12 February 2019

JEL classification:
C72
C73
D81
D83

Keywords:
Sequential games with feedback
Smooth ambiguity
Self-confirming equilibrium

ABSTRACT

We consider a game in extensive form recurrently played by agents who are randomly drawn from large populations and matched. We assume that preferences over actions at any information set admit a smooth-ambiguity representation in the sense of Klibanoff et al. (2005), which may induce dynamic inconsistencies. We take this into account in our analysis of self-confirming equilibrium (SCE) given players' feedback about the path of play. Battigalli et al. (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 SCE.

1. Introduction

Self-confirming equilibrium When a game is played recurrently and the learning dynamic has reached a rest point, each agent chooses a (one-period) best reply to his subjective belief, which may be incorrect and yet confirmed by the evidence available to him. A profile of strategies and beliefs with this property is a self-confirming equilibrium (henceforth, SCE).[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.[3] This is particularly relevant when agents only have scarce evidence of opponents' behavior and face therefore strategic uncertainty. Following Battigalli et al. (2015, henceforth BCMM), we analyze the SCE concept in games played by ambiguity averse agents. Unlike BCMM,

☆ This is an abridged version of the working paper by Battigalli et al. (2017) with the same title. We are grateful to three anonymous referees for their detailed and insightful observations. We also thank Sarah Auster, Federico Bobbio, Roberto Corrao, Enrico De Magistris, Nicodemo De Vito, David Ruiz, Thomas J. Sargent, 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), ESAM 2017 (Hong Kong), and ESEM 2017 (Lisbon), and an anonymous associate editor for useful comments. The study has been funded by ERC grant 324219, by the Guido Cazzavillan Scholarship, and within the framework of the Basic Research Program at the National Research University Higher School of Economics (HSE) and by the Russian Academic Excellence Project '5-100.

* Corresponding author.
E-mail address: pierpaolo.battigalli@unibocconi.it (P. Battigalli).

[1] This terminology is due to Fudenberg and Levine (1993).
[2] For a perspective on ambiguity attitudes that applies to this paper, see Cerreia-Vioglio et al. (2013) and the survey by Marinacci (2015).
[3] See, for example, the survey by Trautmann and van de Kuilen (2016).

https://doi.org/10.1016/j.geb.2019.02.005
0899-8256/

who essentially restrict their attention to simultaneous-move games, we consider games with sequential moves, represented in extensive form. 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 our paper adds to the previous literature in general and to BCMM in particular.

SCE and model uncertainty 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 et al. (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.[4]

SCE is defined as follows: Let I denote the set of player roles (e.g., buyer and seller) and S_i 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 an SCE 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 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 (e.g., one's own monetary payoff) are observed by each agent after each play, BCMM prove a **monotonicity result**: higher ambiguity aversion entails a larger set of equilibria. Intuitively, for each agent, the strategy played repeatedly in equilibrium yields known risks because the agent observes its long-run distribution of payoffs, while deviations are untested and may be perceived as ambiguous; therefore, higher ambiguity aversion penalizes deviations relative to 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.

Our contribution 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"[5]—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.[6] Thus, the fact that agents in sequential games cannot (irreversibly) choose strategies must be faced and dealt with explicitly. We assume that agents are sophisticated: They understand their future contingent incentives and choose actions in early stages predicting how such incentives determine their actions in later stages, that is, they plan by "folding back" given their subjective conditional beliefs about the strategies of other players. Hence, they execute "unimprovable" strategies, that is, those that satisfy the one-deviation property. 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.

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 monotonicity result extend to games with sequential moves? To elaborate on the first question, fix a sequential game with feedback (Γ, f), where f represents what each player observes ex post about the path of play. This determines the normal (or strategic) form $(G, F) = \mathcal{N}(\Gamma, f)$, where G is given by the strategic-form payoff functions, and F is the strategic-form feedback derived from f, so that two strategy profiles inducing the same path are necessarily indistinguishable. Under subjective expected utility maximization, which is dynamically consistent, SCE in the extensive form (Γ, f) is equivalent to SCE in the normal form (G, F). 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.

This leads us to the second question. Since 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. 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. Indeed, as we demonstrate in examples, there may be instances of non-monotonicity when the SCE strategies at some baseline level of ambiguity aversion are such that some agents at some reachable information sets would be willing to pay to commit—if they only could—on a different course of action involving multiple sequential deviations from the folding-back plan. We prove a preliminary result, Lemma 1, stating that such willingness to commit is necessary for non-monotonicity.

[4] See the survey by Marinacci (2015) and the discussion of Cubitt et al. (2019), who test experimentally the different behavioral implications of the smooth-ambiguity model and other models of ambiguity attitudes.

[5] See the survey by Brandts and Charness (2011) and the references therein.

[6] Unlike *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.

Lemma 1 allows us to 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) there cannot be multiple sequential deviations from folding-back planning. In case (ii) we show that, loosely speaking, in equilibrium ambiguity aversion is not distinguishable from the risk aversion of expected utility maximizers, who are dynamically consistent. Lemma 1 also implies the following 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.

Related literature This is the first paper that analyzes SCE with non-neutral ambiguity attitudes in sequential games. BCMM and Battigalli et al. (2016b) only consider games played in strategic form. The extensive form game from which the strategic form is derived only affects feedback: two strategy profiles that yield the same terminal node necessarily yield the same feedback to every player. Furthermore, Battigalli et al. (2016b) focus on "maxmin SCE" a notion of equilibrium with extreme ambiguity aversion (see Gilboa and Schmeidler, 1989), whereas here—like BCMM—we apply the smooth ambiguity criterion of KMM. Hereafter we call "standard SCE" the self-confirming equilibrium concept with ambiguity neutrality, i.e., subjective expected utility maximization.

A version of the standard SCE concept was first put forward in the undergraduate thesis of Battigalli (1987)[8] and called "conjectural equilibrium." The first definition in English appears in a working paper by Battigalli and Guaitoli (1988), who also define and apply an extensive-form rationalizability refinement. Fudenberg and Levine (1993) independently put forward a definition of standard SCE under the maintained assumption that agents perfectly observe ex-post the path of play. Unlike Battigalli (1987) and Battigalli and Guaitoli (1988), the analysis of Fudenberg and Levine applies to large population games. This is important because it justifies the assumption that agents maximize their short-run expected utility even if they are patient. The reason is that in a large population game there are no incentives to affect the future behavior of current opponents, who are almost surely different from the future ones.[9] For further references and details see the discussion in Section IV of BCMM and the survey by Battigalli et al. (1992).

Fudenberg and Levine (1993) note that standard SCE is not strategic-form invariant, arguing that the strategic—or normal—form is therefore insufficient to characterize SCE. 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 strategic form). Therefore, when they compare standard SCE in the extensive (i.e., sequential) and strategic form, they change what players can observe ex post about the behavior of others: only on-path actions in the extensive form, and complete strategy profiles in the strategic form. If feedback about the behavior of others is instead kept constant—as we do when we look at strategic-form feedback—standard SCE in the extensive form is equivalent to standard SCE in the strategic form (see Remark 3). On the other hand, we show that SCE with ambiguity aversion is *not* strategic-form invariant (see Example 4).

As mentioned above, BCMM proved that in simultaneous-move games the SCE correspondence is monotone with respect to players' degree of ambiguity aversion. Here we prove partial versions of this result for sequential games. A similar monotonicity result can be proved for comparative risk aversion in the particular case of *pure* standard SCE (see the lecture notes of Battigalli, 2018), and for comparative risk or ambiguity aversion in the case of rationalizable strategies (Battigalli et al., 2016a). Note that the intuition and proof in the case of rationalizability are very different from the SCE case.[10]

Next we compare our contribution to other papers on equilibria with ambiguity aversion in sequential games. Unlike us, these papers analyze versions or refinements of the Nash equilibrium concept, that is, they assume that each player's conjecture about the "decision rules" of other players is correct. Thus, they consider different sources of ambiguity. Lo (1999) analyzes a notion of Nash equilibrium whereby players are mini-maximizers given belief sets in the sense of Gilboa and Schmeidler (1989). Assuming Bayesian updating[11] for each measure in each player's belief set, preferences at different information sets may be dynamically inconsistent. As Lo points out, this implies that the proposed equilibrium concept is not strategic-form invariant. We differ from Lo (1999) in several ways: we adopt the smooth ambiguity criterion, we analyze a version of SCE rather than Nash equilibrium, and we define rational planning as unimprovability (folding back), whereas Lo imposes a form of sequential—or interim—optimality (mini-maximization with respect to continuation strategies at each reachable information set). For similar reasons, our work differs from Riedel and Sass (2014, 2017), who apply the maxmin criterion to analyze equilibria of games where players may delegate the selection of actions to random devices with unknown outcome probabilities ("Ellsberg urns"), and put forward conditions on the extensive form that ensure the dynamic consistency of the ex ante optimal strategy. Hanany et al. (2018) use the smooth ambiguity criterion to analyze sequential games with incomplete information. Their work differs from ours in two important ways. First, they analyze a notion of perfect Bayesian equilibrium, which assumes that players' conjectures about the type-dependent strategies of other players

[7] We 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.

[8] Written in Italian.

[9] Furthermore, incentives to experiment vanish in the long run.

[10] Interestingly, in the working paper we rely on Battigalli et al. (2016a) to prove (in some special cases) monotonicity of the rationalizable SCE correspondence.

[11] Or maximum likelihood updating.

are correct, while beliefs about the distribution of types are not disciplined; we instead allow for incorrect beliefs about strategies, the confirmed-beliefs requirement partially disciplines these beliefs, and the limited observability of opponents' behavior is the source of the residual ambiguity. Second, Hanany et al. model subjectively rational planning in a different way. We maintain that players' beliefs satisfy the standard rules of conditional probability, which imply Bayes rule; given this, each player adopts an unimprovable strategy that—by the dynamic inconsistency of ambiguity averse preferences—may be suboptimal according to his initial preferences, or the interim preferences he holds at some information set. Hanany et al. instead assume that each player's equilibrium strategy is sequentially optimal, that is, optimal according to his preferences at each information set. Ambiguity aversion implies that this cannot always be reconciled with Bayesian updating. However, they show it is without loss of generality to restrict attention to beliefs generated using a particular generalization of Bayesian updating, where more ambiguity aversion implies larger deviations from standard Bayesian updating. Eichberger et al. (2017) analyze a notion of equilibrium with ambiguity aversion in sequential games that combines Choquet-expected utility maximization (Schmeidler, 1989) with a generalized form of Bayesian updating. Unlike Lo (1999) and Hanany et al. (2018), and like us, they assume folding-back planning. To conclude the comparison, we emphasize that our approach can easily accommodate incomplete information and commitment to randomizations with unknown probabilities, as we explain in Section 7. Yet, from our perspective, the source of ambiguity is the residual uncertainty that persists in the long run because of imperfect learning. Thus, for example, the probabilities of the states of nature in a game with incomplete information played recurrently remain unknown only if players' feedback does not reveal the realized state of nature (cf. Dekel et al., 2004). Similar considerations apply to randomizations with unknown probabilities.

Outline 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; finally, Section 7 discusses the relevance of some assumptions and equilibrium concepts, and provides hints for generalization and extension of the analysis. The main text contains some intuitive arguments, but all formal proofs are collected in Appendix A.

2. Framework

We analyze an agent with non-neutral ambiguity attitudes who plays a game with sequential moves. Such moves occur in one period of time and consequences realize at the end of the period. Our interpretation is that the game is played recurrently for infinitely many periods in a population-game scenario. 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 reached. We also assume that he is sophisticated and therefore he takes this into account when he plans how to play the game. We need *not* assume that the game is common knowledge. We postpone comments on this aspect until after our definition of equilibrium, so that they can be better understood.

Consider an agent who plays in role $i \in I$ of a *finite* **extensive-form game** Γ with *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 $\varnothing$ denote the root of the game; then $\{\varnothing\} \in H_i$ if and only if i is a first mover.[12] We endow H_i with the weak (respectively, strict) precedence relation $\preceq$ ($\prec$) inherited from the game tree.[13] 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, $s_i = \left(s_{i,h}\right)_{h \in H_i}$. We let $S := \times_{j \in I} S_j$ denote the set of pure strategy profiles of the players.

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" $\bar{\beta}_0 \in \times_{h \in H_0} \Delta(A_0(h))$, with $\bar{\beta}_0(a_0|h) > 0$ for every $h \in H_0$ and $a_0 \in A_0(h)$; $\bar{\beta}_0$ induces the "mixed strategy" $\bar{\sigma}_0 \in \Delta(S_0)$ such that

$$\bar{\sigma}_0(s_0) = \prod_{h \in H_0} \bar{\beta}_0\left(s_{0,h}|h\right) > 0$$

for every $s_0 \in S_0$. Thus, the outcome distributions respectively induced by $\bar{\beta}_0$ and $\bar{\sigma}_0$ coincide for every strategy profile of the true players (see Kuhn, 1953). We let $I_0 := I \cup \{0\}$ denote the set of true player-roles *and* chance, and we let $-i := I_0 \backslash \{i\}$. In particular, $S_{-i} := \times_{j \in I_0 \backslash \{i\}} S_j$ denotes the set of pure strategy profiles of the opponents including chance.

[12] Our preferred representation of games in extensive form starts from sequences of actions, that correspond to the nodes of the game tree (e.g., Chapters 6 and 11 of Osborne and Rubinstein, 1994). This affects the way we draw pictures and describe examples, but it is otherwise irrelevant for the analysis of the paper, because our notation is consistent with all standard representations of games.

[13] Perfect 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 $h \prec h'$, 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 $\prec$ is $\preceq$, an antisymmetric and transitive relation that makes H_i a directed forest. If $\{\varnothing\} \in H_i$, then $(H_i, \preceq)$ is a directed tree.

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**

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

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**

$$\gamma : Z \to C$$

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 $c = (c_j)_{j\in I} \in C$ 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 **von Neumann–Morgenstern utility** function

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

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

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

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

$$S_{I_0}(x) := \{(s_0, s) \in S_0 \times S : x \prec \zeta(s_0, s)\}$$

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

$$S_{I_0}(h) := \bigcup_{x\in h} S_{I_0}(x)$$

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

$$S_{-i}(h) := \mathrm{proj}_{S_{-i}} S_{I_0}(h) = \{s_{-i} \in S_{-i} : \exists\, (x, s_i) \in h \times S_i, x \prec \zeta(s_i, s_{-i})\}$$

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

$$S_i(h) := \mathrm{proj}_{S_i} S_{I_0}(h) = \{s_i \in S_i : \exists (x, s_{-i}) \in h \times S_{-i}, x \prec \zeta(s_i, s_{-i})\}$$

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

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

$$\forall h \in H_i, S_{I_0}(h) = S_i(h) \times S_{-i}(h).$$

In words, the information about behavior encoded in $h \in H_i$ can be decomposed into information about own behavior and information about the opponents' behavior, because i remembers what he knew and did at earlier nodes.

Players are drawn at random from large populations of agents, each playing a pure strategy. The **distribution of pure strategies in population** $j \in I$ is denoted by $\sigma_j \in \Delta(S_j)$ and is unknown to the other players. For any profile of distributions $\left(\sigma_j\right)_{j\in I}$, we let[14]

$$\sigma_{-i} := \bar{\sigma}_0 \times \left(\times_{j\in I\setminus\{i\}}\sigma_j\right) \in \Delta(S_{-i})$$

denote the product distribution over strategies of chance and of i's opponents. From the viewpoint of an agent in role i, $\sigma_{-i}(s_{-i})$ is the unknown *objective* probability of facing opponents and chance playing pure strategy profile s_{-i}.[15] We let

$$\Sigma_{-i} := \left\{\sigma_{-i} \in \Delta\left(S_{-i}\right) : \exists \left(\sigma_j\right)_{j\in I\setminus\{i\}} \in \times_{j\in I\setminus\{i\}}\Delta\left(S_j\right), \sigma_{-i} = \bar{\sigma}_0 \times \left(\times_{j\in I\setminus\{i\}}\sigma_j\right)\right\}$$

denote the set of these product distributions. 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 $\mathcal{B}(\Sigma_{-i})$.

[14] We use symbol $\times$ to denote both the Cartesian product of sets and the product of measures.

[15] Statistical independence follows from random matching: For each $i \in I$, let P_i denote the set of agents playing in role i, and let $\varsigma_i : P_i \to S_i$ 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. Taking into account that player i has a known deterministic plan, knows chance probabilities, and knows that opponents are randomly matched, we obtain the product formula in the text.

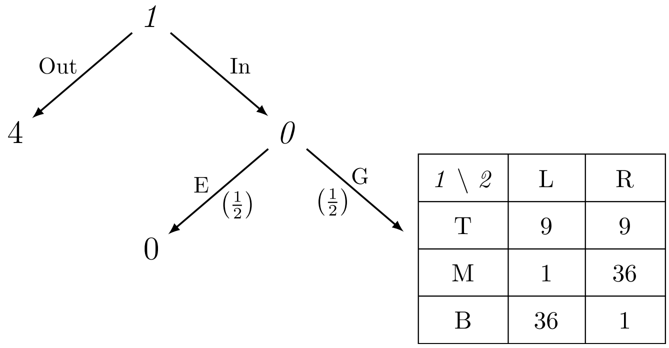

Fig. 1. Running example.

At each point of the game, the agent playing in role i has some **belief** $\mu_i \in \Delta(\Sigma_{-i})$. Note that by the definition of Σ_{-i}, the marginal of μ_i on Σ_0 always assigns probability 1 to $\bar{\sigma}_0$. 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 $s_{-i} \in S_{-i}$,

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

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

Primitive elements	Terminology
$i, j \in I$ ($I_0 := I \cup \{0\}$)	players ($0 \notin I$ denotes chance)
$h \in H_i$	information sets of i
$\preceq$ ($\prec$)	(strict) precedence relation of Γ
$(H_i, \preceq)$	directed forest of information sets of i
$z \in Z$	terminal histories/nodes
$\gamma : Z \to C$	consequence function
$v_i : C \to R$	von Neumann–Morgenstern utility function of i
$a_i \in A_i(h)$	i's actions at $h \in H_i$

Derived elements	Terminology
$s_i \in S_i := \times_{h \in H_i} A_i(h)$	strategies of i
$s \in S$ ($s_{-i} \in S_{-i}$)	strategy profiles (of $-i := I_0 \backslash \{i\}$)
$S_{I_0}(h)$	strategy profiles (including 0) reaching h
$S_i(h) := \mathrm{proj}_{S_i} S_{I_0}(h)$	strategies of i allowing for h
$S_{-i}(h) := \mathrm{proj}_{S_{-i}} S_{I_0}(h)$	strategy profiles of $-i$ (including 0) 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}, with $\mathrm{marg}_{\Delta(S_0)}\sigma_{-i} = \bar{\sigma}_0$
$\mu_i \in \Delta(\Sigma_{-i})$	beliefs of i
$p_{\mu_i} \in \Delta(S_{-i})$	predictive probabilities implied by μ_i
$\zeta : S_0 \times S \to Z$	outcome function
$u_i := v_i \circ \gamma : Z \to \mathbb{R}$	payoff function of i

Example 1 *(Running example, notation).* The game depicted in Fig. 1 is a two-person, multistage game where $I = \{1, 2\}$, 0 is chance, and the numbers at terminal histories/nodes, including the boxes in the matrix subgame, are the payoffs of Player 1. Player 1 is participating to a game show. He has already won an equivalent of 4 utils and can decide to stop (Out) or accept (In) a bet with known odds that can destroy his win (E) or multiply it (G) to an amount worth 9 utils. If he accepts the bet and the bet is successful, he can decide to stop (T) or play a coordination game with Player 2, who qualified to the final stage of the game as well.[16] If the two players coordinate ((M, R) or (B, L)), they win an amount worth 36 utils, otherwise ((M, L) or (B, R)) they go home with a consolation prize worth 1 util. If we identify nodes with histories,[17] information sets, actions sets and terminal histories/nodes are as follows:

[16] It does not matter whether at this stage player 1 chooses among three actions or makes two sequential binary choices (T or ¬T, with the latter followed by a 2×2 subgame). We assume the former for simplicity.
[17] See, e.g., Chapter 11 in Osborne and Rubinstein (1994).

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

The probabilities of chance moves are $\bar{\sigma}_0(\text{E}) = \bar{\sigma}_0(\text{G}) = \frac{1}{2}$.

Assume that the agent in role 1 has the following belief about the behavior of player 2[18]:

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

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 $p_{\mu_1}(\text{E}, \text{L}) = p_{\mu_1}(\text{E}, \text{R}) = p_{\mu_1}(\text{G}, \text{R}) = p_{\mu_1}(\text{G}, \text{L}) = \frac{1}{4}$. ▲

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 contingent 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., his contingent plan is deterministic.

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,

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

Note that whether $\sigma_{-i} \in \Sigma_{-i}(h)$ or not depends only on $(\sigma_j)_{j \in I \setminus \{i\}}$ because $\sigma_0 = \bar{\sigma}_0$ is strictly positive. We now define the **strategic-form von Neumann–Morgenstern conditional expected utility** function

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

where, given $\sigma_{-i} \in \Sigma_{-i}(h)$, $\sigma_{-i}(\cdot|h)$ denotes the objective conditional distribution on the opponents' strategy profiles consistent with h.

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 $U_i(s_i, \sigma_{-i}|h)$.[19]

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 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; that is,

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

where $\alpha_i(h', h)$ is the action chosen at $h' \prec h$ in order to reach h.[20] Finally, we let $s_{i|h}$ denote the replacement plan obtained when action $s_{i,h}$ is played at h:

[18] In the examples, we will always specify beliefs about the true opponents only, taking for granted that $\bar{\sigma}_0$ is commonly known.

[19] One 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 s_i and σ_{-i}.

[20] By perfect recall, α_i is well defined.

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

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 von Neumann–Morgenstern utility function $v_i : C \to \mathbb{R}$ 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

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

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. Function ϕ_i captures DM$_i$'s ambiguity attitudes. It is analogous to the utility of money in the theory of choice under risk, with objective expected utility values replacing monetary values. In particular, DM$_i$ is ambiguity averse if ϕ_i is concave. 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[21]:

$$\bar{V}_i(a_i|h; s_i, \mu_i, \phi_i) := \phi_i^{-1}\left(\int_{\Sigma_{-i}(h)} \phi_i\left(U_i\left((s_{i|h}, a_i), \sigma_{-i}|h\right)\right)\mu_i(\mathrm{d}\sigma_{-i})\right). \tag{1}$$

Condition $\mu_i(\Sigma_{-i}(h)) = 1$ suggests that, in eq. (1), 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

$$\bar{V}_i(a_i|h; s_i, \mu_i, \phi_i) = \phi_i^{-1}\left(\phi_i\left(U_i\left((s_{i|h}, a_i), \sigma_{-i}|h\right)\right)\right) = U_i\left((s_{i|h}, a_i), \sigma_{-i}|h\right).$$

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 (1) boils down to the classical subjective expected utility formula if ϕ_i is linear (ambiguity neutrality), hence equivalent to the identity function $\mathrm{Id}_{\mathbb{V}_i}$. We emphasize in our notation only the dependence of values of i's actions on parameter ϕ_i, not on the von Neumann–Morgenstern 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.

Example 2 *(Running example, value).* In the game of Fig. 1,

$$U_1(\text{In.B}, \bar{\sigma}_0 \times \delta_{\text{L}} | \{(\text{In}, \text{G})\}) = 36,$$
$$U_1(\text{In.B}, \bar{\sigma}_0 \times \delta_{\text{R}} | \{(\text{In}, \text{G})\}) = 1,$$

and

$$U_1(\text{In.B}, \bar{\sigma}_0 \times \delta_{\text{L}} | \{\varnothing\}) = \frac{1}{2} \cdot 36 = 18,$$
$$U_1(\text{In.B}, \bar{\sigma}_0 \times \delta_{\text{R}} | \{\varnothing\}) = \frac{1}{2} \cdot 1 = \frac{1}{2}.$$

Assume $\phi_1(u) = \sqrt{u}$ and consider the (marginal) belief of Example 1, that is

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

then

$$\bar{V}_1(\text{B}| \{(\text{In}, \text{G})\}; s_1, \mu_1, \phi_1) = \left(\frac{1}{2}\sqrt{36} + \frac{1}{2}\sqrt{1}\right)^2 = (3.5)^2 = 12.25$$

for each $s_1 \in S_1$. ▲

[21] For an axiomatization of the smooth ambiguity criterion, see Cerreia-Vioglio et al. (2013). Note that their value function does not include the normalization operation ϕ^{-1}. Obviously, the two criteria are behaviorally equivalent. We use (1) because it allows for a more direct comparison of the values of ambiguous and unambiguous actions.

4. Conditional beliefs and unimprovability

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

$$\forall (s_{-i}, E_{-i}) \in S_{-i} \times \mathcal{B}(\Sigma_{-i}), \qquad \pi_{\mu_i}(\{s_{-i}\} \times E_{-i}) := \int_{E_{-i}} \sigma_{-i}(s_{-i})\mu_i(\mathrm{d}\sigma_{-i}).$$

This probability measure represents the joint belief about the actual strategy profile played by the agent(s) matched with i and the overall distribution in population(s) $-i$. Note that

$$\pi_{\mu_i}(S_{-i} \times E_{-i}) = \mu_i\left(E_{-i}\right),$$

and

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

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

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

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

$$\begin{aligned}\mu_i(E_{-i}|h) &= \pi_{\mu_i}(S_{-i} \times E_{-i}|S_{-i}(h) \times \Sigma_{-i}(h)) \\ &= \frac{\pi_{\mu_i}((S_{-i} \cap S_{-i}(h)) \times (E_{-i} \cap \Sigma_{-i}(h)))}{\pi_{\mu_i}(S_{-i}(h) \times \Sigma_{-i}(h))} \\ &= \frac{\int_{E_{-i}} \sigma_{-i}(S_{-i}(h))\mu_i(\mathrm{d}\sigma_{-i})}{p_{\mu_i}(S_{-i}(h))}.\end{aligned} \tag{2}$$

For example, if μ_i has *finite support*:

$$\mu_i(\sigma_{-i}|h) = \frac{\mu_i(\sigma_{-i}) \cdot \sigma_{-i}(S_{-i}(h))}{\sum_{\sigma'_{-i} \in \mathrm{Supp}\mu_i} \mu_i(\sigma'_{-i}) \cdot \sigma'_{-i}(S_{-i}(h))}.$$

With this, we consider the profile of conditional beliefs $(\mu_i(\cdot|h))_{h \in H_i(\mu_i)}$ derived from prior belief μ_i, we let

$$V_i\left(a_i|h; s_i, \mu_i, \phi_i\right) := \overline{V}_i(a_i|h; s_i, \mu_i(\cdot|h), \phi_i) = \phi_i^{-1}\left(\int_{\Sigma_{-i}(h)} \phi_i\left(U_i\left((s_{i|h}, a_i), \sigma_{-i}|h\right)\right)\mu_i(\mathrm{d}\sigma_{-i}|h)\right)$$

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 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

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

Unimprovability is a kind of "one-deviation property" requiring that the strategy of player i is an *intrapersonal* equilibrium between the "agents" of i playing at different reachable information sets.[22] By finiteness of the game, unimprovability is equivalent to **folding-back planning**: 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 successor in $H_i(\mu_i)$, and to each

[22] As we note below, game theory textbooks consider the analogous property for *strategy profiles* of $|I|$ players. They call "**one-deviation principle**" the theorem stating that in finite games played by expected utility maximizers a given profile is a sequential, equilibrium if and only if it satisfies the one-deviation property (the result extends to most infinite games of interest in applications). We instead consider the subjective sequential decision problem of a *single player* for given beliefs about the co-players. Furthermore, we explain below that the foregoing principle does not apply to such decision problems due to the dynamic inconsistency implied by ambiguity aversion.

one of them he assigns an action $s_{i,h}$ that maximizes $V_i(a_i|h; s_i, \mu_i, \phi_i)$.[23] Then he folds back, considering the information sets $h \in H_i(\mu_i)$ such that every successor has no further successors; for every such successor, 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. This folding-back procedure—possibly breaking some ties arbitrarily—computes an unimprovable strategy for every (μ_i, ϕ_i).[24]

Of course, we could define beliefs, and thus impose optimality requirements, also at information sets in $H_i \backslash H_i(\mu_i)$, but we are not interested in doing so: The moves of DM$_i$ at each $h \in H_i \backslash H(\mu_i)$ are immaterial for the equilibrium outcomes (by ex post perfect recall and confirmed beliefs), and impossible to predict for the opponents as long as we do not assume that they know the payoff function of i.[25]

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,

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

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

We let $[s_i]$ denote the set of 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: $\phi_i = \mathrm{Id}_{\mathbb{V}_i}$. Then, by a classical dynamic programming result, unimprovability is equivalent to "global" (ex ante) subjective EU-maximization. This is an intrapersonal-equilibrium version of the one-deviation principle[26]:

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 $(\mu_i, \mathrm{Id}_{\mathbb{V}_i})$-unimprovable strategy,*
(2) $s_i^* \in \arg\max_{s_i \in S_i} \sum_{s_{-i}} p_{\mu_i}(s_{-i}) u_i(\zeta(s_i, s_{-i}))$.

We introduce the following strengthening of unimprovability:

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

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

[23] Note that, in this particular case, $V_i(a_i|h; s_i, \mu_i, \phi_i)$ does not depend on s_i.

[24] A further refinement can be obtained 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.

[25] In the working paper, we put forward a stronger notion of optimal planning, called "full unimprovability," which considers beliefs and imposes optimality requirements also at information sets that cannot be reached under the initial belief. Full unimprovability, per se, does not refine the set of SCE outcome distributions (see Proposition 10 in the working paper), but it is instrumental for the analysis of strategic reasoning through rationalizable SCE.

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

If DM_i has dynamically consistent preferences over strategies, sequential optimality is equivalent to the traditional definition, which requires maximization also at the root even if DM_i is inactive at the beginning of the game. This equivalence does not hold under dynamic inconsistency, and we find it convenient to use the weaker definition given above.[27] We also point out that, under dynamic inconsistency, a sequentially optimal strategy may not exist, as illustrated in Example 3 below.

If DM_i is ambiguity neutral, unimprovability coincides with sequential optimality. This is an intrapersonal-equilibrium version of the one-deviation principle[28]:

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

The following example illustrates the well known dynamic inconsistency of preferences of decision makers with non-neutral attitudes towards ambiguity.[29,30]

Example 3 *(Running example, dynamic inconsistency).* Consider the game of Fig. 1 and the (marginal) belief of Example 1:

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

Then

$$H_1(\mu_1) = H_1 = \{\{\varnothing\}, \{(\mathrm{In}, \mathrm{G})\}\}.$$

The induced belief on $S_{-1} \times \Sigma_{-1}$ is

$$\pi_{\mu_1}(s_{-1}, \sigma_{-1}) = \begin{cases} \frac{1}{4} & \text{if } (s_{-1}, \sigma_{-1}) \in \{\mathrm{E,G}\} \times \{(\mathrm{L}, \bar{\sigma}_0, \delta_{\mathrm{L}}), (\mathrm{R}, \bar{\sigma}_0, \delta_{\mathrm{R}})\}, \\ 0 & \text{otherwise.} \end{cases}$$

For an ambiguity neutral player 1 with belief μ_1, the value of M and B at $(\mathrm{In}, \mathrm{G})$ is $\left(\frac{1}{2} \cdot 36 + \frac{1}{2} \cdot 1\right) > 9$, so it is higher than the value of T. Therefore, by folding-back, the $(\mu_1, \mathrm{Id}_{\mathbb{V}_1})$-unimprovable strategies are In.M and In.B.

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

$$\bar{\phi}_1(u) = \sqrt{u} \text{ if } 1 \le u \le 36,$$
$$\bar{\phi}_1\left(\frac{1}{2}\right) = -1.$$

At $(\mathrm{In}, \mathrm{G})$, player 1 still prefers M (or B) over T, because:

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

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

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

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

[27] Sequential optimality appears in the hypotheses of some of our results, but never in the theses. Thus, using a weak definition makes our results more general.

[28] Cf., for example, Fudenberg and Tirole (1991, pp. 108–110), and Osborne and Rubinstein (1994, Lemma 98.2).

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

[30] Note 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., 2018). 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.

$$V_1(\text{In}|\{\varnothing\};\text{In.T},\mu_1,\bar{\phi}_1)=\left(\sqrt{\frac{1}{2}\cdot 9}\right)^2>4=V_1(\text{Out}|\{\varnothing\};\text{Out}.a_1,\mu_1,\bar{\phi}_1), \tag{3}$$

for all $a_1 \in \{\text{T, M, B}\}$. So, player 1 would commit to take the first bet (In) and then to stop before the coordination game (T) if he only could. This implies that neither Out.M nor Out.B are sequentially optimal. ▲

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, $f=(f_i : Z \to M)_{i\in I}$ 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 $\phi=(\phi_i : \mathbb{V}_i \to \mathbb{R})_{i\in I}$ is a profile of continuous and strictly increasing functions capturing players' attitudes toward ambiguity.

To relate to BCMM it is convenient to define the normal form of a game with feedback (Γ, f). We first define the strategic form payoff function and feedback function:

$$U_i := u_i \circ \zeta : S_0 \times S \to \mathbb{R},$$
$$F_i := f_i \circ \zeta : S_0 \times S \to M.$$

With this, the **normal-form** (expected) **payoff function** of player i is

$$\begin{aligned}\bar{U}_i :\ & S \ \to\ \mathbb{R},\\ & s \ \mapsto\ \textstyle\sum_{s_0\in S_0}\bar{\sigma}_0(s_0)U_i(s_0,s).\end{aligned}$$

Similarly, we define **normal-form feedback function** $\bar{F}_i : S \to \Delta(M)$ 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,

$$\forall (s,m) \in S \times M, \bar{F}_i(s)(m) = \sum_{s_0 : F_i(s_0,s)=m} \bar{\sigma}_0(s_0).$$

With this, the normal form of (Γ, f) is $\mathcal{N}(\Gamma, f) := \left(S_i, \bar{U}_i, \bar{F}_i\right)_{i\in I}$. The equilibrium concept of BCMM applies to $(\mathcal{N}(\Gamma, f), \phi) = \left(S_i, \bar{U}_i, \bar{F}_i, \phi_i\right)_{i\in I}$ 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) also satisfies **"ex post perfect recall"**[31]:

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[32]:

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:

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

[31] See 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. Otherwise, one should use the statement of this paper.

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

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, $f_i = \mathrm{Id}_Z$. We call this hypothesis **perfect feedback**.

To define our main equilibrium concept, it is convenient to introduce the **pushforward** distribution of messages $\hat{F}_i(s_i, \sigma_{-i}) \in \Delta(M)$ induced by strategy s_i and the profile of strategy distributions σ_{-i}. Specifically:

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

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

1. (*rationality*) $\bar{s}_i$ is $(\mu_{\bar{s}_i}, \phi_i)$-unimprovable,
2. (*confirmed beliefs*) $\mu_{\bar{s}_i}\left(\left\{\sigma_{-i} \in \Sigma_{-i} : \hat{F}_i(\bar{s}_i, \sigma_{-i}) = \hat{F}_i(\bar{s}_i, \bar{\sigma}_{-i})\right\}\right) = 1.$[33]

An SCE $(\bar{\sigma}_i)_{i \in I}$ 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.

When profiles $(\bar{\sigma}_i)_{i \in I}$ and $\left(\mu_{\bar{s}_i}\right)_{i \in I, \bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i}$ satisfy the foregoing SCE conditions, we say that $(\bar{\sigma}_i)_{i \in I}$ is **justified by confirmed beliefs** $\left(\mu_{\bar{s}_i}\right)_{i \in I, \bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i}$. The set of self-confirming equilibria of (Γ, f, ϕ) is denoted by $SCE(\Gamma, f, \phi)$, while $symSCE(\Gamma, f, \phi)$ denotes the set of symmetric SCEs.

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, that is, $\hat{F}_i(\bar{s}_i, \bar{\sigma}_{-i})$.[34] Every profile of distributions that yields the same distribution of observations is empirically indistinguishable from the true one, and hence it cannot be objectively rejected.

The extensive form Γ is sufficient to define traditional solution concepts that do not try to capture steady states and implicitly assume that either ambiguity is absent or players are ambiguity neutral. With this, we first observe that every sequential equilibrium of Γ corresponds to an SCE of any extension of Γ with feedback and ambiguity attitudes:

Proposition 3. *For every (Γ, f, ϕ), every sequential equilibrium in behavioral strategies $\left(\bar{\beta}_i\right)_{i \in I}$ of Γ corresponds to a realization-equivalent SCE $(\bar{\sigma}_i)_{i \in I}$ of (Γ, f, ϕ).*

Intuitively, for every sequential equilibrium in behavioral strategies $\left(\bar{\beta}_i\right)_{i \in I}$ of Γ, we can consider the corresponding realization-equivalent mixed strategy profile $(\bar{\sigma}_i)_{i \in I}$. Let $\mu_{\bar{s}_i} = \delta_{\bar{\sigma}_{-i}}$ for every i and $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$; since no ambiguity is perceived, agents with these beliefs behave as expected utility maximizers. With this, it can be shown that each $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$ is $(\mu_{\bar{s}_i}, \phi_i)$-unimprovable, because $(\bar{\sigma}_i)_{i \in I}$ corresponds to a sequential equilibrium of Γ.

Thus, we obtain existence of SCE from well known results of existence for the sequential equilibrium concept (see, e.g., Theorem 8.5 in Fudenberg and Tirole, 1991):

Corollary 4. *For every (Γ, f, ϕ), the set $SCE(\Gamma, f, \phi)$ is not empty. Furthermore, if Γ has perfect information also $symSCE(\Gamma, f, \phi)$ is not empty.*

The second statement in Corollary 4 follows from Proposition 3 and the well known fact that every finite game with perfect information has a pure strategy subgame perfect equilibrium, which is also 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:

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

33 Recall that $\bar{\sigma}_{-i}$ includes $\bar{\sigma}_0$.

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

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). Therefore, the set of self-confirming equilibria *à la* BCMM of a sequential game (Γ, f, ϕ) is $SCE(\mathcal{N}(\Gamma, f), \phi)$. Here, instead, we require agents to maximize the KMM value over actions at every information set they deem reachable.

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

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

If agents instead are ambiguity averse, the SCEs of a game are not realization-equivalent to SCEs of its strategic-form representation. Indeed, the two sets of SCE outcomes may be non-nested.

Example 4 *(Running example, strategic form).* 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, for instance, (Out.M, L) 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 justified by any confirmed belief for any $\bar{\phi}_2$; the latter condition holds vacuously because player 2 is not reached. However, note that inequality (3) implies that (Out.M, L) does not belong to $SCE\left(\mathcal{N}(\Gamma, f), \bar{\phi}\right)$. Specifically, (3) implies that, for every belief μ_1 and action $a_1 \in \{\mathrm{T}, \mathrm{M}, \mathrm{B}\}$, strategy Out.$a_1$ is not ex-ante optimal, and thus it does not satisfy the best reply condition of BCMM. Hence, for every feedback f, Out is an SCE outcome of (Γ, f), but it is not an SCE outcome of $\mathcal{N}(\Gamma, f)$. Now suppose that players just observe their monetary payoff (that is, $f = u$). Then it can be checked that, for every α_2, (In.T, α_2) is an SCE of $\mathcal{N}(\Gamma, f)$ supported by the same belief μ_1 as before,[35] because player 1 prefers to commit to T and this prevents him from observing the long-run frequencies of L and R. Yet, In.T is not an SCE strategy of (Γ, f) because—as argued in Example 3—player 1 with prior μ_1 would deviate from T in the subgame despite the fact that μ_1 is the belief that minimizes the maximum value of deviating from T.[36] ▲

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 $\zeta : S \to Z$ and $\gamma : Z \to C$), and his feedback function $f_i : Z \to M$. Therefore, in an 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.[37]

6. Monotonicity of selfconfirming equilibrium

In this section we analyze changes in the set of equilibria when ambiguity attitudes are modified with respect to some baseline $\bar{\phi}$. In particular, we would like to prove an extension for sequential games of the following *monotonicity theorem* of BCMM: under observable payoffs, the SCE correspondence is monotone with respect to ambiguity aversion, that is, as ambiguity aversion increases, the set of SCEs expands.[38] Formally:

Definition 6. We 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, and 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

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

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

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

[35] The belief of player 2 is immaterial.

[36] See Lemma 8, or Lemma 6 in BCMM.

[37] In the working paper, we analyze a notion of rationalizable SCE that is appropriate when there is common knowledge of (Γ, f, ϕ). See Section 7.

[38] In their working paper BCMM also show that when feedback is less informative, the set of SCE's is larger. This intuitive monotonicity result holds for our SCE concept as well.

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.[39]

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

This result is a corollary of Lemma 5 proved in Appendix A. We provide a sketch of proof, because it helps to understand how dynamic (in)consistency matters for SCE analysis. Fix $(\bar{\sigma}_i)_{i\in I} \in SCE(\Gamma, f, \bar{\phi})$ and consider a strategy $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$ 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 $\left(\mu_{\bar{s}_i}, \phi_i\right)$-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 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, $h \notin H_i\left(\mu_{\bar{s}_i}\right)$) and unimprovability does not impose any optimality requirement on $\bar{s}_{i,h}$ (see Definition 1). Hence, one can find a $\left(\mu_{\bar{s}_i}, \phi_i\right)$-unimprovable and realization-equivalent strategy $s_i \in [\bar{s}_i]$. The following example illustrates this intuition.

Example 5 *(Sequential optimality).* Let Γ' be the our game show, but suppose that Player 1 has already won an amount worth 5 utils instead of 4 that he gets if he opts Out immediately. 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:

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

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:

$$\phi_1\left(\frac{1}{2}\right) = \bar{\phi}_1\left(\frac{1}{2}\right) = -1,$$
$$\phi_1(u) = \bar{\phi}_1(u) = \sqrt{u} \quad \text{if } u \in [1, 9],$$
$$\bar{\phi}_1(36) = 6 > k = \phi_1(36) > \phi_1(9) = 3.$$

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 $\{(\mathrm{In}, \mathrm{G})\}$. This implies that Out.$a_1^{\phi}$ is (μ_1, ϕ_1)-unimprovable and $\left(\mathrm{Out}.a_1^{\phi}, \mathrm{L}\right)$ is a symmetric SCE of (Γ', f, ϕ) realization equivalent to (Out.M, L). ▲

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 $\left(\Gamma, \bar{\phi}, f\right)$ 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

[39] See Definition 2 and Remark 1.

that, given his belief and some action a_i' at h', a_i has a higher value than the SCE expected payoff. Yet, since 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 *(Running example, no monotonicity).* Let $\bar{\phi}_1$ and ϕ_1 be the ambiguity attitudes described in Examples 4 and 5. In Appendix A.2 we show that, if ϕ_1 (36) is sufficiently small, Out does not belong to the set of SCE outcomes of (Γ, f, ϕ), although it is an SCE outcome of $(\Gamma, f, \bar{\phi})$. 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 $(\mu_1(\cdot|\{(\text{In}, \text{G})\}) = \frac{1}{2}\delta_{\text{L}} + \frac{1}{2}\delta_{\text{R}}$, cf. Lemma 8). At the root, he anticipates this and, scared by the implied ex-ante objective expected reward of 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. ▲

Thus, the monotonicity result of BCMM does not extend to all sequential games, even if we restrict our attention to distributions of outcomes.[40] 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** 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. For example, games with simultaneous moves, leader-follower games and signaling games have this property. In this class of games, at any information set $h \in H_i$, the agent does not move again after h. Then, the value of an action at $h \in H_i$ does not depend on i's strategy, thus unimprovability coincides with sequential optimality.[41] Moreover, no information set of i is prevented by any strategy of i, thus strategies (or strategy distributions) are realization-equivalent if and only if they coincide. Hence, Lemma 1 implies the following result:

Corollary 5. *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*

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

Since in games with simultaneous moves no player moves more than once, Corollary 5 yields the monotonicity result of BCMM as a special case.

6.2. Symmetric self-confirming equilibria

Consider now a *symmetric* SCE $\bar{s}$ in a game with observable payoffs and *without chance moves*. We show that with higher ambiguity aversion there is a symmetric SCE equilibrium s with the same outcome.[42]

Theorem 6. *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:*

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

The intuition for this result is as follows. The symmetric 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

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

[41] Indeed, Proposition 2 can be extended to any ϕ_i in one-move games. Yet, Proposition 1 cannot; the reason is that, if ϕ_i is not linear and i is not a first mover, sequential optimality does not imply ex ante value maximization. Indeed, if sequential optimality was defined to include ex ante maximization, Lemma 1 would not be applicable to one-move games.

[42] Note that we can identify $symSCE\left(\Gamma, f, \bar{\phi}\right)$ with a subset of S. Hence, it makes sense to write $\zeta\left(symSCE\left(\Gamma, f, \bar{\phi}\right)\right)$.

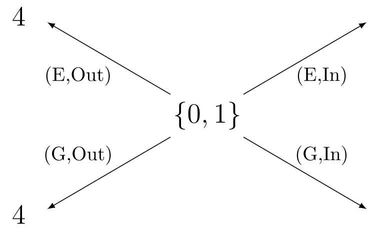

1 \ 2	L	R
T	9	9
M	1	36
B	36	1

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

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 8 in Appendix A and it is based on the following intuition: For any belief μ_i, the corresponding belief supported by Dirac models $\hat{\mu}_i$ with the same predictive as μ_i is a "mean-preserving spread" of μ_i, and the concavity of ϕ_i implies that player i feels worse-off with such spread. Second, absent chance moves, beliefs supported by Dirac models cannot entail dynamic inconsistencies of preferences over strategies, because the payoff under each model is deterministic, thus it is not influenced by a partial resolution of the uncertainty. This is shown by Lemma 7 in Appendix A. Therefore, the symmetric SCEs that these beliefs justify are in sequentially optimal strategies. 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 aversion. As for the role of symmetry (pure equilibrium), consider the following example.

Example 7 *(No chance).* Let Γ' be the following modification of the game of Fig. 1, represented in Fig. 2: 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). Assume perfect feedback and consider the same $\bar{\phi}_1$ and ϕ_1 of Example 6. The SCE $(\sigma_0, \text{Out.M}, \text{L})$ of $(\Gamma', f, \bar{\phi})$, where $\sigma_0(\text{E}) = \sigma_0(\text{G}) = \frac{1}{2}$, yields outcomes (E, Out) and (G, Out) with probability $1/2$ (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 δ_{σ_0}. Thus, no SCE of (Γ', f, ϕ) yields outcomes (E, Out) and (G, Out) with probability $1/2$. ▲

6.3. Ambiguity aversion versus ambiguity neutrality

Fix a sequential game with feedback (Γ, f) and let $\text{Id}_{\mathbb{V}} := \left(\text{Id}_{\mathbb{V}_i}\right)_{i \in I}$ denote the profile of players' identity functions characterizing their neutrality toward ambiguity. We relate the set $SCE(\Gamma, f, \text{Id}_{\mathbb{V}})$ of SCEs of Γ given feedback functions f and neutral ambiguity attitudes with the set $SCE(\Gamma, f, \phi)$ with non-neutral ambiguity attitudes ϕ. We start with a preliminary observation[43]:

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

To see this, note that ambiguity neutrality and the convexity of $\Sigma_{-i} = \Delta(S_{-i})$ in two-person games allow to replace the justifying confirmed beliefs supporting $(\bar{\sigma}_i)_{i \in I}$ 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, $(\bar{\sigma}_i)_{i \in I}$ 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 $\mu_{\bar{s}_i}$ 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. 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

43 Cf. footnote 23 of BCMM.

Proposition 2), the result about the comparison between ambiguity neutrality and ambiguity aversion is a corollary of Lemma 1.

Corollary 7. *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.*

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.

7. Discussion

We discuss some extensions and possible modifications of the SCE concept.

7.1. Heterogeneous populations

Our analysis can be easily extended to the case of populations with heterogeneous personal traits affecting preferences, or feedback. Let $\theta_i \in \Theta_i$ be the possible, privately known, personal traits of agents in population i. Tastes and risk attitudes (v_{θ_i}), ambiguity attitudes (ϕ_{θ_i}), and feedback (f_{θ_i}) are affected by θ_i. 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_i. For every profile of distributions $\left(\sigma_{\theta_j}\right)_{\theta_j \in \Theta_j} \in \Delta\left(S_j\right)^{\Theta_j}$ and every $s_j \in S_j$ the aggregate fraction of agents playing s_j is $\sum_{\theta_j \in \Theta_j} \tau_j\left(\theta_j\right)\sigma_{\theta_j}\left(s_j\right)$. With this, we can extend SCE to take into account that unimprovability and belief confirmation depend on personal traits (cf. Ch. 7 in Battigalli, 2018, and Dekel et al., 2004): A profile of distributions $\left(\left(\sigma_{\theta_i}\right)_{\theta_i \in \Theta_i}\right)_{i \in I}$ is a selfconfirming equilibrium if there is a profile of beliefs $\left(\left(\mu_{\theta_i, s_i}\right)_{\theta_i \in \Theta_i, s_i \in \mathrm{Supp}\sigma_{\theta_i}}\right)_{i \in I}$ such that, for all $i \in I$, $\theta_i \in \Theta_i$, and $s_i \in \mathrm{Supp}\sigma_{\theta_i}$

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

Our results extend seamlessly to this more general environment.

7.2. Incomplete information, ambiguous randomizations

We already explained that the SCE concept is relevant for incomplete-information environments with "private values," where players know their own utility (and feedback) function, but not those of co-players (see Section 5). We can extend the analysis to other forms of incomplete information. In particular, it is straightforward to relax the assumption that the probabilities of chance moves are (commonly) known: Chance can be analyzed as an indifferent player; 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. This notational trick also allows us to cover the case where players can delegate their choice of actions to random devices with unknown probabilities of outcomes ("Ellsberg urns"), as in Riedel and Sass (2014, 2017).

More generally, we can parametrize all the features of the game that are not commonly known with a parameter profile $\theta = \left(\theta_j\right)_{j \in I_0} \in \Theta$, where θ_0 parameterizes unknown chance probabilities and other non personal aspects of the game such as some features 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 (cf. Ch. 7 in Battigalli, 2018; Battigalli and Guaitoli, 1988). Also in this case, our results extend seamlessly to this more general environment.[44]

7.3. Rationalizable selfconfirming equilibrium

The SCE concept does not capture strategic reasoning: In an SCE, agents' beliefs may allow for irrational, or at least "strategically unsophisticated" behavior of other agents. Indeed, as we explained in Section 5, the SCE concept makes sense even if agents do not know the preferences and the feedback functions of other agents. In the working paper version we refine our notion of SCE so as to capture a form of strategic reasoning under complete information. Specifically, we put forward a relatively simple definition of **rationalizable SCE** that can be derived from the epistemic assumptions of (a)

[44] Of course, Theorem 6 holds vacuously if the games has chance moves.

subjective rationality, (b) belief confirmation, and (c) common initial belief of (a) and (b).[45] In the special case of ambiguity neutrality, our refinement is similar to the "partition-confirmed rationalizable SCE" of Fudenberg and Kamada (2015, 2018). We prove that analogs of the monotonicity results of this paper also hold for rationalizable SCE.

7.4. Monotonicity and restrictions on beliefs

In some applications, it is natural to suppose that the players have some structural knowledge, or common belief about the possible distributions of strategies in the populations of opponents, e.g., that agents use rationalizable strategies (see Section 7.3). More generally, we may add to the parameters describing the game a profile of **restricted belief sets** $R = (R_i)_{i\in I}$ with $R_i \subseteq \Delta(\Sigma_{-i})$ for each $i \in I$, and consider the equilibrium set $SCE(\Gamma, f, \phi, R)$ obtained by requiring that each player i's belief belongs to his restricted set. By inspection of the definition of SCE, one can check that stronger restrictions yield a smaller set of equilibria, that is, $\times_{i\in I}\bar{R}_i \subseteq \times_{i\in I}R_i$ implies $SCE\left(\Gamma, f, \phi, \bar{R}\right) \subseteq SCE(\Gamma, f, \phi, R)$. Hence, if restrictions on beliefs are due to restrictions on the possible probability models and there is more ambiguity (i.e., larger sets of possible probability models), the SCE set is larger.

Another natural question is whether our monotonicity results about increasing aversion to ambiguity extend to the case where the beliefs supporting any SCE are required to belong to the restricted sets. More formally, when is it the case that, if ambiguity aversion increases from a baseline level (represented by) $\bar{\phi}$ to a higher level ϕ, then $SCE\left(\Gamma, f, \bar{\phi}, R\right) \subseteq SCE(\Gamma, f, \phi, R)$? An inspection of the statement of Lemma 1 shows that the results for one-move games and the comparison between ambiguity neutrality and ambiguity aversion still hold when belief sets are restricted. In a one-move game or under ambiguity neutrality, given an SCE $(\bar{\sigma}_i)_{i\in I}$ justified by the confirmed beliefs $\left(\mu_{\bar{s}_i}\right)_{i\in I,\bar{s}_i\in\mathrm{Supp}\bar{\sigma}_i}$ with $\mu_{\bar{s}_i} \in R_i$, every $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$ is $(\mu_{\bar{s}_i}, \bar{\phi}_i)$-sequentially optimal.[46] Therefore, by Lemma 1, under increased ambiguity aversion there is a realization-equivalent SCE justified by the *same* confirmed beliefs, hence beliefs that continue to satisfy the restrictions.

As for symmetric SCEs, by inspection of the proof of Theorem 6 one can show that a sufficient condition for the result to hold is that each restricted belief set R_i has the following property: for every belief μ_i in R_i also the corresponding belief

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

belongs to R_i, where $\delta_{\delta_{s_{-i}}} \in \Delta(\Sigma_{-i})$ is the Dirac belief that assigns probability one to the Dirac model $\delta_{s_{-i}} \in \Sigma_{-i}$ (that is, player i is certain that all agents in population j play strategy s_j). For example, suppose we analyze a repeated game with impatient players, rather than a recurrent population game. Impatience implies that players only care about their one-period payoffs. With this, the definition of pure SCE is the relevant one to characterize steady states. Suppose also that the one-period game has no chance moves, which excludes the possibility to commit to randomizations and fits the hypothesis of Theorem 6. It is then natural to assume that each player believes that the co-players play pure strategies. In this case, R_i is the set of beliefs supported by Dirac models (isomorphic to $\Delta(S_{-i})$), the foregoing sufficient condition is obviously satisfied, and Theorem 6 holds.[47] This also clarifies that we can reinterpret this theorem as a result about repeated games played by impatient players.

7.5. Infinite games

As most of the literature on SCE in sequential games, we analyze finite games. In this discussion, without much loss of generality, we restrict our attention to games with a multistage structure and we consider infinite games. It is quite straightforward to extend the formal analysis to games with finitely many stages where players have finitely many feasible actions at nonterminal information sets, but may have a continuum of feasible actions in the last stage.[48] Yet, our analysis of ambiguity aversion and our notion of SCE rely on the informal assumptions that the sequential play occurs in *one period*, consequences realize at the end, and the game is played—with changing, randomly drawn opponents—for infinitely many periods (see Section 2). The possibility of infinite duration, instead, requires a different analysis for games with one or more stages in each period that may last for infinitely many periods, such as infinitely repeated games. It does not make much sense to assume that games with infinite duration are played infinitely many times with changing opponents. We should instead consider a notion of SCE as a limit that can be approached while playing the same game, similar to the "subjective equilibrium" of Kalai and Lehrer (1995). The key difference concerning the belief-confirmation property is that in this context beliefs concern multi-period strategies, e.g., grim strategies in the infinitely-repeated Prisoners Dilemma. As for preferences, one should consider streams of consequences that realize at the end of each period, and a smooth ambiguity

---

[45] In the working paper we offer a preliminary discussion of the learning foundations of rationalizable SCE and we relate them to alternative epistemic assumptions.

[46] Clearly, $\bar{\phi}_i$ is linear in the case of ambiguity neutrality.

[47] Another example of a restriction on beliefs that satisfies the aforementioned condition is the one imposed by Peer-Confirming Equilibrium, a network-based refinement of RSCE (see Lipnowski and Sadler, 2018).

[48] A continuum of actions in earlier stages instead presents technical difficulties in the analysis of conditional beliefs. For related reasons, there is no agreed-upon definition of sequential equilibrium for such games. See Myerson and Reny (2018).

criterion that takes this into account, possibly allowing for a form of discounting (Battigalli et al. 2016c take some steps in this direction). Our informed conjecture is that appropriate versions of Lemma 1, Theorem 6, and Corollary 7 can be proved for this different framework.

Appendix A

To ease notation, throughout the appendix we let σ denote a profile of distributions $(\sigma_i)_{i\in I}$, whereas σ_{-i} still denotes a product distribution in Σ_{-i} (thus including $\bar{\sigma}_0$) as in the main text.

A.1. Proof of Proposition 3

Fix any sequential equilibrium in behavioral strategies $\bar{\beta} = \left(\bar{\beta}_i\right)_{i\in I_0}$ of game Γ (chance has been included for notational convenience). For each $i \in I_0$, let $\bar{\sigma}_i$ denote the mixed strategy associated with $\bar{\beta}_i$ according to Kuhn's (1953) transformation:

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

Let $\mu_{\bar{s}_i} = \delta_{\bar{\sigma}_{-i}}$ for each $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$. By construction, these beliefs are correct, hence confirmed. To show that $\bar{\sigma} = (\bar{\sigma}_j)_{j\in I}$ is an SCE justified by these confirmed beliefs, we must prove that, for each $i \in I$, each $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$ 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 $\prod_{h'\in H_i} \bar{\beta}_i\left(\bar{s}_{i,h'}|h'\right) > 0$, that is, for every $\bar{s}_i \in \mathrm{Supp}\bar{\sigma}_i$,

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

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

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

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

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

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

A.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

$$\phi_1\left(\frac{1}{2}\right) = \bar{\phi}_1\left(\frac{1}{2}\right) = -1$$

$$\phi_1\left(u\right) = \bar{\phi}_1(u) = \sqrt{u} \quad \text{if } u \in [1, 9],$$

$$\bar{\phi}_1(36) = 6 > \phi_1\left(36\right) > \phi_1\left(9\right) = 3$$

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

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

$$\mathrm{e_M}(\sigma_2) = 1 \cdot \sigma_2\left(\mathrm{L}\right) + 36 \cdot \sigma_2\left(\mathrm{R}\right)$$

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

$$G_\nu(\sigma_2) := \nu\left(\left\{\sigma_2' \in \Delta(\{\mathrm{L}, \mathrm{R}\}) : \mathrm{e_M}(\sigma_2') \le \mathrm{e_M}(\sigma_2)\right\}\right).$$

Player 1 prefers M to T only if

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

Let $\nu = \mu_1(\cdot|\{(\text{In}, \text{G})\})$. Then, for every $\sigma_2 \in \Delta(\{\text{L}, \text{R}\})$,[49]

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

that is,

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

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

$$G_\nu(\sigma_2) \cdot \phi_1(\text{e}_\text{M}(\sigma_2)) + (1 - G_\nu(\sigma_2)) \cdot k \geq 3.$$

Let $\hat{\sigma}_2$ be the solution of $\text{e}_\text{M}(\sigma_2) = 17/2$, then

$$G_\nu(\hat{\sigma}_2) \leq \frac{k-3}{k - \sqrt{17/2}}. \tag{A.1}$$

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

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

because the belief about the co-player cannot change upon observing one's own move, or a chance move. DM_i 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 $\{\varnothing\}$.
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 (A.1) 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 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 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[50]:

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

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

A.3. Monotonicity of the SCE correspondence

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

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

[49] 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 $(1 - G_\nu(\sigma_2))$.

[50] The 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 $\frac{k-3}{k-\sqrt{17/2}}$.

Lemma 2. *Ex post perfect recall implies that, for all $i \in I$, $h \in H_i$, and $s_i \in S_i(h)$, the collection of subsets $\left\{F_{i,s_i}^{-1}(m) : m \in M_i(h)\right\}$ is a partition of $S_{-i}(h)$.*

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

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

$$S_{I_0}(h) = S_i(h) \times S_{-i}(h),$$

and that $m' \neq m''$ implies $F_{i,s_i}^{-1}(m') \cap F_{i,s_i}^{-1}(m'') = \emptyset$. We must show that

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

We first prove that

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

Fix any $s_{-i} \in S_{-i}(h)$; since $s_i \in S_i(h)$ and $S_{I_0}(h) = S_i(h) \times S_{-i}(h)$, then $(s_i, s_{-i}) \in S_{I_0}(h)$, that is, $x \prec \zeta(s_i, s_{-i})$ for some $x \in h$. Thus, by definition of $M_i(h)$, $f_i(\zeta(s_i, s_{-i})) \in M_i(h)$. Hence, $s_{-i} \in F_{i,s_i}^{-1}(m)$ for some $m \in M_i(h)$.

Next we prove by contraposition that the converse

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

is implied by *ex post* perfect recall. Suppose that we can find some $m \in M_i(h)$ and $s'_{-i} \in F_{i,s_i}^{-1}(m) \setminus S_{-i}(h)$. 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 $s_{-i} \in \mathrm{proj}_{S_{-i}} \zeta^{-1}(z)$, so that $(s_i, s_{-i}) \in S_i(h) \times S_{-i}(h) = S_{I_0}(h)$ for some $s_i \in S_i(h)$. Let $z = \zeta(s_i, s_{-i})$ and $z' = \zeta\left(s_i, s'_{-i}\right)$. Then, by choice of s_{-i} and s'_{-i}, $f_i(z) = m = f_i\left(z'\right)$, 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 $f_i(z) = m = f_i\left(z'\right)$, thus violating ex post perfect recall. ■

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

Lemma 3. *Let $\bar{\sigma}$ be an SCE of the game with observable payoffs (Γ, f, ϕ) justified by the confirmed beliefs $\left(\mu_{s_i}\right)_{i \in I, s_i \in \mathrm{Supp}\bar{\sigma}_i}$. For every $i \in I$ and $s_i \in \mathrm{Supp}\bar{\sigma}_i$ 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 2,

$$S_{-i}(h) = \bigcup_{m \in M_i(h)} F_{i,s_i}^{-1}(m) \tag{A.2}$$

for each $h \in H_i$, where the sets $F_{i,s_i}^{-1}(m)$ ($m \in M_i(h)$) are disjoint. Let

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

denote 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,

$$\forall \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}),$$
$$\sigma_{-i}(S_{-i}(h)) \overset{((A.2),\text{obs.eq})}{=} \bar{\sigma}_{-i}(S_{-i}(h)) \overset{(\text{conf})}{=} p_{\mu_{s_i}}(S_{-i}(h)) > 0, \tag{A.3}$$

where the first equality follows from eq. (A.2) and the fact that σ_{-i} is observationally equivalent to $\bar{\sigma}_{-i}$, the second equality follows from the first and belief confirmation (conf), that is $\mu_{s_i}\left(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i})\right) = 1$, 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 $u_i(\zeta(s_i, s_{-i})) = u^m$ for all $s_{-i} \in S_{-i}$ with $f_i(\zeta(s_i, s_{-i})) = m$. Then, observable payoffs and eqs. (A.2)–(A.3) imply that

$$\forall \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i}), U_i(s_i, \sigma_{-i}|h) = U_i(s_i, \bar{\sigma}_{-i}|h). \tag{A.4}$$

Indeed, for all $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i})$,

$$\begin{aligned} U_i(s_i, \sigma_{-i}|h) &\overset{(\text{def}, (A.3))}{=} \sum_{s_{-i} \in S_{-i}(h)} \frac{\sigma_{-i}(s_{-i})}{\sigma_{-i}(S_{-i}(h))} u_i(\zeta(s_i, s_{-i}))) \\ &\overset{((A.2), \text{obs.p})}{=} \sum_{m \in M_i(h)} \frac{\sigma_{-i}(F_{i,s_i}^{-1}(m))u^m}{\sigma_{-i}(S_{-i}(h))} \overset{(\text{obs.eq}, (A.3))}{=} \sum_{m \in M_i(h)} \frac{\bar{\sigma}_{-i}(F_{i,s_i}^{-1}(m))u^m}{\bar{\sigma}_{-i}(S_{-i}(h))} \\ &\overset{((A.2), \text{obs.p})}{=} \sum_{s_{-i} \in S_{-i}(h)} \frac{\bar{\sigma}_{-i}(s_{-i})}{\bar{\sigma}_{-i}(S_{-i}(h))} u_i(\zeta(s_i, s_{-i}))) \overset{(\text{def})}{=} U_i(s_i, \bar{\sigma}_{-i}|h). \end{aligned}$$

Confirmed beliefs ($\mu_{s_i}\left(\hat{\Sigma}_{-i}(s_i, \bar{\sigma}_{-i})\right) = 1$), the updating formula (2), and eqs. (A.3)–(A.4) yield

$$\begin{aligned} V_i(s_{i,h}|h; s_i, \mu_{s_i}, \phi_i) &= \phi_i^{-1}\left(\int_{\Sigma_{-i}(h)} \phi_i\left(U_i(s_i, \sigma_{-i}|h)\right) \mu_{s_i}(\mathrm{d}\sigma_{-i}|h)\right) \\ &\overset{(2)}{=} \phi_i^{-1}\left(\int_{\Sigma_{-i}(h)} \phi_i\left(U_i(s_i, \sigma_{-i}|h)\right) \frac{\sigma_{-i}(S_{-i}(h))}{p_{\mu_{s_i}}(S_{-i}(h))} \mu_{s_i}(\mathrm{d}\sigma_{-i})\right) \\ &\overset{((A.4), \text{conf.})}{=} \phi_i^{-1}\left(\phi_i\left(U_i(s_i, \bar{\sigma}_{-i}|h)\right) \int_{\Sigma_{-i}(h)} \frac{\sigma_{-i}(S_{-i}(h))}{p_{\mu_{s_i}}(S_{-i}(h))} \mu_{s_i}(\mathrm{d}\sigma_{-i})\right) \\ &\overset{(A.3)}{=} \phi_i^{-1}\left(\phi_i\left(U_i(s_i, \bar{\sigma}_{-i}|h)\right)\right) = U_i(s_i, \bar{\sigma}_{-i}|h). \qquad \blacksquare \end{aligned}$$

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

Lemma 4. *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:

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

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

$$\begin{aligned} V_i(a_i|h; s_i, \mu_i, \bar{\phi}_i) &= \bar{\phi}_i^{-1}\left(\mathbb{E}_{\mu_i(\cdot|h)}\left[\bar{\phi}_i \circ U_{s_i,h,a_i}\right]\right), \\ V_i(a_i|h; s_i, \mu_i, \phi_i) &= \phi_i^{-1}\left(\mathbb{E}_{\mu_i(\cdot|h)}\left[\phi_i \circ U_{s_i,h,a_i}\right]\right). \end{aligned} \tag{A.5}$$

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

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

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

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

By eq. (A.5), $V_i(a_i|h; s_i, \mu_i, \bar{\phi}_i) \geq V_i(a_i|h; s_i, \mu_i, \phi_i)$. ■

The following result implies Lemma 1 of the main text. Furthermore, it is used to prove Theorem 6.

Lemma 5. *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 $\left(\mu_{s_i}\right)_{i\in I, s_i \in \mathrm{Supp}\bar{\sigma}_i}$. Suppose that for each $i \in I$, every $s_i \in \mathrm{Supp}\bar{\sigma}_i$ is $(\mu_{s_i}, \bar{\phi}_i)$-sequentially optimal. Then, there exist maps $\bar{\mathrm{s}}_i : \mathrm{Supp}\bar{\sigma}_i \to S_i$, $i \in I$, such that*

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

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

Proof. For every $i \in I$ and $s_i \in \mathrm{Supp}\bar{\sigma}_i$, we construct a (μ_{s_i}, ϕ_i)-unimprovable strategy $\bar{\mathrm{s}}_i(s_i)$. Map $\bar{\mathrm{s}}_i : \mathrm{Supp}\bar{\sigma}_i \to S_i$ is such that (1) $\bar{\mathrm{s}}_i(s_i)_h = s_{i,h}$ for all $h \in H_i(s_i)$, and (2) $\bar{\mathrm{s}}_i(s_i)$ is derived by folding back on $H_i\left(\mu_{s_i}\right) \backslash H_i(s_i)$ given μ_{s_i} and ϕ_i.[51] Therefore, by construction, for every $h \in H_i\left(\mu_{s_i}\right) \backslash H_i(s_i)$,

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

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

$$\begin{aligned} V_i(\bar{\mathrm{s}}_i(s_i)_h|h; \bar{\mathrm{s}}_i(s_i), \mu_{s_i}, \phi_i) &= V_i(s_{i,h}|h; s_i, \mu_{s_i}, \phi_i) \\ &\overset{\text{(Lemma 3)}}{=} U_i(s_i, \bar{\sigma}_{-i}|h) \\ &\overset{\text{(Lemma 3)}}{=} V_i(s_{i,h}|h; s_i, \mu_{s_i}, \bar{\phi}_i) \\ &\overset{\text{(s.opt)}}{\geq} V_i(a_i|h; \bar{\mathrm{s}}_i(s_i), \mu_{s_i}, \bar{\phi}_i) \\ &\overset{\text{(Lemma 4)}}{\geq} V_i(a_i|h; \bar{\mathrm{s}}_i(s_i), \mu_{s_i}, \phi_i), \end{aligned}$$

where the first equality follows from condition (i) in the statement of the lemma, which implies that $\bar{\mathrm{s}}_i(s_i)$ and s_i are realization equivalent, the second and the third equalities from Lemma 3, the first inequality from sequential optimality (s.opt), and the second inequality from Lemma 4. This shows that $\bar{\mathrm{s}}_i(s_i)$ is (μ_{s_i}, ϕ_i)-unimprovable.

To conclude, note that the profile $\left(\bar{\sigma}_i \circ \bar{\mathrm{s}}_i^{-1}, (\mu_{\hat{s}_i})_{\hat{s}_i \in \mathrm{Supp}(\bar{\sigma}_i \circ \bar{\mathrm{s}}_i^{-1})}\right)_{i\in I}$, where $\mu_{\hat{s}_i} = \mu_{s_i}$ for some s_i with $\hat{s}_i = \bar{\mathrm{s}}_i(s_i)$, also satisfies the confirmed beliefs condition, because profiles

$$\left(\bar{\mathrm{s}}_i(s_i), (\bar{\sigma}_0, (\bar{\sigma}_j \circ \bar{\mathrm{s}}_j^{-1})_{j\in I\backslash\{i\}})\right) \text{ and } (s_i, \bar{\sigma}_{-i})$$

induce the same distribution over terminal nodes. Therefore $(\bar{\sigma}_i \circ \bar{\mathrm{s}}_i^{-1})_{i\in I}$ is an SCE of (Γ, f, ϕ). ■

A.4. Symmetric equilibria

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

$$\zeta : S \to Z,$$

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

and

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

The proof of Theorem 6 requires some preliminary results. Since, for every $i \in I$, the augmented collection

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

satisfies ex post perfect recall, 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 $m \in f_i(Z)$,

[51] Note 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.

$$S\left(f_i^{-1}(m)\right) = F_i^{-1}(m) = \text{proj}_{S_i} F_i^{-1}(m) \times \text{proj}_{S_{-i}} F_i^{-1}(m)$$
$$= S_i\left(f_i^{-1}(m)\right) \times S_{-i}\left(f_i^{-1}(m)\right).$$

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:

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

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 $V_i\left(a_i^m | h; s_i, \mu_i, \phi_i\right) = u_i(z)$ for every $z \in h$, independently of (s_i, μ_i, ϕ_i).

Remark 5. By ex post perfect recall, for every $i \in I$, $h \in H_i$ and $a_i \in A_i(h)$, $\left\{S_{-i}(h')\right\}_{h' \in \hat{H}_i(h, a_i)}$ is a partition of $S_{-i}(h)$.

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

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

denote the strategies of i allowing for h and choosing a_i at h. Then, by definition of $\hat{H}_i(h, a_i)$, $S_i\left(h'\right) = S_i(h, a_i)$ 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,

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

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

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

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

Note that

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

because $h \prec h'$, hence $S_{-i}\left(h'\right) \subseteq S_{-i}(h)$, 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 $\left\{S_{-i}(h')\right\}_{h' \in \hat{H}_i(h, a_i)}$ is a partition. ■

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

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

Then

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

and

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

Therefore, if μ_i is supported by Dirac models,

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

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

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

and we can focus on the latter expected value. Under a belief supported by Dirac models, $\phi_i\left(V_i\left(s_{i,h}|h;s_i,\mu_i,\phi_i\right)\right)$ 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 6. *Fix any $i\in I$ and a belief supported by Dirac models $\hat{\mu}_i$. Then, for every $s_i\in S_i$ and $h\in H_i\left(\hat{\mu}_i\right)$,*

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

Proof. Fix $s_i\in S_i$ and $h\in H_i\left(\hat{\mu}_i\right)$ arbitrarily. By Remark 5 and taking into account that $\hat{\mu}_i$ is supported by Dirac models, we can write

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

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

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

Furthermore, for every $h'\in\hat{H}_i(h,s_{i,h})$ and $s_{-i}\in S_{-i}\left(h'\right)$, $U_i\left(s_{i|h},s_{-i}\right)=U_i\left(s_{i|h'},s_{-i}\right)$, because, by definition and by choice of h', replacements $s_{i|h}$ and $s_{i|h'}$ are the same strategy. Thus,

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

and we obtain the desired result. ■

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

Corollary 8. *Fix $i\in I$ and a belief $\hat{\mu}_i$ supported by Dirac models. Then, for every $s_i\in S_i$ and $h\in H_i\left(\hat{\mu}_i\right)$,*

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

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 7. *For every $i\in I$ and $\hat{\mu}_i\in\Delta\left(\Sigma_{-i}\right)$, if $\hat{\mu}_i$ is a belief supported by Dirac models, then the set of $\left(\hat{\mu}_i,\phi_i\right)$-unimprovable and $\left(\hat{\mu}_i,\phi_i\right)$-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:

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

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 $h \in H_i(\mu_i) = H_i\left(\hat{\mu}_i\right)$, 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 $\left(\hat{\mu}_i, \phi_i\right)$-unimprovable strategy. Formally:

Lemma 8. *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,*

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

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 ($h \in \left\{f_i^{-1}(m) : m \in f_i(Z)\right\}$). Furthermore, $\hat{H}_i(\mu_i) = \hat{H}_i\left(\hat{\mu}_i\right)$ 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

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

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

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

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

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

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

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

Next we prove that

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

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, $\hat{H}_i(h, a_i, \sigma_{-i}) = \hat{H}_i\left(h, a_i, \delta_{\sigma_{-i}}\right)$). By Remark 5, 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 6, and the Inductive Hypothesis,

$$\begin{aligned}
&\phi_i\left(V_i\left(s_{i,h}|h; s_i^*, \mu_i, \phi_i\right)\right) \\
&\overset{(R5,\text{def})}{=} \int_{\Sigma_{-i}(h)} \phi_i\left(\sum_{h' \in \hat{H}_i\left(h, s_{i,h}\right)} \sum_{s_{-i} \in S_{-i}(h')} \sigma_{-i}(s_{-i}|h)\, U_i\left(s_i^h, s_{-i}\right)\right) \mu_i(\mathrm{d}\sigma_{-i}|h) \\
&= \int_{\Sigma_{-i}(h)} \phi_i\left(\sum_{h' \in \hat{H}_i\left(h, s_{i,h}, \sigma_{-i}\right)} \sigma_{-i}\left(S_{-i}\left(h'\right)|h\right) \sum_{s_{-i} \in S_{-i}(h')} \frac{\sigma_{-i}(s_{-i}|h)}{\sigma_{-i}(S_{-i}(h')|h)} U_i\left(s_i^h, s_{-i}\right)\right) \mu_i(\mathrm{d}\sigma_{-i}|h)
\end{aligned}$$

$$\overset{\text{(Jen,ch)}}{\geq} \int_{\Sigma_{-i}(h)} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \sigma_{-i})} \sigma_{-i}\left(S_{-i}\left(h'\right)|h\right) \phi_i \left(\sum_{s_{-i} \in S_{-i}(h')} \sigma_{-i}\left(s_{-i}|h'\right) U_i\left(s_i^h, s_{-i}\right) \right) \mu_i\left(\mathrm{d}\sigma_{-i}|h\right)$$

$$\overset{\text{(def)}}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} \int_{\Sigma_{-i}(h')} \sigma_{-i}\left(S_{-i}\left(h'\right)|h\right) \phi_i\left(U_i\left(s^*_{i|h'}, \sigma_{-i}|h'\right)\right) \mu_i\left(\mathrm{d}\sigma_{-i}|h\right)$$

$$= \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}\left(S_{-i}\left(h'\right)|h\right) \int_{\Sigma_{-i}(h')} \frac{\sigma_{-i}\left(S_{-i}\left(h'\right)|h\right)}{p_{\mu_i}\left(S_{-i}\left(h'\right)|h\right)} \phi_i\left(U_i\left(s^*_{i|h'}, \sigma_{-i}|h'\right)\right) \mu_i\left(\mathrm{d}\sigma_{-i}|h\right)$$

$$\overset{\text{(Bayes)}}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}\left(S_{-i}\left(h'\right)|h\right) \int_{\Sigma_{-i}(h')} \phi_i\left(U_i\left(s^*_{i|h'}, \sigma_{-i}|h'\right)\right) \mu_i\left(\mathrm{d}\sigma_{-i}|h'\right)$$

$$\overset{\text{(def)}}{=} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}\left(S_{-i}\left(h'\right)|h\right) \phi_i\left(V_i(s^*_{i,h'}|h'; s^*_i, \mu_i, \phi_i)\right)$$

$$\overset{\text{(I.H.)}}{\geq} \sum_{h' \in \hat{H}_i(h, s_{i,h}, \mu_i)} p_{\mu_i}\left(S_{-i}\left(h'\right)|h\right) \phi_i\left(V_i(s_{i,h'}|h'; s_i, \hat{\mu}_i, \phi_i)\right)$$

$$\overset{\text{(Lemma 6)}}{=} \phi_i\left(V_i\left(s_{i,h}|h; s_i, \hat{\mu}_i, \phi_i\right)\right).$$

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

Proof. Proof of Theorem 6 Let $z^* \in \zeta\left(symSCE\left(\Gamma, f, \bar{\phi}\right)\right)$. Then, there is a symmetric SCE $\left(s^*_i\right)_{i \in I}$ of $(\Gamma, f, \bar{\phi})$ such that $z^* = \zeta(\left(s^*_i\right)_{i \in I})$. Let $\left(s^*_i\right)_{i \in I}$ 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 $\hat{\mu}_i \in \Delta(\Sigma_{-i})$ be the belief supported by Dirac models with the same predictive as μ_i, so that $H_i(\mu_i) = H_i\left(\hat{\mu}_i\right)$. Compute by folding back a $\left(\hat{\mu}_i, \bar{\phi}_i\right)$-unimprovable strategy $\hat{s}^*_i$; 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 $\hat{s}^*_i$ elsewhere, that is,

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

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

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

Therefore, each belief in profile $\left(\hat{\mu}_i\right)_{i \in I}$ is confirmed under both $\left(s^*_i\right)_{i \in I}$ and $\left(\hat{s}_i\right)_{i \in I}$. We are going to prove that each $\hat{s}_i$ is $\left(\hat{\mu}_i, \bar{\phi}_i\right)$-sequentially optimal, which implies that $\left(\hat{s}_i\right)_{i \in I}$ is an SCE of $(\Gamma, f, \bar{\phi})$ and that Lemma 5 applies. Therefore, there is an equivalent symmetric SCE $(s_i)_{i \in I}$ of the more ambiguity averse game (Γ, f, ϕ) such that $\zeta\left((s_i)_{i \in I}\right) = z^*$, as desired.

Sequential optimality of $\hat{s}_i$ under $\left(\hat{\mu}_i, \bar{\phi}_i\right)$. We show that $\hat{s}_i$ is $\left(\hat{\mu}_i, \bar{\phi}_i\right)$-*unimprovable.* Since $\hat{\mu}_i$ is supported by Dirac models, Lemma 7 implies that $\hat{s}_i$ is also $\left(\hat{\mu}_i, \bar{\phi}_i\right)$-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 $\left(\hat{s}_i, \hat{\mu}_i\right)$, that is,

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

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

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

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

and

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

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

$$\forall h \in H_i\left(\hat{\mu}_i\right) \backslash H_i\left(\hat{s}_i\right), \hat{s}_{i,h} \in \arg\max_{a_i \in A_i(h)} V_i\left(a_i|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i\right). \tag{A.6}$$

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 $\hat{\mu}_i$) and payoffs observability imply

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

By definition of $\hat{s}_i$,

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

By Lemma 8,

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

Hence:

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

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

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

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

$$\forall h \in H_i\left(\hat{\mu}_i\right) \cap H_i\left(\hat{s}_i\right), \hat{s}_{i,h} \in \arg\max_{a_i \in A_i(h)} V_i\left(a_i|h; \hat{s}_i, \hat{\mu}_i, \bar{\phi}_i\right). \tag{A.7}$$

Eqs. (A.6)–(A.7) imply that $\hat{s}_i$ is $\left(\hat{\mu}_i, \bar{\phi}_i\right)$-unimprovable, hence—by Lemma 7—sequentially optimal and the thesis follows from Lemma 5. ■

References

Battigalli, P., 1987. Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. Unpublished thesis. Università Bocconi.
Battigalli, P., 2018. Game Theory: Analysis of Strategic Thinking. Typescript. Bocconi University.
Battigalli, P., Guaitoli, D., 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).
Battigalli, P., Gilli, M., Molinari, M.C., 1992. Learning and convergence to equilibrium in repeated strategic interaction. Res. Econ. 46, 335–378.
Battigalli, P., Catonini, E., Lanzani, G., Marinacci, M., 2017. Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games. IGIER Working Paper 607. Bocconi University.
Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2015. Self-confirming equilibrium and model uncertainty. Am. Econ. Rev. 105, 646–677.
Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2016a. A note on comparative ambiguity aversion and justifiability. Econometrica 84, 1903–1916.
Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2016b. Analysis of information feedback and selfconfirming equilibrium. J. Math. Econ. 66, 40–51.
Battigalli, P., Francetich, A., Lanzani, G., Marinacci, M., 2016c. Learning and Self-Confirming Long Run Biases. IGIER Working Paper 588. Bocconi University.
Blume, L., Easley, D., 2006. If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets. Econometrica 74, 929–966.
Brandts, J., Charness, G., 2011. The strategy versus the direct-response method: a first survey of experimental comparisons. Exp. Econ. 14, 375–398.
Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2013. Ambiguity and robust statistics. J. Econ. Theory 148, 974–1049.
Cubitt, R., van de Kuilen, G., Mukerji, S., 2019. Discriminating between models of ambiguity attitude: a qualitative test. J. Eur. Econ. Assoc., forthcoming.
Dekel, E., Fudenberg, D., Levine, D., 2004. Learning to play Bayesian games. Games Econ. Behav. 46, 282–303.
Eichberger, J., Grant, S., Kelsey, D., 2017. Ambiguity and the centipede game: strategic uncertainty in multi-stage games. Mimeo.
Epstein, L.G., Schneider, M., 2007. Learning under ambiguity. Rev. Econ. Stud. 74, 1275–1303.
Fudenberg, D., Kamada, Y., 2015. Rationalizable partition-confirmed equilibrium. Theor. Econ. 10, 775–806.
Fudenberg, D., Kamada, Y., 2018. Rationalizable partition-confirmed equilibrium with heterogeneous beliefs. Games Econ. Behav. 109, 364–381.
Fudenberg, D., Levine, D.K., 1993. Self-confirming equilibrium. Econometrica 61, 523–545.
Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press, Cambridge MA.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior. J. Math. Econ. 18, 141–153.
Hanany, E., Klibanoff, P., 2009. Updating ambiguity averse preferences. B.E. J. Theor. Econ. 9.
Hanany, E., Klibanoff, P., Mukerji, S., 2018. Incomplete information games with ambiguity averse players. Mimeo.
Kalai, E., Lehrer, E., 1995. Subjective games and equilibria. Games Econ. Behav. 8, 123–163.
Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. Econometrica 73, 1849–1892.
Kuhn, H.W., 1953. Extensive games and the problem of information. In: Kuhn, H.W., Tucker, A.W. (Eds.), Contributions to the Theory of Games II. Princeton University Press, Princeton NJ, pp. 193–216.
Lipnowski, E., Sadler, E., 2018. Peer-confirming equilibrium. Mimeo.
Lo, C.K., 1999. Extensive form games with uncertainty averse players. Games Econ. Behav. 28, 256–270.
Marinacci, M., 2015. Model uncertainty. J. Eur. Econ. Assoc. 13, 998–1076.
Myerson, R., Reny, P., 2018. Perfect conditional ε-equilibria of multi-stage games with infinite sets of signals and actions. Mimeo.
Osborne, M., Rubinstein, A., 1994. A Course in Game Theory. MIT Press, Cambridge MA.
Riedel, F., Sass, L., 2014. Ellsberg games. Theory Decis. 76, 469–509.
Riedel, F., Sass, L., 2017. Kuhn's theorem for extensive form Ellsberg games. J. Math. Econ. 68, 26–41.
Schelling, T.C., 1960. The Strategy of Conflict. Harvard University Press, Cambridge MA.
Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica 57, 571–587.
Siniscalchi, M., 2011. Dynamic choice under ambiguity. Theor. Econ. 6, 379–421.
Trautmann, S.T., van de Kuilen, G., 2016. Ambiguity attitudes. In: Keren, G., Wu, G. (Eds.), The Wiley-Blackwell Handbook of Judgment and Decision Making.