Journal of Mathematical Economics 66 (2016) 40–51

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Analysis of information feedback and selfconfirming equilibrium

P. Battigalli, S. Cerreia-Vioglio, F. Maccheroni *, M. Marinacci

Università Bocconi, Italy
IGIER, Italy

ARTICLE INFO

Article history:
Received 6 October 2015
Received in revised form
21 May 2016
Accepted 15 July 2016
Available online 27 July 2016

Keywords:
Selfconfirming equilibrium
Conjectural equilibrium
Information feedback
Ambiguity aversion
Partially specified probabilities

ABSTRACT

Recent research emphasizes the importance of information feedback in situations of recurrent decisions and strategic interaction, showing how it affects the uncertainty that underlies selfconfirming equilibrium (e.g., Battigalli et al., 2015, Fudenberg and Kamada, 2015). Here, we discuss in detail several properties of this key feature of recurrent interaction and derive relationships. This allows us to elucidate different notions of selfconfirming equilibrium, showing how they are related to each other given the properties of information feedback. In particular, we focus on Maxmin selfconfirming equilibrium, which assumes extreme ambiguity aversion, and we compare it with the partially-specified-probabilities (PSP) equilibrium of Lehrer (2012). Assuming that players can implement any randomization, symmetric Maxmin selfconfirming equilibrium exists under either "observable payoffs," or "separable feedback." The latter assumption makes this equilibrium concept essentially equivalent to PSP-equilibrium. If observability of payoffs holds as well, then these equilibrium concepts collapse to mixed Nash equilibrium.

1. Introduction

In a selfconfirming equilibrium (SCE), agents best respond to confirmed, but possibly incorrect, beliefs. The notion of SCE captures the rest points of dynamics of strategies and beliefs in games played recurrently (see, e.g., Fudenberg and Levine, 1993b; Fudenberg and Kreps, 1995; and Gilli, 1999). Battigalli et al. (2015) (henceforth BCMM) define a notion of selfconfirming equilibrium whereby agents have non-neutral attitudes toward model uncertainty, or ambiguity.[1] The SCE concept of BCMM encompasses the traditional notions of conjectural equilibrium (Battigalli, 1987; Battigalli and Guaitoli, 1988) and selfconfirming equilibrium (Fudenberg and Levine, 1993a) as special cases, taking as given the specification of an ex post information structure, or information feedback. Specifically, the information feedback function describes what an agent can observe *ex post*, at the end of the stage game which is being played recurrently. The properties of information feedback determine the type of partial-identification problem faced by a player who has to infer the co-players' strategies from observed data. This, in turn, shapes the set of selfconfirming equilibria.

We define several properties of information feedback, we study their relationships, and we illustrate them through the analysis of equilibrium concepts. Specifically, we focus on Maxmin SCE, which assumes an extreme form of ambiguity aversion, and its relationships with other equilibrium concepts. We also deviate from BCMM by allowing players to delegate their choices to arbitrary randomization devices. Three properties of information feedback play a prominent role in our analysis: (i) "observable payoffs" means that each player observes his own realized utility, (ii) "own-strategy independence of feedback" means that inferences about the strategy profile played by the opponents do not depend on one's own adopted strategy; (iii) "separable feedback" is a strengthening of own-strategy independence: inferences about the strategy of each opponent are independent of how other agents play. While (i) is a natural property that holds in many applications, we argue that (ii)–(iii) are very strong properties of feedback. BCMM show that, if payoffs are observable, then the traditional ambiguity-neutral SCE concept is a refinement of Maxmin SCE; hence, ambiguity aversion (weakly) expands the equilibrium set. On the other hand, under observable payoffs and own-strategy independence of feedback every SCE concept is equivalent to mixed Nash equilibrium.

We show that all games with separable feedback have a "symmetric" Maxmin SCE in mixed strategies.[2] We observe that

* Correspondence to: Università Bocconi, Via Roentgen, 1, 20136, Milan (MI), Italy.
E-mail address: fabio.maccheroni@unibocconi.it (F. Maccheroni).

[1] For a discussion on the literature of choice under ambiguity, see the surveys of Gilboa and Marinacci (2013), and Marinacci (2015).

[2] "Symmetric" refers to the population-game scenario that we use to interpret the SCE concept: all the agents in the same player role use the same strategy.

http://dx.doi.org/10.1016/j.jmateco.2016.07.002
0304-4068/

games with separable feedback have a canonical representation as games with partially specified probabilities (PSP) in the sense of Lehrer (2012). Under this representation, our symmetric Maxmin SCE is equivalent to the equilibrium concept put forward by Lehrer. We also show that – under the (strong) assumption that all randomizations are feasible – Maxmin SCE is a refinement of ambiguity-neutral SCE. Our results imply that, in the canonical representation of a game with separable feedback, Lehrer's PSP equilibrium is a refinement of ambiguity-neutral SCE, and that under observability of payoffs it is equivalent to mixed Nash equilibrium.

The rest of the paper is organized as follows. Section 2 describes games with feedback and the partial identification correspondence; Section 3 analyzes the properties of information feedback and their consequences for the identification of co-players' behavior; Section 4 defines SCE concepts and relates them to each other and to Nash equilibrium; Section 4.3 analyzes existence of Maxmin SCE; Section 5 relates information feedback to partially specified probabilities; Section 6 further discusses our assumptions and results, and relates to the literature.

2. Games with feedback and partial identification

Throughout the analysis, we consider agents who play the strategic form of a finite game Γ in extensive form with perfect recall and no chance moves. The extensive-form structure shapes feedback and is relevant for the analysis of its properties. In this section, we introduce the key elements of games with feedback (2.1), we present the population-game backdrop of our analysis (2.2), and we define the identification correspondence (2.3).

2.1. Games with feedback

To define games with feedback, we start from a *finite* game in extensive form Γ. We defer the details of the extensive-form representation to Section 3, where we analyze the properties of feedback. Here we use only the following key primitive and derived elements of the game:

- I is the set of players *roles* in the game;
- Z is the finite set of *terminal nodes*;
- $u_i : Z \to \mathbb{R}$ is the *payoff* (vNM utility) *function* of player i;
- $S = \times_{i\in I} S_i$ is the finite set of *pure-strategy* profiles;
- $\zeta : S \to Z$ is the *outcome function*;
- $(I, (S_i, U_i)_{i\in I})$ is the *strategic form* of Γ, that is, for each $i \in I$ and $s \in S$, $U_i(s) = u_i(\zeta(s))$; as usual, U_i is multi-linearly extended to $\times_{j\in I} \Delta(S_j)$.

Following Battigalli (1987), we specify, for each player role $i \in I$, a *feedback function*

$$f_i : Z \to M,$$

where M is a *finite* set of "messages" representing what i can observe ex post about the path of play.[3] For instance, suppose that u_i is a monetary payoff function (or a strictly increasing function of the monetary payoff of i) and that i only observes ex post how much money he got, then $M \subseteq \mathbb{R}$ is a set of monetary outcomes and $f_i = u_i$. This example shows that, in our setup, the feedback function f_i does not necessarily reflect what a player remembers about the game just played; but we will introduce a property, called "ex post perfect recall", that requires just this. Another example is the feedback function assumed by Fudenberg and Levine (1993a): Each player i observes ex post the whole path of play. In this case, f_i is any injective function (e.g., the identity on Z).

A *game with feedback* is a tuple

$$(\Gamma, f) = (\Gamma, (f_i)_{i\in I}).$$

The *strategic-form feedback function* of i is $F_i = f_i \circ \zeta : S \to M$. This, in turn, yields the pushforward map $\hat{F}_i : \times_{j\in I} \Delta(S_j) \to \Delta(M)$ defined by

$$\hat{F}_i(\sigma)(m) = \sum_{s\in F_i^{-1}(m)} \prod_{j\in I} \sigma_j(s_j),$$

which gives the probability that i observes message m as determined by the mixed-strategy profile σ. We (informally) assume that each player i knows (1) the game tree and information structure (which determine S and ζ), (2) his feedback function f_i (hence his strategic-form feedback function F_i), and (3) his payoff function u_i. On the other hand, common knowledge of (Γ, f) is not relevant for our analysis, because SCE is not meant to capture inferences based on strategic reasoning.

2.2. Random matching and feasible strategies

We assume (informally) that the strategic form of game with feedback (Γ, f) is played recurrently by a large population of agents, partitioned according to the player roles $i \in I$ (male or female, buyer or seller, etc.). Agents drawn from different sub-populations are matched at random, play, get feedback according to f, and then are separated and re-matched to play again. The large-populations backdrop of our analysis is important to justify the assumption that, in steady state, non-myopic agents maximize their instantaneous expected payoff.[4] Furthermore, we assume that agents can covertly and credibly commit to play any mixed strategy.[5]

The exact details of the matching process are not important as long as the following condition is satisfied: If everyone keeps playing the same strategy, the co-players' strategy profile faced by each agent at each stage is an i.i.d. draw with probabilities given by the statistical distribution of strategies in the co-players' sub-populations. This is consistent with Nash's mass action interpretation of equilibrium (Weibull, 1996).

Our assumptions about (covert) commitment are instead important and restrictive. According to the main equilibrium concept formally defined below, Maxmin SCE, agents are ambiguity averse. Two issues arise. First, it is known that ambiguity-averse agents are dynamically inconsistent; therefore, they may be unwilling to implement contingent choices implied by some pure strategy they deem ex ante optimal (see the discussion in BCMM and the references therein). Second, even in a simultaneous-move game, an ambiguity-averse agent may not want to implement the realization of an ex ante optimal mixed strategy. Therefore, we are assuming that agents truly and irreversibly delegate their choices in Γ to some device implementing a pure or mixed strategy. We maintain such strong assumptions only for expositional simplicity. We want to focus on the properties of feedback and their consequences. The fact that players obtain feedback about the terminal node of an extensive-form game gives structure and makes the analysis more interesting. Taking dynamic incentive

3 See also Battigalli and Guaitoli (1988). Assuming a common finite set of messages is without loss of generality: let $M = \bigcup_{i\in I} f_i(Z)$.

4 Alternatively, with a fixed set of players, we should either assume perfect impatience, or look at equilibria of repeated games with imperfect monitoring, in the spirit of Kalai and Lehrer (1995). See the discussion in the survey by Battigalli et al. (1992).

5 Commitment is covert because it is not observed by other agents. Covert commitment is relevant when agents are dynamically inconsistent.

compatibility constraints into account would complicate the analysis, therefore, we do not address this issue here (see Battigalli et al., 2014, and the discussion in BCMM). As for mixed strategies, we could take as given a subset $\Sigma_i \subseteq \Delta(S_i)$ of implementable mixed strategies and spell out in each one of our formal results the assumptions about such sets Σ_i $(i \in I)$ that we need to prove the thesis; but this would complicate the exposition. We will instead point out where assuming the feasibility of all mixed strategies is crucial. These assumptions about commitment and randomization also make our analysis more directly comparable to some related papers. We further discuss the feasibility of randomizations in Section 6.

2.3. Identification correspondence

Building on Fudenberg and Levine (1993a,b) and BCMM, we analyze notions of equilibrium that capture stationary states of repeated anonymous interaction with learning from personal experience. In equilibrium, each mixed strategy played by a positive fraction of agents has to be a best response given the feedback obtained by those agents. In general, such feedback is partial and does not allow to identify the (average) mixed strategies of the other players, which allows ambiguity aversion to play a role.

The agents in each role-population $i \in I$ are distributed according to their adopted mixed strategies. Such distribution is a kind of "second-order mixed strategy" $\varsigma_i \in \Delta(\Delta(S_i))$ with predictive probabilities $\bar{\varsigma}_i(s_i) = \int_{\Delta(S_i)} \sigma_i(s_i)\varsigma_i(\mathrm{d}\sigma_i)$, $\bar{\varsigma}_i \in \Delta(S_i)$. The expected payoff of each agent depends on his own strategy and on the average mixed strategy of the agents in other populations. Yet, just looking at average behavior is insufficient, because different (mixed) strategies in the support of ς_i may be justified by different feedback (see BCMM). Thus, for the sake of simplicity, here we focus on the *symmetric* case where, for each role $i \in I$, each agent in population i plays the same mixed strategy.[6]

An agent who keeps playing mixed strategy σ_i^*, while the opponents play σ_{-i}^*, obtains a distribution of observations $\hat{F}_i(\sigma_i^*, \sigma_{-i}^*) \in \Delta(M)$. We (informally) assume that each agent observes the realization s_i of his mixed strategy σ_i^*. Therefore an agent playing σ_i^* "observes" the profile of conditional distributions of messages $(\hat{F}_i(s_i, \sigma_{-i}^*))_{s_i \in \mathrm{supp}\sigma_i^*}$ in the long run[7]; collecting these conditional distributions, such an agent infers that the opponents' mixed strategy profile belongs to the set

$$\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*) = \{\sigma_{-i} \in \times_{j\neq i} \Delta(S_j) : \forall s_i \in \mathrm{supp}\sigma_i^*, \hat{F}_i(s_i, \sigma_{-i}) = \hat{F}_i(s_i, \sigma_{-i}^*)\}.$$

We call $\hat{\Sigma}_{-i}(\cdot, \cdot)$ the (partial) *identification correspondence*. This correspondence is non-empty (because $\sigma_{-i}^* \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$) and compact-valued (in two-person games, it is also convex-valued). Identification is partial because, typically, the set $\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$ is not a singleton.

Lemma 1. *For each $i \in I$, $\sigma_i^* \in \Delta(S_i)$, and $\sigma_{-i}^* \in \times_{j\neq i} \Delta(S_j)$,*

$$\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*) \subseteq \{\sigma_{-i} \in \times_{j\neq i} \Delta(S_j) : \hat{F}_i(\sigma_i^*, \sigma_{-i}) = \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)\}. \quad (1)$$

The inclusion may be strict.

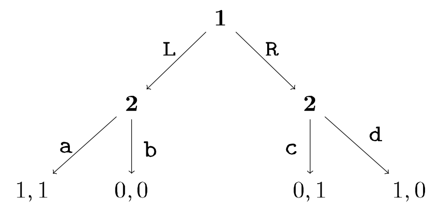

Fig. 1. A perfect information game.

Proof. The inclusion follows from equation $\hat{F}_i(\sigma_i^*, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i^*(s_i)\hat{F}_i(s_i, \sigma_{-i})$. To see that the inclusion may be strict, it is sufficient to exhibit an example. Consider the perfect information game of Fig. 1. Suppose that *the feedback of player* 1 *only reveals his own payoff*, that is, $f_1 = u_1$. Then

$$\hat{\Sigma}_{-1}\left(\frac{1}{2}L + \frac{1}{2}R, a.c\right) = \{a.c\},$$

and

$$\left\{\sigma_2 : \hat{F}_1\left(\frac{1}{2}L + \frac{1}{2}R, \sigma_2\right) = \hat{F}_1\left(\frac{1}{2}L + \frac{1}{2}R, a.c\right)\right\}$$
$$= \left\{\sigma_2 : \frac{1}{2}(\sigma_2(a.c) + \sigma_2(a.d)) + \frac{1}{2}(\sigma_2(a.d) + \sigma_2(b.d)) = \frac{1}{2}\right\}$$
$$= \{\sigma_2 : \sigma_2(a.c) + 2\sigma_2(a.d) + \sigma_2(b.d) = 1\}. \quad \blacksquare$$

In words, if in the foregoing example player 1 plays a totally mixed strategy, then each information set of player 2 is visited infinitely often. If player 1 takes into account the realizations of his mixed strategy, he infers from his own realized payoff the action taken by player 2 (a or b after L, c or d after R). Therefore, he finds out that 2 is playing $a.c$. But if player 1 does not take into account the realizations of his mixed strategy, then he only observes that he "wins" 50% of the times, which is consistent with many strategies of player 2, including pure strategy $b.d$ (which minimizes the payoff of 2) besides $a.c$.

Note that the r.h.s. of (1) would be the partially identified set under the alternative assumption that an agent only knows the mixed strategy σ_i^* he keeps playing, but does not observe the realizations of σ_i^* in the (infinitely) many plays of Γ. We show below in Proposition 1(b) that the inclusion in (1) can be strict only if the information feedback function f_i violates "ex post perfect recall," as in the example above. Intuitively, ex post perfect recall implies that the information about others given by ex post message m is equivalent to the information about others given by m and the pure-strategy realization s_i.

3. Extensive form and properties of feedback

We relate the properties of information feedback to some details of the extensive form of game Γ. The reader can refer to any standard definition of game in extensive form with perfect recall as a labeled tree with nonterminal nodes $x \in X$ and terminal nodes Z (see, e.g., Chapters 6 and 11 in Osborne and Rubinstein, 1994). Recall that $F_i = f_i \circ \zeta : S \to M$ denotes the strategic-form feedback function of i derived from $f_i : Z \to M$. We list in Table 1 the primitive and derived terms we need and the corresponding notation.

The last two lines of the table deserve further explanation. When i plays s_i, the message he receives is a function $m = F_{i,s_i}(s_{-i})$ of the co-players' strategies, the s_i-section of F_i. The collection of sets of pre-images $F_{i,s_i}^{-1}(m)$ of messages $m \in F_{i,s_i}(S_{-i})$ is a

6 This also facilitates the comparison with Lehrer (2012).

7 We denote by $\mathrm{supp}\sigma_i$ the support of measure σ_i. Whenever no confusion may arise, s_i is identified with δ_{s_i}, the Dirac measure supported by s_i.

Table 1
Additional notation.

Notation	Definition
$x \prec y$ $(x \preceq y)$	Node x precedes (weakly) node y
$H_i(x)$	Information set of i containing x
$a_i(x \to y)$	Action of i at x leading to y
$Z(h) = \{z \in Z : \exists x \in h, x \preceq z\}$	Terminal successors of nodes in set h
$S(h) = \{s \in S : \exists x \in h, x \preceq \zeta(s)\}$	Strategy profiles reaching h
$F_{i,s_i}(\cdot) = F_i(s_i, \cdot) : S_{-i} \to M$	s_i-section of F_i
$\mathcal{F}_{-i}(s_i) = \{C_{-i} \subseteq S_{-i} : \exists m \in M, F_{i,s_i}^{-1}(m) = C_{-i}\}$	Strategic feedback given s_i

partition $\mathcal{F}_{-i}(s_i)$ of S_{-i} that describes the feedback about co-players' strategies given i's own-strategy s_i.

We consider the following properties. A game with feedback (Γ, f) satisfies:

1. *perfect feedback* if, for every $i \in I$, f_i is one-to-one (injective);
2. *observable payoffs* if, for every $i \in I$, $u_i : Z \to \mathbb{R}$ is f_i-measurable; that is,[8]

   $$\forall z', \ z'' \in Z, \quad f_i(z') = f_i(z'') \Rightarrow u_i(z') = u_i(z'');$$
3. *own-strategy independence of feedback* if, for every $i \in I$, and $s_i, t_i \in S_i$, the sections F_{i,s_i} and F_{i,t_i} of F_i induce the same partition of pre-images on S_{-i}; that is, if $\mathcal{F}_{-i}(s_i) = \mathcal{F}_{-i}(t_i)$;
4. *ex post perfect recall* if, for every $i \in I$, and $z', z'' \in Z$, whenever there are decision nodes x', x'' of i such that $x' \prec z', x'' \prec z''$, and either $Z(H_i(x')) \cap Z(H_i(x'')) = \emptyset$ or $a_i(x' \to z') \neq a_i(x'' \to z'')$, then $f_i(z') \neq f_i(z'')$;
5. *ex post observable deviators* if, for every $i \in I$ and $m \in F_i(S)$,

   $$F_i^{-1}(m) = \times_{j \in I} \operatorname{proj}_{S_j} F_i^{-1}(m); \tag{2}$$
6. *separable feedback* if, for every $i \in I$, there are profiles of sets $\left(M_{i,j}\right)_{j \neq i}$ and onto functions $(F_{i,j} : S_j \to M_{i,j})_{j \neq i}$ such that, for each $s_i \in S_i$,

   $$\mathcal{F}_{-i}(s_i) = \{C_{-i} \subseteq S_{-i} : \exists (m_j)_{j \neq i} \in \times_{j \neq i} M_{i,j}, \quad C_{-i} = \times_{j \neq i} F_{i,j}^{-1}(m_j)\}. \tag{3}$$

In words, *perfect feedback* means that each player observes (ex post) the complete path of play. *Observable payoffs* weakens perfect feedback by just requiring that each player observes his realized vNM utility. This is a natural property that holds in many, if not most, applications. It does not hold, for example, when agents have other-regarding preferences and cannot perfectly observe *ex post* the consumption, or monetary payoff of other agents. BCMM point out an important consequence of the observability of payoffs: holding one's own strategy fixed, objective expected utility is constant over the partially identified set of opponents' mixed strategies.[9]

Lemma 2 (*Cf. BCMM, Lemma 1*)**.** *If (Γ, f) is a game with* observable payoffs*, then, $U_i(\sigma_i^*, \sigma_{-i}) = U_i(\sigma_i^*, \sigma_{-i}^*)$ for all $\left(\sigma_i^*\right)_{i \in I}$, $i \in I$, and $\sigma_{-i} \in \hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right)$.*

Own-strategy independence of feedback means that each player is an "information taker," that is, his ex post observations about his co-players' strategies are independent of the strategy he plays. For example, in a quantity-setting oligopoly with known demand schedule, even if a firm just observes the market price, it can infer the total output of the competitors from the observation of the price and the knowledge of its own output.

Remark 1. If own-strategy independence of feedback holds, the identification correspondence can be written as

$$\begin{aligned}\hat{\Sigma}_{-i}(\sigma_{-i}^*) &= \{\sigma_{-i} \in \times_{j \neq i} \Delta(S_j) : \exists s_i \in S_i, \hat{F}_i(s_i, \sigma_{-i}) = \hat{F}_i(s_i, \sigma_{-i}^*)\} \\ &= \{\sigma_{-i} \in \times_{j \neq i} \Delta(S_j) : \forall s_i \in S_i, \hat{F}_i(s_i, \sigma_{-i}) = \hat{F}_i(s_i, \sigma_{-i}^*)\}.\end{aligned}$$

Ex post perfect recall means that each player remembers at the end of the game the information he acquired while playing, and the actions he took. Note that players are assumed to remember during the play their previously acquired information and their own previous actions, because Γ is a game with perfect recall (see Kuhn, 1953). Therefore, it makes sense to assume that they remember also after the play. Ex post perfect recall requires that the feedback function f_i reflect this.

The *ex post observable deviators* property requires that each i obtains separate pieces of information about the strategy of each player j. Therefore, if i is "surprised" by a message m, he can observe who deviated from the set of paths $f_i^{-1}(m)$. The observable deviators property is defined for standard extensive-form information structures in independent work of Fudenberg and Levine (1993a) and Battigalli (1994). Using the definition of the latter, a game Γ has observable deviators if $S(h) = \times_{j \in I} \operatorname{proj}_{S_j} S(h)$ for each player i and information set h of i, where $S(h)$ is the set of strategy profiles reaching h.[10] Battigalli (1997) proves that this definition is equivalent (for games without chance moves) to the definition given by Fudenberg and Levine. Eq. (2) extends the observable deviators property to terminal information sets.

Finally, *separable feedback* means that each i obtains a separate signal about the strategy of every co-player j that depends only on what j does.

Own-strategy independence of feedback is a strong property. It can be shown that every game with this property and with perfect feedback can be transformed into a "realization-equivalent" simultaneous-moves game[11] by iteratively interchanging simultaneous moves, and coalescing sequential moves by the same player.[12] The intuition is that, if j moves after observing something of what i did, then there are information sets h' and h'' of j and a move of i that determines whether h' or h'' is reached; if strategy s_i makes h' (h'') reachable, then i cannot observe ex post what j would have chosen at h'' (h'). Therefore, if i's feedback about the co-players is not trivial, it must depend on i's strategy. We omit the formal statement and proof, which involve lengthy and tedious details. The following example illustrates.

Example 1. Consider the extensive form in Fig. 2 and assume *perfect feedback*. It can be verified that this non-simultaneous game satisfies own-strategy independence of feedback:

$$\mathcal{F}_{-1}(L.a) = \mathcal{F}_{-1}(L.b) = \mathcal{F}_{-1}(R.a) = \mathcal{F}_{-1}(R.b) = \{\{l\}, \{r\}\},$$
$$\mathcal{F}_{-2}(l) = \mathcal{F}_{-2}(r) = \{\{L.a\}, \{L.b\}, \{R.a, R.b\}\}.$$

---

8 Hence, $U_i : S \to \mathbb{R}$ is F_i-measurable.

9 In the Appendix, we provide proofs of results from BCMM for the sake of completeness, and also because BCMM consider versions of these results whereby agents can only choose pure strategies.

10 Battigalli (1994) uses this property in an analysis of structural consistency and stochastic independence for systems of conditional probabilities. Under observable deviators, structural consistency is weaker than stochastic independence; without observable deviators the two properties are unrelated.

11 We call two games Γ' and Γ'' "realization-equivalent" if there is an isomorphism between the quotients under the realization-equivalence relation of the strategy sets of Γ' and Γ'', and a corresponding isomorphism between the sets of terminal nodes of Γ' and Γ'', so that the map from strategy profiles and terminal nodes is preserved. Leonetti (2014) shows that two games are realization-equivalent if and only if one can be obtained from the other by a sequence of transformations of two types, interchanging simultaneous moves and coalescing/splitting moves by the same player.

12 In Example 3, we show a perfect-information game where the first mover's feedback about followers is completely trivial (thus violating perfect feedback) and own-strategy independence holds.

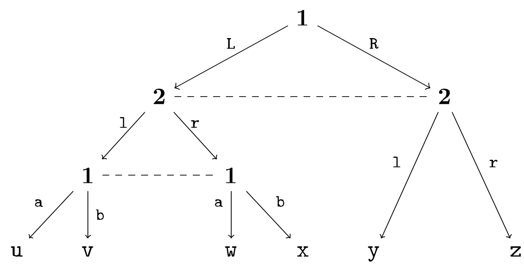

Fig. 2. An extensive form with perfect feedback.

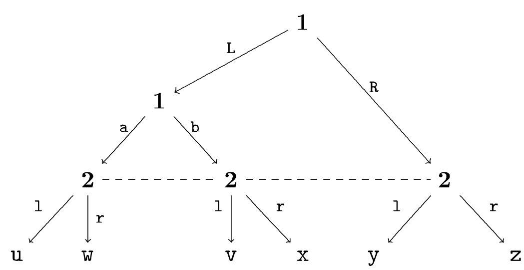

Fig. 3. A game obtained from Fig. 2.

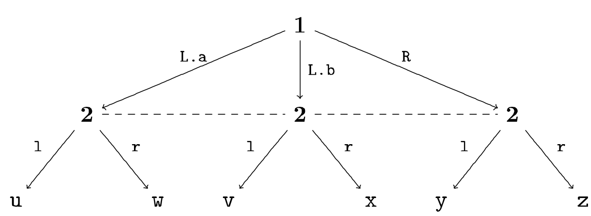

Fig. 4. A game obtained from Fig. 3.

The games with perfect feedback depicted in Figs. 3 and 4 are obtained from the game of Fig. 2 by first interchanging the second-stage simultaneous moves and then coalescing the sequential moves of player 1. ▲

The next proposition, the main result of this note, clarifies the relationships between the properties of information feedback functions just introduced.

Proposition 1. (a) *Perfect feedback implies ex post perfect recall, observable payoffs and ex post observable deviators.*

(b) *Ex post perfect recall implies that*

$$F_i^{-1}(m) = \mathrm{proj}_{S_i} F_i^{-1}(m) \times \mathrm{proj}_{S_{-i}} F_i^{-1}(m) \tag{4}$$

for each $i \in I$ and $m \in M$. Therefore, in two-person games, ex post perfect recall implies ex post observable deviators. Furthermore, if Eq. (4) holds for each $m \in M$, then the identification correspondence can be written as

$$\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*) = \{\sigma_{-i} \in \times_{j \neq i} \Delta(S_j) : \hat{F}_i(\sigma_i^*, \sigma_{-i}) = \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)\}. \tag{5}$$

(c) *Separable feedback implies own-strategy independence of feedback and ex post observable deviators.*

(d) *In two-person games, own-strategy independence is equivalent to separability of feedback.*

Proof. (a) Fix any i and suppose that f_i is one-to-one (perfect feedback). Then, ex post perfect recall and observable payoffs are obviously satisfied. To check that ex post observable deviators also holds, note the following two facts. First, for each node y (including terminal nodes), $S(y) = \times_{j \in I} \mathrm{proj}_{S_j} S(y)$, where $S(y)$ denotes the set of pure strategy profiles reaching y. This follows from the observation that $\mathrm{proj}_{S_j} S(y)$ is the set of j's strategies selecting action $a_j(x \to y)$ for each node x of j preceding y. Therefore, if we pick $s_j \in S_j(y)$ for each $j \in I$, the path induced by $(s_j)_{j \in I}$ must reach y. Second, for each $m \in f_i(Z)$, $f_i^{-1}(m)$ is a singleton by assumption. Therefore,

$$F_i^{-1}(m) = S(f_i^{-1}(m)) = \times_{j \in I} \mathrm{proj}_{S_j} S(f_i^{-1}(m))$$
$$= \times_{j \in I} \mathrm{proj}_{S_j} F_i^{-1}(m).$$

(b) The following is a well-known property of perfect-recall games: for each player i and information set h of i, $S(h) = \mathrm{proj}_{S_i} S(h) \times \mathrm{proj}_{S_{-i}} S(h)$. Sets of preimages of messages $f_i^{-1}(m) \subseteq Z$ are just like information sets of i, and – under the assumption of ex post perfect recall – the aforementioned result applies to such ex post information sets as well. Therefore, Eq. (4) holds for each m. This implies that two-person games with ex post perfect recall must have ex post observable deviators. Now, suppose that (4) holds for each m. By Lemma 1, we only have to show that the right hand side of Eq. (5) is contained in the left hand side, that is, if σ_{-i} is such that $\hat{F}_i(\sigma_i^*, \sigma_{-i}) = \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)$, then $(\delta_{s_i} \times \sigma_{-i})(F_i^{-1}(m)) = (\delta_{s_i} \times \sigma_{-i}^*)(F_i^{-1}(m))$ for each m and $s_i \in \mathrm{supp}\sigma_i^*$. To ease notation, let $S_{i,i}(m) = \mathrm{proj}_{S_i} F_i^{-1}(m)$ (respectively, $S_{-i,i}(m) = \mathrm{proj}_{S_{-i}} F_i^{-1}(m)$) denote the sets of strategies of i (respectively, strategy profiles of $-i$) that allow for message m for i. Then $F_i^{-1}(m) = S_{i,i}(m) \times S_{-i,i}(m)$. Since $\hat{F}_i(\sigma_i^*, \sigma_{-i}) = \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)$ by assumption,

$$\begin{aligned}
\sigma_i^*(S_{i,i}(m)) \times \sigma_{-i}(S_{-i,i}(m)) &= (\sigma_i^* \times \sigma_{-i})(F_i^{-1}(m)) \\
&= \hat{F}_i(\sigma_i^*, \sigma_{-i})(m) \\
&= \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)(m) \\
&= (\sigma_i^* \times \sigma_{-i}^*)(F_i^{-1}(m)) \\
&= \sigma_i^*(S_{i,i}(m)) \times \sigma_{-i}^*(S_{-i,i}(m))
\end{aligned}$$

for every m. Therefore $\sigma_{-i}(S_{-i,i}(m)) = \sigma_{-i}^*(S_{-i,i}(m))$ for every m with $\sigma_i^*(S_{i,i}(m)) > 0$. Now, pick any $s_i \in \mathrm{supp}\sigma_i^*$. If $s_i \in S_{i,i}(m)$, then $\sigma_i^*(S_{i,i}(m)) \geq \sigma_i^*(s_i) > 0$ and $\delta_{s_i}(S_{i,i}(m)) = 1$. Therefore, the previous argument implies

$$(\delta_{s_i} \times \sigma_{-i})(F_i^{-1}(m)) = 1 \times \sigma_{-i}(S_{-i,i}(m))$$
$$= 1 \times \sigma_{-i}^*(S_{-i,i}(m)) = (\delta_{s_i} \times \sigma_{-i}^*)(F_i^{-1}(m)).$$

If $s_i \notin S_{i,i}(m)$ then $\delta_{s_i}(S_{i,i}(m)) = 0$ and

$$(\delta_{s_i} \times \sigma_{-i})(F_i^{-1}(m)) = 0 \times \sigma_{-i}(S_{-i,i}(m))$$
$$= 0 \times \sigma_{-i}^*(S_{-i,i}(m)) = (\delta_{s_i} \times \sigma_{-i}^*)(F_i^{-1}(m)).$$

(c) The right hand side of Eq. (3) – which defines separable feedback – is independent of s_i. Hence, separable feedback implies own-strategy independence of feedback. Fix any message m and pure strategy profile $s \in F_i^{-1}(m)$. Separable feedback implies that there is a profile of subsets $(C_j)_{j \neq i}$ such that

$$F_i^{-1}(m) = \mathrm{proj}_{S_i} F_i^{-1}(m) \times (\times_{j \neq i} C_j),$$

hence the ex post observable deviators property holds.

(d) Suppose that (Γ, f) satisfies own-strategy independence of feedback. Let $j = -i$ be the opponent of i. Then there is a partition $\mathcal{F}_{i,j}$ of $S_j = S_{-i}$ such that $\mathcal{F}_{-i}(s_i) = \mathcal{F}_{i,j}$ for each s_i. With this, we can construct a function $F_{i,j}$ so that Eq. (3) holds: Let $M_{i,j} = \mathcal{F}_{i,j}$ and $F_{i,j}(s_j) = S_{i,j}(s_j)$ for each s_j, where $S_{i,j}(s_j)$ is the atom of $\mathcal{F}_{i,j}$ containing s_j. ■

It can be shown by example that none of the converses of the implications in Proposition 1 is valid. Here we focus on parts (b) and (c).

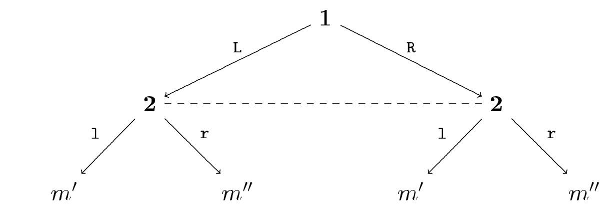

Fig. 5. Ex post perfect recall fails.

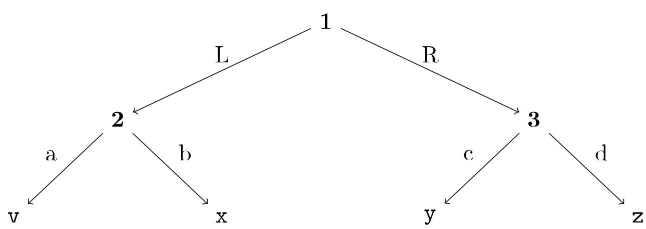

Fig. 6. A three-person extensive form.

Example 2. Consider the extensive form of Fig. 5, where $M = \{m', m''\}$, $F_i(L, l) = F_i(R, l) = m'$, $F_i(L, r) = F_i(R, r) = m''$ for each i. This is a two-person game with ex-post observable deviators and own-strategy independence of feedback, such that Eq. (4) holds:

$$\mathcal{F}_{-1}(L) = \mathcal{F}_{-1}(R) = \{\{l\}, \{r\}\},$$
$$\mathcal{F}_{-2}(l) = \mathcal{F}_{-2}(r) = \{\{L, R\}\} = \{S_1\}.$$

Hence, also Eq. (5) must hold, as stated in Proposition 1(b). Yet, ex post perfect recall fails, because $f_1 = F_1$ does not reveal whether action L or R is chosen. ▲

Example 3. Consider the 3-person perfect information extensive form in Fig. 6. Assume that f_1 reveals only player 1's action, $f_1(v) = f_1(x)$ and $f_1(y) = f_1(z)$, while *players* 2 *and* 3 *have perfect feedback*. This is a game with own-strategy independence of feedback and ex post observable deviators, but separable feedback fails. Formally,

$$\begin{cases} (f_1 \circ \zeta)^{-1}(f_1(v)) = (f_1 \circ \zeta)^{-1}(f_1(x)) = \{L\} \times S_2 \times S_3, \\ (f_1 \circ \zeta)^{-1}(f_1(y)) = (f_1 \circ \zeta)^{-1}(f_1(z)) = \{R\} \times S_2 \times S_3, \end{cases}$$
$$\Longrightarrow \mathcal{F}_{-1}(L) = \{\{S_2 \times S_3\}\} = \mathcal{F}_{-1}(R);$$

hence separable feedback holds trivially for player 1. However,

$$\begin{cases} \zeta^{-1}(v) = \{L\} \times \{a\} \times S_3, \zeta^{-1}(x) = \{L\} \times \{b\} \times S_3, \\ \zeta^{-1}(y) = \{R\} \times S_2 \times \{c\}, \zeta^{-1}(z) = \{R\} \times S_2 \times \{d\}, \end{cases}$$
$$\Longrightarrow \begin{cases} \mathcal{F}_{-2}(a) = \{\{L\} \times S_3, \{R\} \times \{c\}, \{R\} \times \{d\}\} = \mathcal{F}_{-2}(b), \\ \mathcal{F}_{-3}(c) = \{\{L\} \times \{a\}, \{L\} \times \{b\}, \{R\} \times S_2\} = \mathcal{F}_{-3}(d); \end{cases}$$

hence separable feedback fails for players 2 and 3. In words, the strategy of player 2 does not affect what he observes about the strategies of his co-players (own-strategy independence of feedback); furthermore he can always identify who caused a deviation from an expected ex post message (ex post observable deviators). Yet, he observes the strategy of player 3 if and only if player 1 chooses R. Therefore player 2 does not have separable feedback about his two co-players. And similarly with 2 and 3 reversed. ▲

4. Feedback and equilibrium concepts

In this section we define variations of the SCE concept and we illustrate the properties of information feedback showing how they affect the set of selfconfirming equilibria. The preliminary concepts and results of (4.1) draw on BCMM and are used in the rest of the paper.[13] In the rest of the section, we provide new results about relationships between equilibrium concepts (4.2) and existence (4.3). These new results depend on the strong assumption that all randomizations are feasible.

4.1. Preliminary definitions and results

The following notion of equilibrium is a variation on the Maxmin SCE concept of BCMM whereby agents are allowed to commit to mixed strategies. Recall that, unlike BCMM here we focus on equilibria where, under the population-game interpretation, all the agents in the same role/population play the same strategy; therefore, such equilibria are called "symmetric."

Definition 1. Fix a game with feedback (Γ, f). A mixed strategy profile σ^* is a symmetric Maxmin selfconfirming equilibrium (symMSCE) if, for each $i \in I$,

$$\sigma_i^* \in \arg\max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}).$$

Intuitively, if everyone keeps playing according to the mixed strategy profile σ^*, then each agent in population i infers from the dataset of his personal experiences that the true mixed strategy profile of the co-players belongs to the set $\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$. If he is extremely ambiguity averse and takes only this objective information into account, i.e. he does not further narrow down the set of distributions he believes possible, then he attaches to each mixed strategy σ_i the value $\min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i})$ and plays σ_i^* because it maximizes this value.[14] The equilibrium is called "symmetric" because, in the population-game scenario, it is implicitly assumed that all agents of the same population use the same (mixed) strategy.

It is useful to compare symMSCE with the original, ambiguity-neutral definition of selfconfirming (or "conjectural") equilibrium due to Battigalli (1987), which can be rephrased as follows.[15]

Definition 2. Fix a game with feedback (Γ, f). A mixed strategy profile σ^* is a symmetric Bayesian selfconfirming equilibrium (symBSCE) if, for each $i \in I$, there exists a belief $p_i \in \Delta(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$ such that

$$\sigma_i^* \in \arg\max_{\sigma_i \in \Delta(S_i)} \int U_i(\sigma_i, \sigma_{-i}) p_i(\mathrm{d}\sigma_{-i}).$$

The following observation is well known:

Lemma 3. *Every Nash equilibrium is also a symBSCE.*

Furthermore, Lemma 2 implies that—under observability of payoffs—Bayesian SCE is a refinement of Maxmin SCE, that is, ambiguity aversion (weakly) expands the equilibrium set.

Proposition 2 (*Cf. BCMM, Theorem 1*)**.** *If* (Γ, f) *is a game with* observable payoffs, *every symBSCE is also a symMSCE; in symbols, symBSCE* $\subseteq$ *symMSCE.*

---

13 For the proofs of these essentially known results, see Appendix.

14 This is the criterion suggested by Wald (1950). If instead we allow agents to subjectively deem impossible some distributions in $\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$, we obtain a more permissive concept consistent with the axioms of Gilboa and Schmeidler (1989): σ^* is an equilibrium if, for each $i \in I$, there exists a compact set $\Sigma_{-i} \subseteq \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$ such that $\sigma_i^* \in \arg\max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \Sigma_{-i}} U_i(\sigma_i, \sigma_{-i})$.

15 Battigalli's "conjectural equilibrium" is not framed within a population game, but it is equivalent to a notion of *symmetric* SCE in a population game. Under the assumption that players observe ex post the whole path of play, this is the "SCE with unitary uncorrelated beliefs" concept of Fudenberg and Levine (1993a).

Intuitively, under observable payoffs, the strategy played by each player in an SCE yields an objective lottery, because the induced distribution of payoffs is "observed" in the long run. On the other hand, alternative "untested" strategies do not necessarily yield an objective lottery. Therefore, increasing players' aversion to ambiguity can only increase their incentives to stick to their equilibrium strategies. Note that this argument is valid regardless of what mixed strategies are feasible.

The following result provides sufficient conditions for the equivalence of SCE and Nash equilibrium:

Proposition 3 (*Cf. BCMM, Proposition 1*)**.** *If* (Γ, f) *is a game with* observable payoffs *and* own-strategy independence of feedback*, then symmetric Bayesian and Maxmin SCE coincide with Nash equilibrium; in symbols, symBSCE = symMSCE = NE.*

Of course, in games without observable payoffs, symMSCE and Nash equilibrium can be very different, as one can easily check in many non-strictly competitive 2×2 games (e.g. the Battle of the Sexes) with trivial feedback.

4.2. Randomization and SCE

We prove a converse of Proposition 2 that, unlike that proposition, crucially relies on the strong assumption that all the mixed strategies in the simplex are implementable, but does not rely on assumptions about feedback. This gives conditions under which ambiguity aversion does not affect equilibrium outcomes. Before we state and prove the converse of Proposition 2, we need a preliminary result. For every subset Y of a Euclidean space, we let $\mathrm{co}(Y)$ denote its convex hull.

Lemma 4. *Let X be a convex and compact subset of Euclidean space $\mathbb{R}^m$ and Y a compact subset of Euclidean space $\mathbb{R}^n$. Let $U : X \times \mathrm{co}(Y) \to \mathbb{R}$ be a continuous function such that* (i) $x \mapsto U(x, y)$ *is quasi-concave for each* $y \in \mathrm{co}(Y)$ *and* (ii) $y \mapsto U(x, y)$ *is affine for each x. Then, for every*

$$x^* \in \arg\max_{x\in X} \min_{y\in \mathrm{co}(Y)} U(x, y),$$

there is a probability measure $p \in \Delta(Y)$ such that

$$x^* \in \arg\max_{x\in X} \int U(x, y)p(\mathrm{d}y).$$

Proof. Fix $x^* \in \arg\max_{x\in X} \min_{y\in\mathrm{co}(Y)} U(x, y)$. By the minimax theorem (Sion, 1958),

$$\max_{x\in X} \min_{y\in \mathrm{co}(Y)} U(x, y) = \min_{y\in \mathrm{co}(Y)} \max_{x\in X} U(x, y),$$

and, for every $y^* \in \arg\min_{y\in\mathrm{co}(Y)} \max_{x\in X} U(x, y)$, (x^*, y^*) is a saddle point. Thus,

$$x^* \in \arg\max_{x\in X} U(x, y^*).$$

Since $y^* \in \mathrm{co}(Y)$, there is a finite set $\{y_1, \dots, y_K\} \subseteq Y$ and vector of weights $(\lambda_k)_{k=1}^K \in \Delta(\{1, \dots, K\})$ such that $y^* = \sum_{k=1}^K \lambda_k y_k$. Then

$$p = \sum_{k=1}^{K} \lambda_k \delta_{y_k} \in \Delta(Y).$$

Since $U(x, y)$ is affine in its second argument

$$\int_Y U(x, y)p(\mathrm{d}y) = \sum_{k=1}^{K} \lambda_k \int_Y U(x, y)\delta_{y_k}(\mathrm{d}y) = \sum_{k=1}^{K} \lambda_k U(x, y_k) = U(x, y^*)$$

for every $x \in X$. The thesis follows. ■

Proposition 4. *Every symMSCE is also a symBSCE; in symbols, symMSCE ⊆ symBSCE.*

Proof. Fix a symMSCE σ^* and any player $i \in I$. The expected utility function U_i restricted to $\Delta(S_i) \times \mathrm{co}(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$ satisfies the assumptions of Lemma 4. Therefore there exists a belief $p_i \in \Delta(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$ such that

$$\sigma_i^* \in \arg\max_{\sigma_i\in\Delta(S_i)} \int U_i(\sigma_i, \sigma_{-i})p_i(\mathrm{d}\sigma_{-i}).$$

Therefore σ^* is a symBSCE. ■

It is easy to show that the inclusion can be strict in games without observable payoffs:

Example 4. Let (Γ, f) be a game with feedback such that (1) each player $i \in I$ has a *unique* and *fully mixed* maxmin strategy σ_i^* (as in Matching Pennies) and (2) each i has a constant feedback function, so that $\hat{\Sigma}_{-i}(\cdot) = \Delta(S_{-i})$. Then $(\sigma_i^*)_{i\in I}$ is the unique symMSCE of (Γ, f). But every mixed strategy profile is a symBSCE of (Γ, f). To see this, pick any $\bar{\sigma}_{-i}$ in the non-empty set $\arg\min_{\sigma_{-i}\in\Delta(S_{-i})} U_i(\sigma_i^*, \sigma_{-i})$, then $U_i(\sigma_i, \bar{\sigma}_{-i}) = U_i(\sigma_i^*, \bar{\sigma}_{-i})$ for all mixed strategies $\sigma_i \in \Delta(\mathrm{supp}\sigma_i^*) = \Delta(S_i)$. Hence, each σ_i is justified by the (trivially) confirmed Dirac belief $\delta_{\bar{\sigma}_{-i}}$. ▲

Propositions 2 and 4 imply that, in games with observable payoffs (and if all mixed strategies are feasible), ambiguity aversion does not affect the selfconfirming equilibrium set[16]:

Corollary 1. *If* (Γ, f) *is a game with* observable payoffs*, then symmetric Bayesian and Maxmin SCE coincide: symBSCE = symMSCE.*

Finally, we report another result relating symMSCE to Nash equilibrium, which follows from earlier results and Corollary 1. Recall that two mixed strategy profiles σ^* and $\bar{\sigma}$ are *realization equivalent* if they induce the same distribution over terminal nodes[17]; that is, if

$$\forall z \in Z, (\times_{i\in I} \sigma_i^*) \left(\zeta^{-1}(z)\right) = (\times_{i\in I} \bar{\sigma}_i) \left(\zeta^{-1}(z)\right).$$

Proposition 5. *If* (Γ, f) *is a* two-person *game with* perfect feedback*, then symmetric Bayesian and Maxmin SCE are realization-equivalent to Nash equilibrium.*

Proof. Since perfect feedback implies observable payoffs, Bayesian and Maxmin SCE coincide (Corollary 1). Every Nash equilibrium $\bar{\sigma}$ is also a symBSCE (Lemma 3). Battigalli (1987) proved that in two-person games with perfect recall and perfect feedback every symBSCE σ^* is realization-equivalent to some Nash equilibrium $\bar{\sigma}$.[18] The thesis follows. ■

4.3. Existence

Every finite game has a (mixed) Nash equilibrium. Therefore, every game with feedback (Γ, f) has a symmetric Bayesian SCE (Lemma 3). By Proposition 2, this implies the following existence result.

Theorem 1. *If a game with feedback* (Γ, f) *has* observable payoffs*, then* (Γ, f) *has a symmetric Maxmin SCE.*

16 Kuzmics (2015) reports a related result: if all randomizations are feasible, the observed choice of an ambiguity averse agent is consistent with expected utility maximization.

17 See Kuhn (1953).

18 For a proof in English, see the survey by Battigalli et al. (1992).

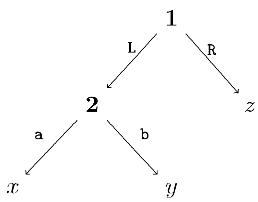

Fig. 7. A two-person PI game.

However, we do not have a general proof of existence of symMSCE. To see the difficulty, one can try to apply standard techniques to show that the correspondence

$$\bar{\sigma} \longmapsto \times_{i\in I} \arg\max_{\sigma_i\in\Delta(S_i)} \min_{\sigma_{-i}\in\hat{\Sigma}_{-i}(\bar{\sigma}_i,\bar{\sigma}_{-i})} U_i(\sigma_i,\sigma_{-i})$$

satisfies the conditions of Kakutani's fixed point theorem, i.e., that it is upper-hemicontinuous and non-empty, convex, compact valued. The problem is to show that the value function

$$V_i(\sigma_i|\bar{\sigma}_{-i}) = \min_{\sigma_{-i}\in\hat{\Sigma}_{-i}(\bar{\sigma}_i,\bar{\sigma}_{-i})} U_i(\sigma_i,\sigma_{-i}) \qquad (6)$$

is continuous in $(\sigma_i, \bar{\sigma}_{-i})$. This would be true if the identification correspondence $\bar{\sigma} \longmapsto \hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$ were continuous. It is easy to show that $\bar{\sigma} \longmapsto \hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$ is upper-hemicontinuous, because the pushforward map $\hat{F}_i$ is continuous.

Lemma 5. *The identification correspondence is upper-hemicontinuous and non-empty compact valued.*

Proof. First note that $\hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$ is non-empty because $\bar{\sigma}_{-i} \in \hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$. Next, we show that the graph of the identification correspondence, $\{(\bar{\sigma}', \sigma'_{-i}) : \hat{F}_i(\bar{\sigma}'_i, \sigma'_{-i}) = \hat{F}_i(\bar{\sigma}')\}$, is closed in the compact space $\Delta(S) \times (\times_{j\neq i} \Delta(S_j))$; this establishes the result. Take any converging sequence in this graph: $(\bar{\sigma}^k, \sigma^k_{-i}) \to (\bar{\sigma}, \sigma_{-i})$ with $\hat{F}_i(\bar{\sigma}^k_i, \sigma^k_{-i}) = \hat{F}_i(\bar{\sigma}^k)$ for each k. Since the pushforward map $\hat{F}_i : \times_{j\in I} \Delta(S_j) \to \Delta(M)$ is continuous, taking the limit for $k \to \infty$ we obtain $\hat{F}_i(\bar{\sigma}_i, \sigma_{-i}) = \hat{F}_i(\bar{\sigma})$. Thus $(\bar{\sigma}, \sigma_{-i}) \in \{(\bar{\sigma}', \sigma'_{-i}) : \hat{F}_i(\bar{\sigma}'_i, \sigma'_{-i}) = \hat{F}_i(\bar{\sigma}')\}$. ■

However, it can be shown by example that the identification correspondence is not necessarily lower-hemicontinuous. One reason is that what a player observes ex post about the strategies of the co-players may depend on his own strategy. As we observed, this is always the case for sequential games with perfect feedback that are not realization-equivalent to simultaneous-move games.

Example 5. Consider the extensive form in Fig. 7 and assume there is *perfect feedback*. The identification correspondence of player 1 is

$$\hat{\Sigma}_{-1}(\bar{\sigma}_1, \bar{\sigma}_2) = \begin{cases} \{\bar{\sigma}_2\}, & \text{if } \bar{\sigma}_1(L) > 0, \\ \Delta(S_2), & \text{if } \bar{\sigma}_1(L) = 0. \end{cases}$$

This correspondence is not lower-hemicontinuous at points $(\bar{\sigma}^0_1, \bar{\sigma}^0_2)$ such that $\bar{\sigma}^0_1(L) = 0$. To see this, consider the sequence $(\bar{\sigma}^n_1, \bar{\sigma}^n_2) \to (\bar{\sigma}^0_1, \bar{\sigma}^0_2)$ with $\bar{\sigma}^n_1(L) = 1/n$ and a mixed strategy $\sigma_2 \neq \bar{\sigma}^0_2$; then $\sigma_2 \in \hat{\Sigma}_{-1}(\bar{\sigma}^0_1, \bar{\sigma}^0_2) = \Delta(S_2)$, but $\sigma^n_2 \in \hat{\Sigma}_{-1}(\bar{\sigma}^n_1, \bar{\sigma}^n_2)$ implies $\sigma^n_2 = \bar{\sigma}^n_2$ for each n. Therefore $\sigma^n_2 \to \bar{\sigma}^0_2 \neq \sigma_2$. ▲

Even if there is own-strategy independence of feedback (hence $\hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$ is independent of $\bar{\sigma}_i$), the identification correspondence may violate lower-hemicontinuity in 3-person games; the reason is that what player i observes about j may depend on the strategy of another player k.

Example 6. Consider the 3-person extensive form in Fig. 6 with the same assumptions about feedback as in Example 3. As shown in that example, own-strategy independence is satisfied. Hence, we can write $\hat{\Sigma}_{-i}(\bar{\sigma}_{-i})$ instead of $\hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})$. The identification correspondence of player 2 is

$$\hat{\Sigma}_{-2}(\bar{\sigma}_1, \bar{\sigma}_3) = \begin{cases} \{\bar{\sigma}_1\} \times \{\bar{\sigma}_3\}, & \text{if } \bar{\sigma}_1(R) > 0, \\ \{\bar{\sigma}_1\} \times \Delta(S_3), & \text{if } \bar{\sigma}_1(R) = 0. \end{cases}$$

As in Example 5, it is easy to show that this correspondence is not lower-hemicontinuous at points $(\bar{\sigma}^0_1, \sigma^0_3)$ with $\bar{\sigma}^0_1(R) = 0$. Consider the sequence $(\bar{\sigma}^n_1, \bar{\sigma}^n_3) \to (\bar{\sigma}^0_1, \bar{\sigma}^0_3)$ with $\bar{\sigma}^n_1(R) = 1/n$ and a mixed strategy $\sigma_3 \neq \bar{\sigma}^0_3$; then, $\sigma_{-2} = (\bar{\sigma}^0_1, \sigma_3) \in \hat{\Sigma}_{-2}(\bar{\sigma}^0_1, \bar{\sigma}^0_3) = \{\delta_L\} \times \Delta(S_2)$, but $\sigma^n_{-2} \in \hat{\Sigma}_{-2}(\bar{\sigma}^n_1, \bar{\sigma}^n_3)$ implies $\sigma^n_3 = \bar{\sigma}^n_3$ for each n. Therefore $\sigma^n_3 \to \bar{\sigma}^0_3 \neq \sigma_3$. ▲

Observe that, as shown in Example 3, in this case separable feedback fails. Indeed, failure of separable feedback is necessary for the discontinuity of the identification correspondence.

Lemma 6. *In a game with separable feedback, the identification correspondence is continuous.*

Proof. Separable feedback implies that, for each player i, there is a profile of correspondences $(\hat{\Sigma}_{i,j}(\cdot))_{j\neq i}$, with $\sigma_j \longmapsto \hat{\Sigma}_{i,j}(\sigma_j) \subseteq \Delta(S_j)$, such that

$$\hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i}) = \times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j) = \times_{j\neq i} \big\{\sigma_j \in \Delta(S_j) : \forall C_j \in \mathcal{F}_{i,j}, \sigma_j(C_j) = \bar{\sigma}_j(C_j)\big\},$$

where $\sigma_j(C_j) = \sum_{s_j\in C_j} \sigma_j(s_j)$, and $\mathcal{F}_{i,j}$ is the partition of pre-images of $F_{i,j} : S_j \to M_{i,j}$. By Lemma 5, we only have to show that each correspondence $\bar{\sigma}_j \longmapsto \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$ is lower-hemicontinuous. Fix a converging sequence $\bar{\sigma}^n_j \to \bar{\sigma}_j$ and a point $\sigma_j \in \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$. To prove lower-hemicontinuity, we construct a selection $(\sigma^n_j)^\infty_{n=1}$ from $(\hat{\Sigma}_{i,j}(\bar{\sigma}^n_j))^\infty_{n=1}$ such that $\sigma^n_j \to \sigma_j$. For every n, s_j, and each atom $C_j \in \mathcal{F}_{i,j}$, let

$$\sigma^n_j(s_j) = \begin{cases} \bar{\sigma}^n_j(s_j), & \text{if } \sigma_j(C_j) = 0, \\ \sigma_j(s_j|C_j)\bar{\sigma}^n_j(C_j), & \text{if } \sigma_j(C_j) > 0. \end{cases}$$

First, note that $\sigma^n_j \in \hat{\Sigma}_{i,j}(\bar{\sigma}^n_j)$; that is, σ^n_j and $\bar{\sigma}^n_j$ assign the same probabilities to the atoms of the partition $\mathcal{F}_{i,j}$: For each $C_j \in \mathcal{F}_{i,j}$, if $\sigma_j(C_j) = 0$

$$\sigma^n_j(C_j) = \sum_{s_j\in C_j} \sigma^n_j(s_j) = \sum_{s_j\in C_j} \bar{\sigma}^n_j(s_j) = \bar{\sigma}^n_j(C_j);$$

if $\sigma_j(C_j) > 0$

$$\sigma^n_j(C_j) = \sum_{s_j\in C_j} \sigma^n_j(s_j) = \bar{\sigma}^n_j(C_j) \sum_{s_j\in C_j} \sigma_j(s_j|C_j) = \bar{\sigma}^n_j(C_j).$$

Thus, $\sigma^n_j(C_j) = \bar{\sigma}^n_j(C_j)$ for each $C_j \in \mathcal{F}_{i,j}$ and

$$\sum_{s_j} \sigma^n_j(s_j) = \sum_{C_j\in\mathcal{F}_{i,j}} \sigma^n_j(C_j) = \sum_{C_j\in\mathcal{F}_{i,j}} \bar{\sigma}^n_j(C_j) = 1.$$

Therefore, $\sigma^n_j \in \hat{\Sigma}_{i,j}(\bar{\sigma}^n_j)$.

Next, we show that $\sigma^n_j(s_j) \to \sigma_j(s_j)$ for each $s_j \in S_j$. Since $\sigma_j \in \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$, σ_j and $\bar{\sigma}_j$ agree on the partition $\mathcal{F}_{i,j}$. Therefore, if $s_j \in C_j$ with $\sigma_j(C_j) = 0$

$$\lim_{n\to\infty} \sigma^n_j(s_j) = \lim_{n\to\infty} \bar{\sigma}^n_j(s_j) = \bar{\sigma}_j(s_j) = 0 = \sigma_j(s_j);$$

if $s_j \in C_j$ with $\sigma_j(C_j) > 0$,

$$\lim_{n\to\infty} \sigma^n_j(s_j) = \sigma_j(s_j|C_j) \lim_{n\to\infty} \bar{\sigma}^n_j(C_j) = \sigma_j(s_j|C_j)\bar{\sigma}_j(C_j) = \sigma_j(s_j|C_j)\sigma_j(C_j) = \sigma_j(s_j). \quad ■$$

Theorem 2. *Every game with separable feedback has a symmetric Maxmin SCE.*

Proof. We prove that each correspondence

$$\bar{\sigma}_{-i} \longmapsto r_i(\times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j)) = \arg\max_{\sigma_i\in\Delta(S_i)} V_i(\sigma_i|\bar{\sigma}_{-i})$$

is non-empty convex compact valued, where $V_i(\sigma_i|\bar{\sigma}_{-i})$ is the value function defined in (6), the minimum of $U_i(\sigma_i,\cdot)$ under constraint $\sigma_{-i} \in \times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$. Since U_i is linear in σ_i, $V_i(\sigma_i|\bar{\sigma}_{-i})$ is concave in σ_i. Hence, $\bar{\sigma}_{-i} \longmapsto r_i(\times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j))$ has convex values. U_i is continuous in σ, and – by Lemma 6 – the identification correspondence $\bar{\sigma}_{-i} \longmapsto \times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$ is continuous; therefore, Berge's (minimum) theorem implies that $V_i(\sigma_i|\bar{\sigma}_{-i})$ is continuous in $(\sigma_i,\bar{\sigma}_{-i})$. By Berge's (maximum) theorem, the correspondence $\bar{\sigma}_{-i} \longmapsto r_i(\times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j))$ is upper-hemicontinuous non-empty compact valued.

Thus, $\bar{\sigma} \longmapsto \times_{i\in I} r_i(\times_{j\neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j))$ satisfies the assumptions of Kakutani's fixed point theorem. Every fixed point is a symmetric MSCE. ■

Theorem 3. *Every two-person game with own-strategy independence of feedback has a symmetric MSCE.*

Proof. By Proposition 1, a two-person game with own-strategy independence of feedback has separable feedback. Hence, Theorem 2 implies that the game has a symmetric MSCE. ■

5. Partially specified probabilities

In the final section of a decision theory paper, Lehrer (2012) defines a kind of mixed-strategy, maxmin selfconfirming equilibrium concept for games with "partially specified probabilities" (PSP). His PSP-equilibrium concept does not rely on the mass action interpretation of mixed strategies: (1) he assumes that each player i commits to a mixed strategy freely chosen from the whole simplex, and (2) he does *not* regard an equilibrium mixed strategy of i as the predictive measure obtained from a distribution over mixed strategies in a population of agents playing in role i. In this respect, Lehrer's PSP-equilibrium is comparable to the symmetric mixed selfconfirming equilibrium of BCMM that we analyze in this paper. We first analyze Lehrer's original definition, which relies on a separability assumption (5.1); then we discuss a generalization which is consistent with Lehrer's decision-theoretic analysis (5.2).

5.1. PSP-equilibrium with separability

Lehrer (2012) postulates the existence of a kind of probabilistic feedback, directly defined on the normal form of the game, that relies on implicit assumptions about information feedback. Specifically, he assumes that, for each player i and co-player j, there is a (finite) set of random variables $\mathcal{Y}_i^j \subseteq \mathbb{R}^{S_j}$ whose expected values are observed by i. The interpretation is that, if σ_j is the true mixed strategy played by j, then i observes (in the long run) the profile of expected values $(\mathbb{E}_{\sigma_j}(Y))_{Y\in\mathcal{Y}_i^j}$. Therefore, the set of partially specified mixed strategies of j (from i's point of view) when j actually plays σ_j^* is[19]

$$\hat{\Sigma}_{i,j}(\sigma_j^*) = \{\sigma_j \in \Delta(S_j) : \forall Y \in \mathcal{Y}_i^j, \mathbb{E}_{\sigma_j}(Y) = \mathbb{E}_{\sigma_j^*}(Y)\}.$$

Once we add sets of random variables $\mathcal{Y}_i^j$ for each $i \in I$ and $j \in I\setminus\{i\}$ to a game in strategic form $(I, (S_i, U_i)_{i\in I})$, we obtain a *game with partially specified probabilities*, $(I, (S_i, U_i, (\mathcal{Y}_i^j)_{j\neq i})_{i\in I})$.

Definition 3. Fix a game in strategic form with partially specified probabilities. A mixed strategy profile σ^* is a *PSP-equilibrium* if, for each $i \in I$,

$$\sigma_i^* \in \arg\max_{\sigma_i\in\Delta(S_i)} \min_{\sigma_{-i}\in\times_{j\neq i} \hat{\Sigma}_{i,j}(\sigma_j^*)} U_i(\sigma_i,\sigma_{-i}).$$

In order to compare SCE with PSP-equilibrium, we have to relate information feedback with Lehrer's partially specified probabilities. A game in extensive form with *separable feedback* (Γ, f) yields a game in strategic form with partially specified probabilities $(I, (S_i, U_i, (\mathcal{Y}_i^j)_{j\neq i})_{i\in I})$ as follows: For each $i \in I, j \in I\setminus\{i\}$ and $m_j \in M_j$, let $Y_{i,m_j} : S_j \to \{0, 1\}$ denote the indicator function of m_j; that is, $Y_{i,m_j}(s_j) = 1$ if and only if $F_{i,j}(s_j) = m_j$. Then $\mathcal{Y}_i^j = \{Y_{i,m_j} : m_j \in F_{i,j}(S_j)\}$. With this, we say that $(I, (S_i, U_i, (\mathcal{Y}_i^j)_{j\neq i})_{i\in I})$ is the *canonical PSP representation of* (Γ, f).

Remark 2. Fix a game with separable feedback (Γ, f) and its canonical PSP-representation $(I, (S_i, U_i, (\mathcal{Y}_i^j)_{j\neq i})_{i\in I})$. Then,

$$\hat{F}_{i,j}(\sigma_j)(m_{i,j}) = \mathbb{E}_{\sigma_j}(Y_{i,m_j})$$

for each player $i \in I$, co-player $j \neq i$, message $m_{i,j} \in M_{i,j}$, and mixed strategy $\sigma_j \in \Delta(S_j)$. Therefore, a mixed strategy profile is a symmetric Maxmin SCE of (Γ, f) *if and only if* it is a PSP-equilibrium of the canonical PSP-representation of (Γ, f).

With this connection between symMSCE and PSP-equilibrium of the canonical PSP-representation of a game with separable feedback, we can use results about the former to obtain results about the latter.

Separable feedback implies own-strategy independence of feedback. Therefore Remark 2 and Propositions 4 and 3 yield the following result.

Corollary 2. *Fix a game with separable feedback (Γ, f) and its canonical PSP-representation. Then*

(a) *every PSP-equilibrium is a symBSCE;*
(b) *if (Γ, f) has observable payoffs, PSP-equilibrium coincides with Nash equilibrium.*

Given Remark 2, Theorem 2 yields the following existence result.

Corollary 3. *The canonical PSP-representation of a game with separable feedback has a PSP-equilibrium.*

Lehrer (2012) states a general existence theorem for strategic-form games with partially specified probabilities, but he omits the proof (he just gives a hint that the result can be proved by standard methods). The analysis of Section 4.3 indicates that everything hinges on proving continuity (in particular lower-hemicontinuity) of the partial identification correspondences

$$\bar{\sigma}_j \longmapsto \{\sigma_j : \forall Y \in \mathcal{Y}_i^j, \mathbb{E}_{\sigma_j}(Y) = \mathbb{E}_{\bar{\sigma}_j}(Y)\} (i \in I, j \in I\setminus\{i\}).$$

The rest can be shown as in the proof of Theorem 2.

5.2. Generalized PSP-equilibrium

Games with PSPs as defined by Lehrer (2012) implicitly presume a sort of separability of feedback, a strong assumption

19 For comparability, we are using notation consistent with BCMM.

that may or may not hold in our analysis. One may therefore be inclined to think that Lehrer's games with PSPs are a less general construct than our games with feedback. But this is not true, the two constructs are not nested. Consider, for example, a simultaneous-moves game, hence a game Γ where $Z = S$. Suppose that, for each player i, there are functions $(F_{i,j} : S_j \to \mathbb{R})_{j \neq i}$ such that $f_i(s) = F_i(s) = (s_i, (F_{i,j}(s_j))_{j \neq i})$. We focused our attention on the PSP-representation of (Γ, f). But there are other meaningful games with PSPs consistent with these data. For example, we can assume that each player i observes in the long run only the first moments $\mathbb{E}_{\sigma_j}(F_{i,j})$, not the distributions $\hat{F}_{i,j}(\sigma_j)$ $(j \neq i)$.

Furthermore, the idea of modeling feedback with a system of PSPs, as suggested by the decision-theoretic axioms of Lehrer (2012), does not require separability. One could modify the definition of game with PSPs as follows. Specify, for each player i, a collection of random variables $\mathcal{Y}_i \subseteq \mathbb{R}^S$, with the interpretation that, for each mixed strategy profile σ, i observes in the long run the expected values $(\mathbb{E}_\sigma(Y))_{Y \in \mathcal{Y}_i}$. Assuming ex post perfect recall (cf. Proposition 1(b)), this yields the partial identification correspondence

$$\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*) = \{\sigma_{-i} \in \times_{j \neq i} \Delta(S_j) : \forall Y \in \mathcal{Y}_i, \mathbb{E}_{\sigma_i^*, \sigma_{-i}}(Y) = \mathbb{E}_{\sigma_i^*, \sigma_{-i}^*}(Y)\},$$

and a related generalization of PSP-equilibrium: σ^* is a *generalized PSP-equilibrium* if, for every player i,

$$\sigma_i^* \in \arg \max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}).$$

With this more general definition, every game with feedback (Γ, f) has a canonical PSP-representation such that the symMSCEs of (Γ, f) coincide with the generalized PSP-equilibria: for each i, let $\mathcal{Y}_i = \{Y_{i,m} : m \in F_i(S)\}$, where $Y_{i,m}(s) = 1$ (resp. $Y_{i,m}(s) = 0$) if and only if $F_i(s) = m$ (resp. $F_i(s) \neq m$).

Since Lehrer assumes that all randomizations are feasible, removing separability from the generalized definition of PSP-equilibrium makes it a refinement of symBSCE and different from Nash equilibrium even if payoffs are observable, that is, even if $U_i \in \mathcal{Y}_i$ for each $i \in I$. However, in this case PSP-equilibrium of the canonical PSP-representation is equivalent to symBSCE, i.e., ambiguity aversion does not affect the equilibrium set. It makes sense to further generalize PSP-equilibrium (as well as symMSCE) by positing a subset $\Sigma_i \subseteq \Delta(S_i)$ of feasible mixed strategies for each player i, as we mentioned in Section 2.2.[20] This allows ambiguity aversion to affect the equilibrium set even under observable payoffs, as shown in BCMM.

6. Discussion

In order to determine the selfconfirming equilibria of a game, one needs to specify the information feedback of each player. We analyze several properties of information feedback, and show how different notions of SCE are related to each other and to Nash equilibrium depending on which of these properties hold. Our analysis mostly focuses on four properties: perfect feedback, observable payoffs, own-strategy independence and separability of feedback. Perfect feedback implies observable payoffs (because each player knows the function associating his payoffs with terminal nodes), and separability implies own-strategy independence of feedback. Perfect feedback means that each player observes ex post the actions taken on the path by his co-players, which is natural in some applications. Observable payoffs is natural in a much wider range of applications, including all games where terminal nodes induce consumption (or monetary) allocations, and players have selfish preferences.[21] On the other hand, our extensive-form analysis clarifies that own-strategy independence, and *a fortiori* separability of feedback, are strong assumptions; in particular, they are hard to satisfy in non-simultaneous games (see Example 1). Games with feedback have a canonical representation in terms of partially specified probabilities, and symmetric Maxmin SCE is equivalent to Lehrer's PSP-equilibrium under this representation (see Remark 2 of Section 5.1 for the separable case and the generalization in Section 5.2).

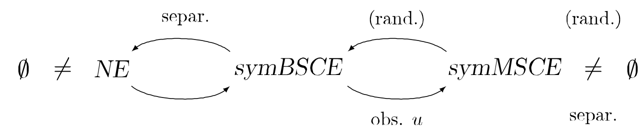

Fig. 8. Arrows denote implications; decorations specify assumptions about feedback (or feasibility of randomizations) under which the shown results are valid.

We show that, in games with observable payoffs, symmetric Bayesian and Maxmin SCE coincide; therefore, ambiguity aversion does not affect selfconfirming equilibrium (Corollary 1). However, this conclusion depends crucially on the strong assumption that agents can implement every mixed strategy, which implies that every symMSCE is also a symBSCE (Proposition 4). In two-person games with perfect feedback (hence with observable payoffs) symMSCE is realization-equivalent to symBSCE and Nash equilibrium (Proposition 5). In games with own-strategy independence of feedback and observable payoffs symMSCE coincides with Bayesian SCE and Nash equilibrium (Proposition 3). Existence of symMSCE is insured by either observability of payoffs, that makes each Nash equilibrium a symMSCE, or separability of feedback. The diagram in Fig. 8 summarizes these results, focusing only on two properties of feedback: separability (which implies own-strategy independence) and observability of payoffs.

The assumption that all randomizations are feasible is crucial for some of our results, including the "equilibrium-irrelevance" of ambiguity aversion when payoffs are observable (Proposition 4). It is well known that this assumption is not truly demanding for ambiguity-neutral decision makers. Indeed, every expected-utility-maximizing randomization σ_i^* can be implemented in an incentive-compatible way by planning to choose a feasible strategy according to the realization of an appropriate random variable, e.g., a sequence of coin tosses. Since all the pure strategies in the support of σ_i^* must be optimal, there is no incentive to deviate from the original plan and choose a different feasible strategy after the realization of the random variable. Suppose now that i is ambiguity-averse, e.g., in the maxmin sense. Then the plan to implement σ_i^* is incentive compatible if and only if it specifies, for each possible realization of the random variable, one of the feasible strategies of i with the highest minimum expected utility, given the set $\hat{\Sigma}_{-i}$ of co-players' mixtures that i deems possible. For example, if only the pure strategies are feasible, then σ_i^* can be implemented if and only if $\text{supp}\sigma_i^* \subseteq S_i^*\left(\hat{\Sigma}_{-i}\right)$, where $S_i^*\left(\hat{\Sigma}_{-i}\right) := \arg\max_{s_i \in S_i} \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}} U_i(s_i, \sigma_{-i})$. In equilibrium, $\hat{\Sigma}_{-i}$ is endogenously determined by the identification correspondence $\hat{\Sigma}_{-i}(\sigma)$. Therefore, we obtain the following modified definition of symMSCE in incentive-compatible mixed strategies: σ^* is an equilibrium

20 This does not conflict with decision-theoretic axiomatizations whereby agents are supposed to have preferences over a rich (convex) set of randomized acts, because the set of acts posited in decision theory is merely a set of *conceivable* alternatives, not the set of alternatives that are actually feasible in a decision problem.

21 See the discussion of observable payoffs in BCMM.

if, for each $i \in I$,

$$\sigma_i^* \in \arg \max_{\sigma_i \in \Delta\left(S_i^*\left(\hat{\Sigma}_{-i}(\sigma^*)\right)\right)} \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma^*)} U_i(\sigma_i, \sigma_{-i}).$$

This differs from Definition 1 because constraint $\sigma_i \in \Delta(S_i)$ is replaced by the incentive-compatibility constraint $\sigma_i \in \Delta\left(S_i^*\left(\hat{\Sigma}_{-i}(\sigma^*)\right)\right)$. We leave the analysis of this alternative concept to future research.

In games with separable feedback, we can compare SCE with Lehrer's (2012) original definition of PSP-equilibrium. Since symMSCE and PSP-equilibrium coincide, PSP-equilibrium refines symBSCE in such games (Corollary 2(a)). Since separability strengthens own-strategy independence of feedback, in games with separable feedback and observable payoffs, PSP-equilibrium coincides with symBSCE and Nash equilibrium (Corollary 2(b)). We suggest that separability of feedback can be removed from Lehrer's game-theoretic analysis, as it is not implied by its decision-theoretic underpinnings. Furthermore, there is no compelling conceptual reason to maintain – in our analysis, or Lehrer's – the assumption that players are able to implement any randomization. Once this assumption is removed, ambiguity aversion may affect the equilibrium set even if realized payoffs are observable.

In the rest of this section we discuss the related literature. We refer to BCMM (Section 4) for a review of the literature on selfconfirming equilibrium and similar concepts. Here we focus on information feedback and ambiguity attitudes. The main difference between the original notion of conjectural equilibrium of an extensive-form game due to Battigalli (1987) and the selfconfirming equilibrium concept of Fudenberg and Levine (1993a) concerns information feedback. Battigalli postulates a general feedback structure described by a profile of partitions of the set of terminal nodes that satisfy ex post perfect recall and observable payoffs, whereas Fudenberg and Levine consider the special case of perfect feedback.[22] We think that a notion of equilibrium whereby players best respond to confirmed beliefs should have the same name regardless of the assumptions about feedback. Therefore, in our terminology, we replaced "conjectural equilibrium" with the more self-explanatory "selfconfirming equilibrium".

To our knowledge, Lehrer (2012) provides the first definition of a concept akin to SCE where agents are not ambiguity neutral. As shown above, his definition implicitly requires a form of feedback separability, and it is equivalent to our symmetric Maxmin SCE when we consider the canonical PSP-representation of a game with separable feedback.

With the exception of separable feedback, the properties of information feedback analyzed here also appear in the previous literature on selfconfirming/conjectural equilibrium. Own-strategy independence of feedback was first introduced (with a different name) in the survey by Battigalli et al. (1992), it plays a prominent role in Azrieli (2009), and it is also emphasized in BCMM and Fudenberg and Kamada (2015).

Battigalli (1987) and Fudenberg and Kamada (2015) explicitly assume that information feedback satisfies ex post perfect recall. Although the analysis of BCMM is in the spirit of this assumption, formally they do not need it. This is related to their restriction of agents' choices to pure strategies. According to BCMM, an agent who plays pure strategy s_i and observes message m infers that the co-players' strategy profile belongs to $F_{i,s_i}^{-1}(m)$, where F_{i,s_i} is the section at s_i of the strategic-form feedback function F_i. Proposition 1(b) implies that, under ex post perfect recall, $F_{i,s_i}^{-1}(m) = \text{proj}_{S_{-i}} F_i^{-1}(m)$ for every strategy s_i consistent with message m.

The first paper where the observable payoffs assumption plays a prominent role is BCMM. Indeed, the main theorem and some other results of BCMM hold under this assumption.

22 Battigalli (1987) is written in Italian. Battigalli and Guaitoli (1988) is the first work in English with a definition of conjectural equilibrium. Fudenberg and Levine (1993a) developed the selfconfirming equilibrium concept independently. Besides the different assumptions about information feedback, Battigalli (1987) makes *stronger* assumptions about beliefs. Therefore the equilibrium concepts are not nested. Formally, under the assumption of perfect feedback, a conjectural equilibrium *à la* Battigalli is an SCE with unitary independent beliefs. For more on this see Battigalli (2012), the annotated extended abstract of Battigalli (1987).

Acknowledgments

We thank two anonymous referees for useful and constructive comments, and Veronica Roberta Cappelli, Enrico De Magistris, Nicodemo De Vito, Alejandro Francetich, Paolo Leonetti and Fabrizio Panebianco for valuable research assistance. The authors gratefully acknowledge the financial support of the European Research Council (BRSCDP-TEA-GA 230367, STRATEMOTIONS-GA 324219) and of the AXA Research Fund.

Appendix

To make the paper self-contained, we provide here the proofs that require minor modifications of those of known results.

A.1. *Proof of Lemma 2*

Fix $(\sigma_i^*)_{i\in I}$ and i arbitrarily. Define a function $v_i : f_i(Z) \to \mathbb{R}$ as follows: for each $m \in f_i(Z)$ pick some $z_m \in f_i^{-1}(m)$ and let $v_i(m) = u_i(z_m)$. By payoff observability, u_i is constant on $f_i^{-1}(m)$ for each $m \in f_i(Z)$; therefore, the choice of z_m in $f_i^{-1}(m)$ is immaterial and $u_i = v_i \circ f_i$. By definition, for each $\sigma_{-i} \in \hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right)$, $\hat{F}_i(\sigma_i^*, \sigma_{-i}) = \hat{F}_i(\sigma_i^*, \sigma_{-i}^*)$, that is, $\left(\widehat{f_i \circ \zeta}\right)(\sigma_i^*, \sigma_{-i}) = \left(\widehat{f_i \circ \zeta}\right)(\sigma_i^*, \sigma_{-i}^*)$. Therefore,

$$\begin{aligned}
\forall \sigma_{-i} \in \hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right), U_i(\sigma_i^*, \sigma_{-i}) \\
&= \sum_m v_i(m) \left(\widehat{f_i \circ \zeta}\right)(\sigma_i^*, \sigma_{-i})(m) \\
&= \sum_m v_i(m) \left(\widehat{f_i \circ \zeta}\right)(\sigma_i^*, \sigma_{-i}^*)(m) = U_i(\sigma_i^*, \sigma_{-i}^*). \quad \blacksquare
\end{aligned}$$

A.2. *Proof of Lemma 3*

Fix a Nash equilibrium σ^*. Then, for each $i \in I$,

$$\begin{aligned}
\sigma_i^* &\in \arg \max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i, \sigma_{-i}^*) \\
&= \arg \max_{\sigma_i \in \Delta(S_i)} \int U_i(\sigma_i, \sigma_{-i}) \delta_{\sigma_{-i}^*}(\mathrm{d}\sigma_{-i}),
\end{aligned}$$

where $p_i = \delta_{\sigma_{-i}^*}$ denotes the Dirac measure that assigns probability one to $\{\sigma_{-i}^*\}$. By definition, $\delta_{\sigma_{-i}^*} \in \Delta(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$. Therefore σ^* is a symBSCE supported by the profile of beliefs $(p_i)_{i\in I} = (\delta_{\sigma_{-i}^*})_{i\in I}$. $\blacksquare$

A.3. *Proof of Proposition 2*

Proof. Let σ^* be a symBCSE and fix $i \in I$ arbitrarily. By Definition 2, there is some prior belief $p_i \in \Delta\left(\hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right)\right)$ such that

$\sigma_i^* \in \arg\max_{\sigma_i \in \Delta(S_i)} \int U_i(\sigma_i, \sigma_{-i}) p_i(d\sigma_{-i})$. Therefore,

$$\forall \sigma_i \in \Delta(S_i), \quad \int U_i(\sigma_i^*, \sigma_{-i}) p_i(d\sigma_{-i})$$

$$\geq \int U_i(\sigma_i, \sigma_{-i}) p_i(d\sigma_{-i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}),$$

where the second inequality holds because $p_i \in \Delta\left(\hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right)\right)$.

By Lemma 2, the observable-payoffs assumption implies that $U_i(\sigma_i^*, \cdot)$ (the section of U_i at σ_i^*) is constant over the set $\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$. Thus,

$$\forall \sigma_i \in \Delta(S_i), \quad \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i^*, \sigma_{-i})$$

$$= \int U_i(\sigma_i^*, \sigma_{-i}) p_i(d\sigma_{-i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}),$$

where the equality holds because $p_i \in \Delta\left(\hat{\Sigma}_{-i}\left(\sigma_i^*, \sigma_{-i}^*\right)\right)$. Therefore, σ^* is a symMSCE. ■

A.4. Proof of Proposition 3

Proof. By Lemma 3 and Proposition 2, we only have to show that every symMSCE is a Nash equilibrium. Let σ^* be a symMSCE. Then, for each $i \in I$,

$$\forall \sigma_i \in \Delta(S_i), U_i(\sigma_i^*, \sigma_{-i}^*) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i^*, \sigma_{-i})$$

$$\geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}). \quad (7)$$

By Lemma 2, the observable-payoffs assumption implies that, for each $\sigma_i \in \Delta(S_i)$, $U_i(\sigma_i, \cdot)$ (the section of U_i at σ_i) is constant over the set $\hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)$. Thus,

$$\forall \sigma_i \in \Delta(S_i), \quad \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}) = U_i(\sigma_i, \sigma_{-i}^*). \quad (8)$$

Next we show that $\hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)$ does not depend on σ_i. Indeed, own-strategy independence of feedback implies that there is a partition $\mathcal{F}_{-i}$ of S_{-i} such that $\mathcal{F}_{-i} = \mathcal{F}_{-i}(s_i)$ for each s_i, where $\mathcal{F}_{-i}(s_i)$ is the partition of pre-images of $F_{i,s_i} : S_{-i} \to M$. Therefore,

$$\hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*) = \left\{\sigma_{-i} : \forall s_i \in \mathrm{supp}\sigma_i, \forall C_{-i} \in \mathcal{F}_{-i}(s_i),\right.$$
$$\left.\sigma_{-i}(C_{-i}) = \sigma_{-i}^*(C_{-i})\right\}$$
$$= \left\{\sigma_{-i} : \forall C_{-i} \in \mathcal{F}_{-i}, \sigma_{-i}(C_{-i}) = \sigma_{-i}^*(C_{-i})\right\}.$$

This implies

$$\forall \sigma_i \in \Delta(S_i), \quad \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i})$$

$$= \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}). \quad (9)$$

Expressions (7), (9) and (8) yield:

$$\forall \sigma_i \in \Delta(S_i), U_i(\sigma_i^*, \sigma_{-i}^*) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i^*, \sigma_{-i})$$

$$\geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i})$$

$$= \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}) = U_i(\sigma_i, \sigma_{-i}^*).$$

Hence σ^* is a Nash equilibrium. ■

References

Azrieli, Y., 2009. On pure conjectural equilibrium with non-manipulable information. Internat. J. Game Theory 38, 209–219.
Battigalli, P., 1987. Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali (thesis), Università Bocconi, Milano, Typescript.
Battigalli, P., 1994. Structural consistency and strategic independence in extensive games. Ric. Econom. 48, 357–376.
Battigalli, P., 1997. Games with observable deviators. In: Battigalli, P., Montesano, A., Panunzi, F. (Eds.), Decisions, Games and Markets. Kluwer Academic Publishers, Dordrecht.
Battigalli, P., 2012. Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. Annotated extended abstract and comments, typescript, Università Bocconi.
Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2014. Selfconfirming equilibrium and sequential rationality with ambiguity averse agents. Mimeo, Università Bocconi.
Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2015. Self-confirming equilibrium and model uncertainty. Amer. Econ. Rev. 105, 646–677.
Battigalli, P., Gilli, M., Molinari, M.C., 1992. Learning and convergence to equilibrium in repeated strategic interaction. Res. Econ. (Ric. Econom.) 46, 335–378.
Battigalli, P., Guaitoli, D., 1988. Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information. In: Quaderni di Ricerca 1988-6, I.E.P.. Università Bocconi (published in Decisions, Games and Markets, Kluwer, Dordrecht, 1997, pp. 97–124).
Fudenberg, D., Kamada, Y., 2015. Rationalizable partition-confirmed equilibrium. Theoret. Econ. 10, 775–806.
Fudenberg, D., Kreps, D., 1995. Learning in extensive games, I: self-confirming equilibria. Games Econom. Behav. 8, 20–55.
Fudenberg, D., Levine, D.K., 1993a. Self-confirming equilibrium. Econometrica 61, 523–545.
Fudenberg, D., Levine, D.K., 1993b. Steady state learning and Nash equilibrium. Econometrica 61, 547–573.
Gilboa, I., Marinacci, M., 2013. Ambiguity and the Bayesian paradigm. In: Acemoglu, D., Arellano, M., Dekel, E. (Eds.), Advances in Economics and Econometrics: Theory and Applications (Vol. 1). Cambridge University Press, Cambridge.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior. J. Math. Econom. 18, 141–153.
Gilli, M., 1999. Adaptive learning in imperfect monitoring games. Rev. Econ. Dynam. 2, 472–485.
Kalai, E., Lehrer, E., 1995. Subjective games and equilibria. Games Econom. Behav. 8, 123–163.
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.
Kuzmics, C., 2015. Abraham Wald's complete class theorem and Knightian uncertainty. Mimeo, Bielefeld University.
Lehrer, E., 2012. Partially-specified probabilities: decisions and games. Amer. Econ. J.: Macroecon. 4, 70–100.
Leonetti, P., 2014. Equivalent extensive form games (M.Sc. thesis), Bocconi University, Typescript.
Marinacci, M., 2015. Model uncertainty. J. Eur. Econ. Assoc. 13, 998–1076.
Osborne, M., Rubinstein, A., 1994. A Course in Game Theory. MIT Press, Cambridge, MA.
Sion, M., 1958. On general minimax theorems. Pacific J. Math. 171–176.
Wald, A., 1950. Statistical Decision Functions. Wiley, New York.
Weibull, J.W., 1996. "The mass action interpretation" in H. Kuhn et al. "The work of John Nash in game theory". J. Econom. Theory 69, 153–185.