

Analysis of Information Feedback and Selfconfirming Equilibrium

P. Battigalli*, S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci
Università Bocconi and IGER[†]

Draft of December 2013

Abstract

Our 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. Here, we discuss in detail the properties of this key feature of recurrent interaction. This allows us to elucidate our notion of Maxmin selfconfirming equilibrium, and to compare it with an equilibrium concept due to Lehrer.

KEYWORDS: Selfconfirming equilibrium, conjectural equilibrium, information feedback, partially specified probabilities.

JEL CLASSIFICATION: C72, D80.

*Corresponding Author. Email: pierpaolo.battigalli@unibocconi.it; Tel. +39 02-58365402; Fax +39 02-58365630.

[†]Via Roentgen, 1; 20136, Milan (MI), Italy.

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 [12, 1993], Fudenberg and Kreps [10, 1995] and Gilli [16, 1999]). Battigalli, Cerreia-Vioglio, Maccheroni and Marinacci [9, 2011] (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, which encompasses the traditional notions of conjectural equilibrium (Battigalli [3, 1987] and Battigalli and Guaitoli [7, 1988]) and selfconfirming equilibrium (Fudenberg and Levine [11, 1993]) as special cases, requires the specification of an ex post information structure, or information feedback. Specifically, the information-feedback function describes the personal experience of an agent 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.

In this paper, we define several properties of information feedback, we study their relationships, and we illustrate them through the analysis of symmetric Maxmin SCE, a special case of the BCMM equilibrium concept. Finally, we note that each game with “separable feedback” has a canonical representation as a game with partially specified probability in the sense of Lehrer [18, 2012]. Under this representation, our symmetric Maxmin SCE is equivalent to the equilibrium concept put forward by Lehrer. With this, our results imply that, in the canonical representation of a game with separable feedback, Lehrer's equilibrium is a refinement of the traditional SCE concept, and under observability of payoffs (i.e., of the realized own utility), it is equivalent to mixed Nash equilibrium.

The rest of the paper is organized as follows. Section 2 defines extensive-form games with feedback and Maxmin SCE; Section 3 analyzes the properties of information feedback; Section 4 relates Maxmin SCE to the traditional SCE concept and Nash equilibrium; Section 5 analyzes existence of Maxmin SCE; Section 6 relates information feedback to partially specified probabilities; Section 7 discusses the related literature and concludes.

2 Games with feedback and selfconfirming equilibrium

Throughout the analysis, we consider a finite game Γ in extensive form with perfect recall and no chance moves. We use the following notation for some 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*;
- $u_i : Z \rightarrow \mathbb{R}$ is the *payoff* (vNM utility) *function* of player i ;
- $S = \times_{i \in I} S_i$ is the finite set of *pure-strategy* profiles;

¹For a discussion on the literature of choice under ambiguity, see the survey of Gilboa and Marinacci [?, 2013].

- $\zeta : S \rightarrow Z$ 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 $\times_{j \in I} \Delta(S_j)$.

In their set up, BCMM specify, for each player role $i \in I$, a *feedback function*

$$f_i : Z \rightarrow M,$$

representing what i can observe ex post about the path of play. 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 [11, 1993]: 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 $F_i = f_i \circ \zeta : S \rightarrow M$. This, in turn, yields the pushforward map $\hat{F}_i : \times_{j \in I} \Delta(S_j) \rightarrow \Delta(M)$ 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 mixed-strategy profile σ .

We assume (informally) that the 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 to play the strategic form of Γ , then get feedback according to f and are separated and re-matched to play again. As in Section 6 of BCMM, we assume that agents can commit to any mixed strategy. 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 [22, 1996]). We also (informally) assume that each agent in role 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 .

For the sake of simplicity, here we focus on the *symmetric* case, where, for each role $i \in I$, each agent in sub-population i plays the same mixed strategy. 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 $(\hat{F}_i(s_i, \sigma_{-i}^*))_{s_i \in \text{supp}\sigma_i^*}$ in the long run;² collecting these conditional distributions, such an agent infers that the opponents' mixed strategy profile belongs to the set

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

We call $\hat{\Sigma}_{-i}(\cdot, \cdot)$ the (partial) *identification correspondence*. By inspection of the definition, one can see that this correspondence is non-empty ($\sigma_{-i}^* \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$) 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.

Proof. The inclusion follows from the following 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, consider the perfect information game of Figure 1:




Figure 1: A perfect information game.

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

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

²Whenever no confusion may arise, s_i is identified with δ_{s_i} , the Dirac measure supported by s_i .

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 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.* We show below (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 .

BCMM define a notion of mixed strategy Maxmin SCE whereby agents do not rule out any σ_{-i} consistent with their conditional distributions of observations and are extremely ambiguity averse. In the context of a population game, different agents in the same population $i \in I$ could play different mixed strategies, justified by different datasets of personal experiences. But, as explained above, here we focus on the *symmetric* case, in which each agent of the same population plays the same mixed strategy in equilibrium (for the more general definition see Section 6 of BCMM).

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 $\min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i})$, as suggested by Wald [21, 1950], and plays σ_i^* because it maximizes this value.³ It is useful to compare symMSCE with the original definition of selfconfirming equilibrium due to Battigalli [3, 1987], which can be rephrased as follows.⁴

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

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

³If 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 [15, 1989]: σ^* is an equilibrium if, for each $i \in I$, there exist a compact set $\Sigma_{-i} \subseteq \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)$ such that $\sigma_i^* \in \arg \max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \Sigma_{-i}} U_i(\sigma_i, \sigma_{-i})$.

⁴Battigalli’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 [11, 1993].

Though the following observation is well known, we provide a proof for the reader's convenience:

Lemma 2 *Every Nash equilibrium is also a symBSCE.*

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

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

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

3 Properties of information feedback

To analyze the properties of information feedback it is convenient to introduce additional notation, summarized by Table 1.

Notation	Definition
$x \prec y$ ($x \preceq y$)	Node x precedes (weakly) node y
$H_i(x)$	Information set of i containing x
$a_i(x \rightarrow y)$	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
$F_{i,s_i}(\cdot) = F_i(s_i, \cdot) : S_{-i} \rightarrow M$	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

Table 1: Additional 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 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 $i \in I$, $u_i : Z \rightarrow \mathbb{R}$ is f_i -measurable; that is,⁵

$$\forall z', z'' \in Z, f_i(z') = f_i(z'') \Rightarrow u_i(z') = u_i(z'');$$

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 $\mathcal{F}_{-i}(s_i) = \mathcal{F}_{-i}(t_i)$;

⁵Hence, $U_i : S \rightarrow \mathbb{R}$ is F_i -measurable.

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 $a_i(x' \rightarrow z') \neq a_i(x'' \rightarrow z'')$, then $f_i(z') \neq f_i(z'')$;
5. *ex post observable deviators* if, for every $i \in I$ and $m \in F_i(S)$,

$$F_i^{-1}(m) = \times_{j \in I} \text{proj}_{S_j} F_i^{-1}(m); \quad (2)$$

6. *separable feedback* if, for every $i \in I$, there are onto functions $(F_{i,j} : S_j \rightarrow M_{i,j})_{j \neq i}$ such that, for each $s_i \in S_i$,

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

In words, *perfect feedback* means that each player observes (ex post) the complete path of play. *Observable payoffs* says that each player observes his realized vNM utility.

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 *Under own-strategy independence of feedback, 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 [17, 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 [11, 1993] and Battigalli [4, 1994]. Using the definition of the latter, a game Γ has observable deviators if $S(h) = \times_{j \in I} \text{proj}_{S_j} S(h)$ for each player i and information set h of i , where $S(h)$ is the set of strategy profiles reaching h .⁶ Battigalli [5, 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.

⁶Battigalli [4, 1994] uses this property in the context of a discussion 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.

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 by iteratively interchanging simultaneous moves, and coalescing sequential moves by the same player.⁷ 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.




Figure 2: An extensive form with perfect feedback.

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

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

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

⁷In 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.




Figure 3: A game obtained from Figure 2.




Figure 4: A game obtained from Figure 3.

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*

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

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

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

- (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), $S(y) = \times_{j \in I} \text{proj}_{S_j} S(y)$, where $S(y)$ denotes the set of pure strategy profiles reaching y . This follows from the observation that $\text{proj}_{S_j} S(y)$ is the set of j 's strategies selecting action $a_j(x \rightarrow y)$ 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,

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

(b) The following is a well-known property of perfect-recall games: for each player i and information set h of i , $S(h) = \text{proj}_{S_i} S(h) \times \text{proj}_{S_{-i}} S(h)$. 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 $s_i \in \text{supp} \sigma_i^*$. To ease notation, let $S_{i,i}(m) = \text{proj}_{S_i} F_i^{-1}(m)$ (respectively, $S_{-i,i}(m) = \text{proj}_{S_{-i}} F_i^{-1}(m)$) 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 $s_i \in \text{supp} \sigma_i^*$. 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

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

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

$$\begin{aligned} (\delta_{s_i} \times \sigma_{-i})(F_i^{-1}(m)) &= 0 \times \sigma_{-i}(S_{-i,i}(m)) = \\ 0 \times \sigma_{-i}^*(S_{-i,i}(m)) &= (\delta_{s_i} \times \sigma_{-i}^*)(F_i^{-1}(m)). \quad \square \end{aligned}$$

(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

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

hence the ex post observable deviators property holds. \square

(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).




Figure 5: Ex post perfect recall fails.

Example 2 Consider the extensive form of Figure 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:

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

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.




Figure 6: A three-person extensive form.

Example 3 Consider the 3-person perfect information extensive form in Figure 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,

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

hence separable feedback holds trivially for player 1. But

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

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 Relationships between equilibrium concepts

In this section we illustrate the properties of information feedback showing how they shape the set of symMSCE. Specifically, we report some results (mostly due to BCMM) characterizing symMSCE under a variety of assumptions about feedback.

We first compare Maxmin to Bayesian self-confirming equilibrium. We show that the former is a refinement of the latter, but under observable payoffs they are equivalent. Indeed, 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. This informal argument provides intuition for the following result.

Proposition 2 (cf. BCMM Theorem 6) *If (Γ, f) is a game with observable payoffs, every symBSCE is also a symMSCE. In symbols, $\text{symBSCE} \subseteq \text{symMSCE}$.*

Note that this result holds also if players’ feasible choices are restricted to a subset of the mixed-strategy simplex, e.g., only the pure strategies.

Before we state and prove a converse, we need a preliminary result. For every subset Y of a Euclidean space, we let $\text{co}(Y)$ denote its convex hull.

Lemma 3 *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 $U : X \times \text{co}(Y) \rightarrow \mathbb{R}$ be a continuous function such that (i) $x \mapsto U(x, y)$ is quasi-concave for each $y \in \text{co}(Y)$ and (ii) $y \mapsto U(x, y)$ is affine for each x . Then, for every*

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

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

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

Proof Fix $x^* \in \arg \max_{x \in X} \min_{y \in \text{co}(Y)} U(x, y)$. By the minimax theorem (Sion, 1953),

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

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

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

Since $y^* \in \text{co}(Y)$, 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

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

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

The following converse of Proposition 2 crucially relies on the assumption that players' feasible set is the whole mixed-strategy simplex.

Proposition 3 (cf. BCMM Proposition 19) *Every symMSCE is also a symBSCE: symMSCE \subseteq symBSCE.*

Proof Fix a symMSCE σ^* and any player $i \in I$. The expected utility function U_i restricted to $\Delta(S_i) \times \text{co}(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$ satisfies the assumptions of Lemma 3. Therefore there exists a belief $p_i \in \Delta(\hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*))$ such that

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

Therefore σ^* is a symBSCE. \blacksquare

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 $\sigma_i \in \Delta(\text{supp} \sigma_i^*) = \Delta(S_i)$. Hence, each σ_i is justified by the (trivially) confirmed Dirac belief $\delta_{\bar{\sigma}_{-i}}$.*

Propositions 2 and 3 imply that, in games with observable payoffs, ambiguity aversion does not affect symmetric SCE:

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

Next, we report results relating symMSCE to Nash equilibrium. The first one concerns the induced distributions of outcomes. Recall that two mixed strategy profiles σ^* and $\bar{\sigma}$ are *realization equivalent* if they induce the same distribution over terminal nodes;⁸ that is, if

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

Proposition 4 (cf. BCMM Proposition 21) *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 2). Battigalli (1987) proved that in two-person games with perfect recall every symBSCE σ^* is realization-equivalent to some Nash equilibrium $\bar{\sigma}$.⁹ The thesis follows. \blacksquare

Another link between symMSCE and Nash equilibrium is given by the following result:

Proposition 5 (cf. BCMM Proposition 10) *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: $\text{symBSCE} = \text{symMSCE} = NE$.*

Proof¹⁰ By Lemma 2 and Proposition 2, we only have to show that every symMSCE is a Nash equilibrium. Let σ^* be a symMSCE. Then, for each $i \in I$,

$$\forall \sigma_i \in \Delta(S_i), U_i(\sigma_i^*, \sigma_{-i}^*) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i^*, \sigma_{-i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}). \quad (6)$$

By Lemma 1 in BCMM, 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,

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

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 $F_{i,s_i} : S_{-i} \rightarrow M$. Therefore,

$$\begin{aligned} \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*) &= \{ \sigma_{-i} : \forall s_i \in \text{supp} \sigma_i, \forall C_{-i} \in \mathcal{F}_{-i}(s_i), \sigma_{-i}(C_{-i}) = \sigma_{-i}^*(C_{-i}) \} \\ &= \{ \sigma_{-i} : \forall C_{-i} \in \mathcal{F}_{-i}, \sigma_{-i}(C_{-i}) = \sigma_{-i}^*(C_{-i}) \}. \end{aligned}$$

⁸See Kuhn [17, 1953].

⁹For a proof in English, see the survey by Battigalli *et al.* [8, 1992].

¹⁰BCMM do not explicitly provide a proof of this version of the result: their Proposition 10 concerns the case where agents are restricted to pure strategies, and mixed strategies only represent statistical distributions of pure strategies.

This implies

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

Expressions (6), (8) and (7) yield:

$$\begin{aligned} \forall \sigma_i \in \Delta(S_i), U_i(\sigma_i^*, \sigma_{-i}^*) &= \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i^*, \sigma_{-i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i^*, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}) \\ &= \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\sigma_i, \sigma_{-i}^*)} U_i(\sigma_i, \sigma_{-i}) = U_i(\sigma_i, \sigma_{-i}^*). \end{aligned}$$

Hence σ^* is a Nash equilibrium. ■

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.

5 Equilibrium existence

Every finite game has a (mixed) Nash equilibrium. Therefore, every game with feedback (Γ, f) has a symmetric Bayesian SCE (Lemma 2). 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.*

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

$$V_i(\sigma_i | \bar{\sigma}_{-i}) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(\bar{\sigma}_i, \bar{\sigma}_{-i})} U_i(\sigma_i, \sigma_{-i}) \quad (9)$$

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 4 *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: $(\bar{\sigma}^k, \sigma_{-i}^k) \rightarrow (\bar{\sigma}, \sigma_{-i})$ with $\hat{F}_i(\bar{\sigma}_i^k, \sigma_{-i}^k) = \hat{F}_i(\bar{\sigma}^k)$ for each

k . Since the pushforward map $\hat{F}_i : \times_{j \in I} \Delta(S_j) \rightarrow \Delta(M)$ is continuous, taking the limit for $k \rightarrow \infty$ we obtain $\hat{F}_i(\bar{\sigma}_i, \sigma_{-i}) = \hat{F}_i(\bar{\sigma})$. Thus $(\bar{\sigma}, \sigma_{-i}) \in \{(\bar{\sigma}', \sigma'_{-i}) : \hat{F}_i(\bar{\sigma}', \sigma'_{-i}) = \hat{F}_i(\bar{\sigma}')\}$. ■

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.




Figure 7: A two person PI game.

Example 5 Consider the extensive form in Figure 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 $(\bar{\sigma}_1^0, \bar{\sigma}_2^0)$ such that $\bar{\sigma}_1^0(L) = 0$. To see this, consider the sequence $(\bar{\sigma}_1^n, \bar{\sigma}_2^n) \rightarrow (\bar{\sigma}_1^0, \bar{\sigma}_2^0)$ with $\bar{\sigma}_1^n(L) = 1/n$ and a mixed strategy $\sigma_2 \neq \bar{\sigma}_2^0$; then $\sigma_2 \in \hat{\Sigma}_{-1}(\bar{\sigma}_1^0, \bar{\sigma}_2^0) = \Delta(S_2)$, but $\sigma_2^n \in \hat{\Sigma}_{-1}(\bar{\sigma}_1^n, \bar{\sigma}_2^n)$ implies $\sigma_2^n = \bar{\sigma}_2^n$ for each n . Therefore $\sigma_2^n \rightarrow \bar{\sigma}_2^0 \neq \sigma_2$.

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 Figure 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 $(\bar{\sigma}_1^0, \bar{\sigma}_3^0)$ with $\bar{\sigma}_1^0(R) = 0$. Consider the sequence $(\bar{\sigma}_1^n, \bar{\sigma}_3^n) \rightarrow (\bar{\sigma}_1^0, \bar{\sigma}_3^0)$ with $\bar{\sigma}_1^n(R) = 1/n$ and a mixed strategy $\sigma_3 \neq \bar{\sigma}_3^0$; then, $\sigma_3 \in \hat{\Sigma}_{-2}(\bar{\sigma}_1^0, \bar{\sigma}_3^0) = \{\delta_L\} \times \Delta(S_3)$, but $\sigma_3^n \in \hat{\Sigma}_{-2}(\bar{\sigma}_1^n, \bar{\sigma}_3^n)$ implies $\sigma_3^n = \bar{\sigma}_3^n$ for each n . Therefore $\sigma_3^n \rightarrow \bar{\sigma}_3^0 \neq \sigma_3$.

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 5 *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 $\sigma_j \mapsto \hat{\Sigma}_{i,j}(\sigma_j) \subseteq \Delta(S_j)$, such that

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

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 $F_{i,j} : S_j \rightarrow M_{i,j}$. By Lemma 4 and [1, Theorem 17.28], we only have to show that each correspondence $\bar{\sigma}_j \mapsto \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$ is lower-hemicontinuous. Fix a converging sequence $\bar{\sigma}_j^n \rightarrow \bar{\sigma}_j$ and a point $\sigma_j \in \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$. To prove lower-hemicontinuity, we construct a selection $(\sigma_j^n)_{n=1}^\infty$ from $(\hat{\Sigma}_{i,j}(\bar{\sigma}_j^n))_{n=1}^\infty$ such that $\sigma_j^n \rightarrow \sigma_j$. For every n , s_j , and each atom $C_j \in \mathcal{F}_{i,j}$, let

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

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

$$\sigma_j^n(C_j) = \sum_{s_j \in C_j} \sigma_j^n(s_j) = \sum_{s_j \in C_j} \bar{\sigma}_j^n(s_j) = \bar{\sigma}_j^n(C_j);$$

if $\sigma_j(C_j) > 0$

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

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

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

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

Next, we show that $\sigma_j^n(s_j) \rightarrow \sigma_j(s_j)$ 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 $\sigma_j(C_j) = 0$

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

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

$$\lim_{n \rightarrow \infty} \sigma_j^n(s_j) = \sigma_j(s_j|C_j) \lim_{n \rightarrow \infty} \bar{\sigma}_j^n(C_j) = \sigma_j(s_j|C_j)\bar{\sigma}_j(C_j) = \sigma_j(s_j|C_j)\sigma_j(C_j) = \sigma_j(s_j).$$

■

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

Proof. We prove that each correspondence

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

is non-empty convex compact valued, where $V_i(\sigma_i | \bar{\sigma}_{-i})$ is the value function defined in (9), 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, $\bar{\sigma}_{-i} \mapsto r_i(\times_{j \neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j))$ has convex values. U_i is continuous in σ , and – by Lemma 5 – the identification correspondence $\bar{\sigma}_{-i} \mapsto \times_{j \neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j)$ 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 $\bar{\sigma}_{-i} \mapsto r_i(\times_{j \neq i} \hat{\Sigma}_{i,j}(\bar{\sigma}_j))$ is upper-hemicontinuous non-empty compact valued.

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

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. \blacksquare

6 Partially specified probabilities

In the final section of a decision theory paper, Lehrer [18, 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.

Lehrer [18, 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^* is¹¹

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

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*, $(I, (S_i, U_i, (\mathcal{Y}_i^j)_{j \neq i})_{i \in I})$.

¹¹For comparability, we are using notation consistent with BCMM.

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 $Y_{i,m_j} : S_j \rightarrow \{0, 1\}$ 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) .

Now that we have established this link 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 3 and 5 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 Lehrer [18, 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 5 indicates that everything hinges on proving continuity (in particular lower-hemicontinuity) of the partial identification correspondences

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

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

7 Discussion

In order to obtain 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 identifies four crucial 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.¹² On the other hand, we argue that own-strategy independence, and *a fortiori* separability of feedback, are strong assumptions. Games with separable 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 (Remark 2).

We show that, in games with observable payoffs, symmetric Bayesian and Maxmin SCE coincide, hence ambiguity aversion does not affect selfconfirming equilibrium (Corollary 1). We observe that this conclusion depends crucially on the strong assumption that agents can commit to any mixed strategy, which implies that every symMSCE is also a symBSCE (Proposition 3). In two-person games with perfect feedback (hence with observable payoffs) symMSCE is realization-equivalent to symBSCE and Nash equilibrium (Proposition 4). In games with own-strategy independence of feedback and observable payoffs symMSCE coincides with Bayesian SCE and Nash equilibrium (Proposition 5).

In games with separable feedback, we can compare SCE with PSP-equilibrium. Since symMSCE and PSP-equilibrium coincide, PSP-equilibrium refines symBSCE (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).

In the rest of this section, we further discuss the relationship between information feedback and PSPs (Section 7.1), and the related literature (Section 7.2).

7.1 Information feedback and partially specified probabilities

Games with PSPs as defined by Lehrer [18, 2012] presume a sort of separability of feedback, a strong assumption that may or may not hold in our analysis. One may be 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 $(F_{i,j} : S_j \rightarrow \mathbb{R})_{j \neq i}$ 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 $\hat{F}_{i,j}(\sigma_j)$ ($j \neq i$).

¹²See the discussion of observable payoffs in BCMM.

Furthermore, the idea of modeling feedback with a system of PSPs 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.¹³ Removing separability from the (generalized) definition of PSP-equilibrium makes it a refinement of symBSCE and different from Nash equilibrium. But if payoffs are observable, PSP-equilibrium is equivalent to symBSCE. On the other hand, we have argued that proving existence is problematic only when payoffs are not observable. Indeed, under observable payoffs, symMSCE/PSP-equilibrium is a coarsening of Nash equilibrium, hence existence is trivially satisfied. But we can provide a proof of existence that does not rely on observable payoffs only if we go back to the separable feedback case for which PSP-equilibrium was originally defined.

7.2 Related literature

We refer to BCMM (Section 7.5) 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 [3, 1987] and the selfconfirming equilibrium concept of Fudenberg and Levine [11, 1993] 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.¹⁴ We think that a notion of equilibrium whereby players best respond to confirmed beliefs should have the same name whatever the assumptions about feedback. Therefore, in our terminology, we replaced “conjectural equilibrium” with the more self-explanatory “selfconfirming equilibrium”.

To our knowledge, Lehrer [18, 2012] provides the first definition of a concept akin to SCE where agents are not ambiguity neutral.¹⁵ As shown above, his definition implicitly requires

¹³For 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. $F_i(m) \neq m$).

¹⁴Battigalli [3, 1987] is written in Italian. Battigalli and Guaitoli [7, 1988] is the first work in English with a definition of conjectural equilibrium. Fudenberg and Levine [11, 1993] developed the selfconfirming equilibrium concept independently. Besides the different assumptions about information feedback, Battigalli [3, 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 [6, 2012], the annotated extended abstract of Battigalli [3, 1987].

¹⁵See also Lehrer and Teper [19, 2011].

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 [2, 2009], and it is also emphasized in BCMM and Fudenberg and Kamada [13, 2011].

Battigalli [3, 1987] and Fudenberg and Kamada [13, 2011] 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 of other results of BCMM hold under this assumption.

Acknowledgements We thank Veronica Roberta Cappelli, Alejandro Francetich and Paolo Leonetti 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.

References

- [1] ALIPRANTIS C.D. AND K. BORDER (2006): *Infinite dimensional analysis*, 3rd ed. Berlin: Springer Verlag.
- [2] AZRIELI Y. (2009): "On pure conjectural equilibrium with non-manipulable information," *International Journal of Game Theory*, 38, 209-219.
- [3] BATTIGALLI P. (1987): *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali* (thesis), typescript, Università Bocconi, Milano, 1987.
- [4] BATTIGALLI P. (1994): "Structural Consistency and Strategic Independence in Extensive Games," *Ricerche Economiche*, 48, 357-376.
- [5] BATTIGALLI P. (1997): "Games with Observable Deviators," in P. Battigalli, A. Montesano and F. Panunzi (Eds.) *Decisions, Games and Markets*. Dordrecht: Kluwer Academic Publishers.
- [6] BATTIGALLI P. (2012): "Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. Annotated extended abstract and comments," typescript, Bocconi University.
- [7] BATTIGALLI P. AND D. GUAITOLI (1988): "Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information," *Quaderni di Ricerca* 1988-6, I.E.P., Università Bocconi (published in *Decisions, Games and Markets*, Kluwer, Dordrecht, 97-124, 1997).

- [8] BATTIGALLI P., M. GILLI AND M.C. MOLINARI (1992): “Learning and Convergence to Equilibrium in Repeated Strategic Interaction,” *Research in Economics (Ricerche Economiche)*, 46, 335-378.
- [9] BATTIGALLI P., S. CERREIA-VIOGLIO, F. MACCHERONI AND M. MARINACCI (2011): “Self-Confirming Equilibrium and Model Uncertainty,” IGIER working paper 428, Bocconi University.
- [10] FUDENBERG D. AND D. KREPS (1995): “Learning in extensive games, I: self-confirming equilibria,” *Games and Economic Behavior*, 8, 20-55.
- [11] FUDENBERG D. AND D.K. LEVINE (1993a): “Self-confirming equilibrium,” *Econometrica*, 61, 523-545.
- [12] FUDENBERG D. AND D.K. LEVINE (1993b): “Steady state learning and Nash equilibrium,” *Econometrica*, 61, 547-573.
- [13] FUDENBERG D. AND Y. KAMADA (2011): “Rationalizable partition-confirmed equilibrium,” mimeo, Harvard University.
- [14] GILBOA I. AND M. MARINACCI (2013): “Ambiguity and the Bayesian paradigm,” in D. Acemoglu, M. Arellano, and E. Dekel (eds.) *Advances in Economics and Econometrics: Theory and Applications (Vol. 1)*, Cambridge: Cambridge University Press.
- [15] GILBOA I. AND D. SCHMEIDLER (1989): “Maxmin expected utility with a non-unique prior”, *Journal of Mathematical Economics*, 18, 141-153.
- [16] GILLI M. (1999): “Adaptive Learning in Imperfect Monitoring Games,” *Review of Economic Dynamics*, 2, 472-485.
- [17] KUHN H.W. (1953): “Extensive games and the problem of information,” in: H.W. Kuhn and A.W. Tucker (Eds.), *Contributions to the Theory of Games II*, Princeton University Press, Princeton NJ, 193-216.
- [18] LEHRER E. (2012): “Partially-specified probabilities: decisions and games,” *American Economic Journals, Microeconomics*, 4, 70-100.
- [19] LEHRER E. AND R. TEPER (2011): “Justifiable preferences,” *Journal of Economic Theory*, 146, 762-774.
- [20] SION M. (1958): “On general minimax theorems,” *Pacific Journal of Mathematics*, 171-176.
- [21] WALD A.(1950): *Statistical Decision Functions*. New York: Wiley.
- [22] WEIBULL J. W. (1996): “The mass action interpretation” in H. Kuhn et al “The work of John Nash in game theory”, *Journal of Economic Theory*, 69, 153-185.