Selfconfirming Equilibrium and Model Uncertainty*

P. Battigalli S. Cerreia-Vioglio F. Maccheroni M. Marinacci†

July 18, 2013

*Part of this research was done while the first author was visiting the Stern School of Business of New York University, which he thanks for its hospitality. We thank Nicodemo De Vito, Ignacio Esponda, Eduardo Feingold, Faruk Gul, Johannes Hörner, Yuchiro Kamada, Margaret Meyer, Sujoy Mukerji, Wolfgang Pesendorfer and Bruno Strulovici for useful discussions, as well as seminar audiences at CISEI Capri, D-TEA 2012, Games 2012, RUD 2012, as well as Chicago, Duke, Georgetown, Gothemburg, Milano-Bicocca, MIT, Napoli, NYU, Oxford, Penn, Princeton, UC Davis, UCLA, UCSD, and Yale. 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.

†Battigalli: Bocconi University, IGIER and Department of Decision Sciences, Via Roentgen, 1, Milano (MI), 20136 (pierpaolo.battigalli@unibocconi.it); Cerreia-Vioglio: Bocconi University, IGIER and Department of Decision Sciences, Via Roentgen, 1, Milano (MI), 20136 (simone.cerreia@unibocconi.it); Maccheroni: Bocconi University, IGIER and Department of Decision Sciences, Via Roentgen, 1, Milano (MI), 20136 (fabio.maccheroni@unibocconi.it); Marinacci: Bocconi University, IGIER and Department of Decision Sciences, Via Roentgen, 1, Milano (MI), 20136 (massimo.marinacci@unibocconi.it).

We analyze a notion of selfconfirming equilibrium with non-neutral ambiguity attitudes that generalizes the traditional concept. We show that the set of equilibria expands as ambiguity aversion increases. The intuition is quite simple: by playing the same strategy in a stationary environment an agent learns the implied distribution of payoffs, but alternative strategies yield payoffs with unknown distributions; increased aversion to ambiguity makes such strategies less appealing. In sum, a kind of "status quo bias" emerges: in the long run, the uncertainty related to tested strategies disappears, but the uncertainty implied by the untested ones does not.

KEYWORDS: Selfconfirming equilibrium, conjectural equilibrium, model uncertainty, smooth ambiguity.

JEL CLASSIFICATION: C72, D80.

Chi lascia la via vecchia per la via nuova,
sa quel che perde ma non sa quel che trova[1]

1 Introduction

In a situation of *model uncertainty*, or *ambiguity*, the decision maker does not know the probabilistic model for the variables affecting the consequences of choices. Such uncertainty is inherent to situations of strategic interaction. This is quite obvious when such situations have been faced only a few times. In this paper, we argue that uncertainty is pervasive also in games played recurrently where agents have had the opportunity to collect a large set of observations and the system has settled into a steady state. Such a situation is captured by the selfconfirming equilibrium concept (also called conjectural equilibrium). In a *selfconfirming equilibrium* (henceforth, SCE) agents best respond to confirmed probabilistic beliefs, where "confirmed" means that their beliefs are consistent with the evidence they can collect, given the strategies they adopt. Of course, this evidence depends on how everybody else plays. We analyze SCE and model uncertainty jointly and show that they are conceptually complementary: the SCE conditions endogenously determine the extent of uncertainty, and uncertainty aversion induces a kind of status quo bias that expands the set of selfconfirming patterns of behavior.

The SCE concept can be framed within different scenarios. A simple scenario is just a repeated game with a fixed set of players. In this context, the constituent game, which is being repeated, may have sequential moves and monitoring may be imperfect. To avoid repeated game effects, it is assumed that players do not value their future payoffs, but simply best respond to their updated beliefs about the current period strategies of the opponents. Here instead we refer to a scenario that is more appropriate for the ideas we want to explore: there is a large society of individuals who play recurrently a given game G, possibly a sequential game with chance moves: for each player/role

[1]Italian proverb "Those who leave the old road for a new one, know what they leave but do not what they will find."

i in G (male or female, buyer or seller, etc.) there is a large population of agents who play in role i. Agents are drawn at random and matched to play G. Then, they are separated and re-matched to play G with (almost certainly) different co-players, and so on. After each play of a game in which he was involved, an agent obtains some evidence on how the game was played. The accumulated evidence is the data set used by the agent to evaluate the outcome distribution associated with each choice. Note, there is an intrinsic limitation to the evidence that an agent can obtain: if the game has sequential moves, at most he can observe the terminal node reached, but often he can observe even less, e.g., only his monetary payoffs. However, what each agent is really interested about is the statistical distribution of strategies in the populations corresponding to opponents' roles, as such distributions determine (*via* random matching) the objective probabilities of different strategy profiles of the opponents with whom he is matched. Typically, this distribution is not uniquely identified by long-run frequencies of observations. This defines the fundamental inference problem faced by an agent, and explains why model uncertainty is pervasive also in steady states. The key difference between SCE and Nash equilibrium is that, in an SCE, agents may have incorrect beliefs because many possible underlying distributions are consistent with the empirical frequencies they observe (see Battigalli and Guaitoli 1988, Fudenberg and Levine 1993a, Fudenberg and Kreps 1995).

According to the traditional SCE concept, agents are Bayesian subjective expected utility (SEU) maximizers. Nevertheless, when a set of underlying distributions is consistent with their information, agents face a condition of model uncertainty, or ambiguity, rather than risk; in such a condition SEU maximization amounts to *ambiguity neutrality.* Several models of choice capturing more general attitudes toward model uncertainty have been studied in decision theory (see Gilboa and Marinacci, 2013). The decision theoretic work which is more germane to our approach distinguishes between objective and subjective uncertainty. Given a set S of states, there is a set $\Sigma \subseteq \Delta(S)$ of possible probabilistic "models."[2] Each model $\sigma \in \Sigma$ specifies the objective

[2]In this context, we call "objective probabilities" the possible probability models (distri-

probabilities of states and, for each action a of the decision maker (DM), it determines a von Neumann-Morgenstern expected utility evaluation $U(a, \sigma)$; the DM is uncertain about the true model σ (see Cerreia-Vioglio *et al.*, 2013a,b). In our framework, a is the action, or strategy, of an agent playing in role i, σ is a distribution of strategies in the population of opponents (or a profile of such distributions in n-person games), and Σ is the set of distributions consistent with the database of the agent. Roughly, an agent is uncertainty averse if he dislikes the uncertainty about $U(a, \sigma)$ implied by uncertainty about the true probability model $\sigma \in \Sigma$. We interchangeably refer to such feature of preferences with the expression "aversion to model uncertainty" or the shorter "ambiguity aversion." For example, an extreme form of ambiguity aversion is captured by the *Maxmin criterion* $\max_a \min_{\sigma \in \Sigma} U(a, \sigma)$ of Gilboa and Schmeidler (1989), in its Waldean interpretation.[3] In this paper we span a large set of ambiguity attitudes using the "smooth ambiguity" model of Klibanoff, Marinacci and Mukerji (2005, henceforth KMM). This latter criterion admits the Maxmin criterion as a limit case and the Bayesian SEU criterion as a special case. In an SCE, agents in each role best respond to their database choosing actions with the highest value, and their database is the one that obtains under the true data generating process corresponding to the actual strategy distributions. The following example shows how this notion of SCE differs from the traditional, or Bayesian, SCE.

butions) over a state space S. These are not to be confused with the objective probabilities stemming from an Anscombe and Aumann setting. For a discussion, see Cerreia-Vioglio et al (2013b).

[3] See Cerreia-Vioglio et al (2013a).

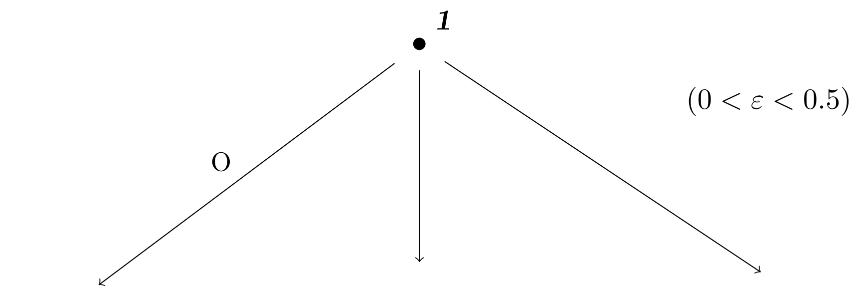

MP^1	h^1	t^1
H^1	2	1
T^1	1	2

MP^2	h^2	t^2
H^2	4	0
T^2	0	4

Figure 1: Matching Pennies with increasing stakes

In the zero-sum game[4] of Figure 1, the first player chooses between an outside option O and two Matching-Pennies subgames, say MP^1 and MP^2. Subgame MP^2 has "higher stakes" than MP^1: it has a higher (mixed) maxmin value ($2 > 1.5$), but a lower minimum payoff ($0 < 1$). In this game there is only one Bayesian SCE outcome,[5] which must be the unique Nash outcome: MP^2 is reached with probability 1 and half of the agents in each population play Head. But we argue informally that moderate aversion to uncertainty makes the low-stakes subgame MP^1 reachable, and high aversion to uncertainty makes the outside option O also possible.[6] Specifically, let $\bar{p}^k$ denote the subjective probability assigned by an ambiguity neutral agent in role 1 to h^k, with $k = 1, 2$. Going to the low-stake subgame MP^1 has subjective value $\max\{\bar{p}^1 + 1, 2 - \bar{p}^1)\} \geq 1.5$ and going to the high-stakes subgame MP^2 has subjective value $\max\{4\bar{p}^2, 4(1 - \bar{p}^2)\} \geq 2$. Thus, O is never an ambiguity neutral best reply and cannot be played by a positive fraction of agents in a Bayesian SCE. Furthermore, also the low-stakes subgame MP^1 cannot be played in a Bayesian SCE. For suppose by way of contradiction that a positive fraction of agents in population 1 played MP^1. In the long run, each one of these agents, and

[4]The zero-sum feature simplifies the example but is inessential.

[5]We call "outcome" a distribution on terminal nodes.

[6]See Section 4 for a rigorous analysis.

all agents in population 2, would learn the relative frequencies of Head and Tail. Since in an SCE agents best respond to confirmed beliefs, the relative frequencies of Head and Tail should be the same in equilibrium, i.e., the agents in population 1 playing MP^1 would learn that its objective expected utility is $1.5 < 2$ and would deviate to MP^2 to maximize their SEU. On the other hand, for agents who are (at least) moderately averse to model uncertainty and keep playing MP^1, having learned the risks involved with the low-stakes subgame confers to reduced-form[7] strategies H^1 and T^1 a kind of "status quo advantage": the objective expected utility of the untried strategies H^2 and T^2 is unknown and therefore they are penalized. Thus, the low-stakes subgame MP^1 can be played by a positive fraction of agents if they are sufficiently averse to model uncertainty. Finally, also the outside option O can be played by a positive fraction of agents in an SCE if they are extremely averse to model uncertainty, as represented by the maxmin criterion. If an agent keeps playing O, he cannot learn anything about the opponents' strategy distribution, hence he deems possible every distribution, or model, σ_2. Therefore, the minimum expected utility of H^1 (resp. T^1) is 1 and the minimum expected utility of H^2 (resp. T^2) is zero, justifying O as a maxmin best reply.[8]

The example shows that, by combining the SCE and ambiguity aversion ideas, a kind of "status quo bias" emerges: in the long run, uncertainty about the expected utility of tested strategies disappears, but uncertainty about the expected utility of the untested ones does not. Therefore, ambiguity averse agents have weaker incentives to deviate than ambiguity neutral agents. More generally, higher ambiguity aversion implies a weaker incentive to deviate from an equilibrium strategy. This explains the main result of the paper: the set of SCE's expands as ambiguity aversion increases. We make this precise by adopting the "smooth ambiguity" model of KMM, which conveniently separates the endogenous subjective beliefs about the true strategy distribution

[7] H^k (resp. T^k) corresponds to the class of realization-equivalent strategies that choose subgame MP^k and then select H^k (resp. T^k).

[8] Note that we are excluding the possibility of mixing through randomization, an issue addressed in Section 5.

from the exogenous ambiguity attitudes, so that the latter can be partially ordered by an intuitive "more ambiguity averse than" relation. With this, we provide a definition of "Smooth SCE" whereby agents "smooth best respond" to beliefs about strategy distributions consistent with their long-run frequencies of observations. The traditional SCE concept is obtained when agents are ambiguity neutral, while a Maxmin SCE concept obtains as a limit case when agents are infinitely ambiguity averse. By our comparative statics result, these equilibrium concepts are intuitively nested from finer to coarser: each Bayesian SCE is also a Smooth SCE, which in turn is also Maxmin SCE.

The rest of the paper is structured as follows. Section 2 gives the setup and our definition of SCE. In Section 3, the core of the paper, we present a comparative statics result and analyze the relationships between equilibrium concepts. Section 4 illustrates our concepts and results with a detailed analysis of a generalized version of the game of Figure 1. Section 5 concludes the paper with a discussion of some important theoretical issues and of the related literature. In the main text we provide informal intuitions for our results. All proofs are collected in the Appendix.

2 Recurrent games and selfconfirming equilibrium

2.1 Preliminaries

Given any measurable space $(X, \mathcal{X})$, we denote by $\Delta(X)$ the collection of all probability measures $\nu : \mathcal{X} \to [0,1]$.[9] When X is finite, say with cardinality k, we assume that $\mathcal{X} = 2^X$ and we identify $\Delta(X)$ with the simplex of $\mathbb{R}^k$.

When the simplex is interpreted as a set of distributions, we consider the measurable space $(\Delta(X), \mathcal{B})$, where $\mathcal{B} = \mathcal{B}(\Delta(X))$ is the Borel sigma-algebra that $\Delta(X)$ inherits as the simplex of $\mathbb{R}^k$. The resulting set of probability

[9] Among them, δ_x denotes the Dirac measure concentrated on $x \in X$, that is, $\delta_x(A) = 1$ if $x \in A$ and $\delta_x(A) = 0$ if $x \notin A$.

measures $\Delta(\Delta(X))$ can be interpreted as the set of subjective beliefs about distributions on X, e.g. the compositions of an urn. Finally, we also endow any measurable subset Σ of $\Delta(X)$ with the relative sigma-algebra inherited from $\Delta(X)$, and we denote by $\Delta(\Sigma)$ the collection of all probability measures defined on such sigma-algebra.

Given $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$, a pair of measurable spaces, we endow the Cartesian product $X \times Y$ with the product sigma-algebra $\mathcal{X} \otimes \mathcal{Y}$. We denote by $\Delta(X) \times \Delta(Y)$ the collection of all *product probability measures*. Moreover, each measurable function $\varphi : X \to Y$ induces the pushforward map $\hat{\varphi} : \Delta(X) \to \Delta(Y)$ defined by

$$\hat{\varphi}(\nu) = \nu \circ \varphi^{-1} \qquad \forall \nu \in \Delta(X).$$

In other words, $\hat{\varphi}(\nu)$ is the distribution of φ with respect to ν, given by the formula $\hat{\varphi}(\nu)(E) = \nu(\varphi^{-1}(E))$ for all $E \in \mathcal{Y}$, sometimes denoted $\nu_{\varphi}(E)$.

2.2 Games with feedback and ambiguity

We consider a finite game played recurrently between agents drawn at random from large populations, one population for each player role. The game may be dynamic, but in this case we assume that the agents play its strategic form, i.e. they simultaneously and irreversibly choose a pure strategy, which is then mechanically implemented by some device. We first describe the key features of rules of the game, then the personal features of the agents playing the game.

The rules of the game determine a *game form with feedback* $(I, (S_i, M_i, F_i)_{i \in I})$ where

- $I = \{1, ...n\}$ is the set of player roles and we call "player i" the agent who in a given instance of the game plays in role $i \in I$;
- S_i is the finite set of strategies of $i \in I$; with this, we let $S = \times_{i \in I} S_i$ and $S_{-i} = \times_{j \neq i} S_j$ respectively denote the set of strategy profiles of i's

co-players' strategy profiles;[10]

- M_i is a set of messages that player i may receive *ex post* and $F_i : S \to M_i$ is an onto feedback function.

For each player role $i \in I$ there is a corresponding population of agents. Agents playing in different roles are drawn at random, hence independently, from the corresponding populations, which do not overlap. Once the game is played by the agents matched at random, the resulting strategy profile s generates a message $m_i = F_i(s)$ for each player $i \in I$. This message encodes all the information about play that player i receives ex post. This information includes, but needs not be limited to the material consequences of interaction observed by i, such as his consumption. To keep our notation simple, we do not make such consequences explicit. If the game is dynamic, a player's feedback is a function of the terminal node $\zeta(s) \in Z$ reached under strategy profile s. In this case $F_i(s) = f_i(\zeta(s))$ where $f_i : Z \to M_i$ is the extensive-form feedback function for player i (see Battigalli *et al.* 2011 for details).

We assume informally that player i knows the feedback function F_i and hence infers from m_i that the strategy profile played must belong to the set of pre-images $\{s' \in S : F_i(s') = m_i\}$. We also informally assume that player i remembers the strategy s_i he just played. Such ex post recall is crucial in our analysis, but we do not have to code it into function F_i. An easy example of feedback function to keep in mind is the following: at the end of the game each player obtains a monetary payoff and player i observes only his own gain. In this case $F_i : S \to M_i \subseteq \mathbb{R}$ specifies the monetary payoff of player i.

[10]To ease notation, we write $j \neq i$ instead of $j \in I \backslash \{i\}$.

$\boldsymbol{F}_1$	$\boldsymbol{H}_2$	$\boldsymbol{T}_2$
$\boldsymbol{H}_1$	1	-1
$\boldsymbol{T}_1$	-1	1

Figure 2

Example 1 *By assumption, in the game Matching Pennies of Figure 2 the feedback function of player 1 specifies at least his monetary gain. Even if this is the only feedback he obtains, he can nonetheless infer the opponent's choice. Indeed, if he wins, he infers that the strategy profile is either* (H_1, H_2) *or* (T_1, T_2)*. On top of this he remembers what he did, hence he further infers that the opponent must have played the matching action. A similar argument holds if he loses.*

Formally, an agent who just played s_i and observes m_i infers that the strategy profile of the co-players belongs to the set

$$F_{i,s_i}^{-1}(m_i) = \{s'_{-i} \in S_{-i} : F_{i,s_i}(s'_{-i}) = m_i\},$$

where $F_{i,s_i} : S_{-i} \to M_i$ denotes the section at s_i of feedback function F_i.[11] To streamline notation, we will write F_{s_i} instead of F_{i,s_i}.[12] Thus, every strategy s_i gives rise to an ex post information partition of the set of co-players' strategy profiles:

$$\mathcal{F}_{s_i} = \{F_{s_i}^{-1}(m_i)\}_{m_i \in M_i}.$$

We say that the strategic game form $(I, (S_i, M_i, F_i)_{i \in I})$ satisfies *own-strategy independence of feedback* if the ex post information partition $\mathcal{F}_{s_i}$ is independent of s_i for every $i \in I$.[13]

This property is very strong and is violated in many interesting cases. For example, the property fails whenever the strategic game form is derived from a non trivial extensive game form where agents infer ex post the terminal node reached, such as the games discussed in the Introduction.

[11] That is, $F_{s_i}(s'_{-i}) = F_i(s_i, s'_{-i})$ for every $s'_{-i} \in S_{-i}$.

[12] Furthermore, we will see that the case in which F_i does not depend on i is important.

[13] This property is called "non manipulability of information" by Battigalli *et al.* (1992) and Azrieli (2009), and "own-strategy independence" by Fudenberg and Kamada (2011).

Example 2 *In the game of Figure 1, assuming that player 1 observes only his monetary payoff, the ex post information partition depends on s_1 as follows:*[14]

$$\begin{aligned} \mathcal{F}_O &= \{S_2\}, \\ \mathcal{F}_{H^1} &= \mathcal{F}_{T^1} = \{\{h^1.h^2, h^1.t^2\}, \{t^1.h^2, t^1.t^2\}\}, \\ \mathcal{F}_{H^2} &= \mathcal{F}_{T^2} = \{\{h^1.h^2, t^1.h^2\}, \{h^1.t^2, t^1.t^2\}\}, \end{aligned}$$

where $a^1.a^2$ denotes the strategy of player 2 that chooses action $a^1 \in \{h^1, t^1\}$ (respectively $a^2 \in \{h^2, t^2\}$) in subgame MP^1 (respectively MP^2).

Next we describe agents' personal features. We assume for notational simplicity that all agents in any given population i have the same attitudes toward risk and the same attitudes toward uncertainty (or ambiguity). The former are represented by a von Neumann-Morgenstern (vNM) utility function

$$v_i : \times_{j \in I} M_j \to \mathbb{R},$$

which gives rise to the strategic-form payoff function

$$U_i(s) = v_i(F_1(s), ..., F_n(s)) \qquad \forall s \in S.$$

With this, we obtain a *game with feedback*

$$G = \left(I, (S_i, M_i, F_i, v_i)_{i \in I}\right).$$

To understand the functional form of the vNM utility functions, let us focus on the case where each M_j is a subset in the real line $\mathbb{R}$ and $m_j = F_j(s)$ is the consumption, or the monetary payoff of player j. In principle, we allow player i to care also about the consumption of any co-player j. But, when agents are selfish,[15] each v_i depends only on m_i, and so $U_i(s)$ only depends on $F_i(s)$. The dependence of $U_i(s)$ solely on $F_i(s)$ holds also when all players

[14] We are coalescing realization-equivalent strategies of player 1.

[15] In this case v_i is strictly increasing and captures risk attitudes in the traditional sense.

receive the same message, that is, $M_i = M$ and $F_i = F$ for all $i \in I$.[16] In general, we say that game G has *observable payoffs* whenever, for each $i \in I$ and all $s \in S$, $U_i(s)$ only depends on $F_i(s)$, that is, the payoff of every player only depends on all the information about play that he receives ex post. Our main results rely on this assumption, that can be formalized as follows.

For each $i \in I$, there exists $u_i : M_i \to \mathbb{R}$ such that

$$\begin{array}{rccc} U_i = u_i \circ F_i : & S & \to & \mathbb{R} \\ & s & \mapsto & u_i\left(F_i\left(s\right)\right). \end{array} \tag{1}$$

For each $i \in I$, the attitudes toward uncertainty (or ambiguity attitudes) of agents in population i are represented by a strictly increasing function $\phi_i : \mathbb{U}_i \to \mathbb{R}$, where $\mathbb{U}_i := [\min_{s\in S} U_i(s), \max_{s\in S} U_i(s)]$ denotes the convex hull of the payoff values. Suppose that player i is uncertain about the true distribution $\sigma_{-i} \in \Delta(S_{-i})$ of strategies in the population of potential co-players[17] and that his uncertainty is expressed by some prior belief $p_i \in \Delta(\Delta(S_{-i}))$, then the value to player i of playing strategy $s_i \in S_i$ is given by the KMM smooth ambiguity criterion:

$$V_i\left(s_i, p_i; \phi_i\right) = \phi_i^{-1}\left(\int_{\Delta(S_{-i})} \phi_i(U_i\left(s_i, \sigma_{-i}\right))p_i(\mathrm{d}\sigma_{-i})\right), \tag{2}$$

where

$$U_i(s_i, \sigma_{-i}) = \sum_{s_{-i}\in S_{-i}} U_i(s_i, s_{-i})\sigma_{-i}(s_{-i})$$

is the objective expected utility of s_i under σ_{-i}. With this, we obtain a *game with feedback and ambiguity attitudes*

$$(G, \phi) = \left(I, (S_i, M_i, F_i, v_i, \phi_i)_{i\in I}\right).$$

[16] For example, G is the strategic form of a game with sequential moves and players observe the terminal node induced by the strategy profile: $M_i = Z$ and $F_i = \zeta$ for each $i \in I$.

[17] In games with three or more players i is facing a profile of strategy distributions $(\sigma_j)_{j\neq i} \in \times_{j\neq i}\Delta(S_j)$. The random matching structure implies that the objective probability of strategy profile s_{-i} is $\sigma_{-i}(s_{-i}) = \prod_{j\neq i}\sigma_j(s_j)$. Thus, $\sigma_{-i} \in \Delta(S_{-i})$ is the product distribution.

On top of the smooth ambiguity criterion, we also consider a special case and a limit case: the standard Bayesian SEU criterion

$$V_i\left(s_i, p_i\right) = V_i\left(s_i, p_i; \mathrm{Id}_{\mathbb{U}_i}\right) = \int_{\Delta(S_{-i})} U_i\left(s_i, \sigma_{-i}\right) p_i(\mathrm{d}\sigma_{-i}), \tag{3}$$

and the (Waldean) Maxmin criterion

$$W_i(s_i, \Sigma_{-i}) = \min_{\sigma_{-i} \in \Sigma_{-i}} U_i(s_i, \sigma_{-i}), \tag{4}$$

where $\Sigma_{-i} \subseteq \Delta(S_{-i})$ is a given restricted set of distributions. If $\Sigma_{-i} =$suppp_i (or, more generally, if the two subsets have the same convex hull) the Maxmin criterion (4) can be obtained as a limit of the KMM criterion (2) when the measure of ambiguity aversion $-\phi_i''/\phi_i'$ converges pointwise to infinity (see KMM for details).

As a matter of interpretation, we have to ascribe to the agents some knowledge of (G, ϕ). In order to make sense of the following analysis, we informally assume that each agent in population i knows S, the random matching structure, and his own feedback, utility and weighting functions F_i, U_i (that is, u_i in the case of observable payoffs) and ϕ_i. No additional knowledge, mutual knowledge or common knowledge of (G, ϕ) is required. In Section 5 we comment extensively on the limitations and possible extensions of our framework. In particular, we discuss our restriction to pure strategies.

Next we describe how an agent who keeps playing a fixed strategy in a stationary environment can partially identify the co-players' strategy distributions. For each strategy s_i, the section at s_i of i's feedback function, F_{s_i}, induces the pushforward map $\hat{F}_{s_i} : \otimes_{j \neq i} \Delta(S_j) \to \Delta(M_i)$, where $\hat{F}_{s_i}(\sigma_{-i})\,(m_i) = \sigma_{-i}((F_{s_i}^{-1}(m_i))$ is the probability that i observes message m_i given the strategy distribution σ_{-i} of his co-players.

If i plays the pure strategy s_i and observes the long-run frequency distribution of messages $\mu_i \in \Delta(M_i)$, then i can compute the set of (product) strategy distributions of the opponents that may have generated μ_i given s_i:

$$\left\{\sigma_{-i} \in \otimes_{j\neq i}\Delta(S_j) : \hat{F}_{s_i}(\sigma_{-i}) = \mu_i\right\}.$$

If $\sigma^*_{-i} = \times_{j\neq i}\sigma^*_j$ is the true strategy distribution of his co-players, the long-run frequency distribution of messages observed by i when playing s_i is the one induced by the objective distribution σ^*_{-i}, that is, $\mu_i = \hat{F}_{s_i}\left(\sigma^*_{-i}\right)$. The set of possible distributions from i's perspective is thus

$$\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i}) = \left\{\sigma_{-i} \in \otimes_{j\neq i}\Delta(S_j) : \hat{F}_{s_i}(\sigma_{-i}) = \hat{F}_{s_i}\left(\sigma^*_{-i}\right)\right\}.$$

The identification correspondence $\hat{\Sigma}_{-i}(s_i, \cdot)$ is nonempty and compact valued; it is also convex valued in two-person games. Our definition of $\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$ reflects the informal assumption that each agent in population i knows he is matched at random with agents from other populations, and hence that – conditional on the true profile of strategy distributions – the strategy played by the agent drawn from population j is independent of the strategy played by the agent drawn from population k. Therefore $\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$ need not be convex in games with three or more players.[18]

In this long-run perspective, if payoffs are observable, then the expected utility function $U_i(s_i, \cdot)$ is constant over $\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$, because the frequency distribution of messages $\hat{F}_{s_i}\left(\sigma^*_{-i}\right)$ identifies the objective probabilities of payoffs associated with strategy s_i:

Lemma 3 *Under observable payoffs, for each s_i and σ^*_{-i},*

$$U_i(s_i, \sigma_{-i}) = U\left(s_i, \sigma^*_{-i}\right) \qquad \forall \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i}).$$

Our interpretation of this result is that, in a stationary environment, there is a kind of status-quo bias in favor of the strategy that has been played for a long time, because this strategy is associated with a known objective expected

[18] If we assumed total ignorance about the matching process, then the partially identified set would be convex, as in the two person case: $\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i}) = \left\{\sigma_{-i} \in \Delta(S_{-i}) : \hat{F}_{s_i}(\sigma_{-i}) = \hat{F}_{s_i}\left(\sigma^*_{-i}\right)\right\}$.

utility. In other words, $\{U_i(s_i, \sigma_{-i}) : \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})\}$ is a singleton, but the set $\{U_i(s'_i, \sigma_{-i}) : \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})\}$ of expected utility levels associated with an untested strategy $s'_i \neq s_i$ may be large.

2.3 Selfconfirming equilibrium

Next we give our definition of selfconfirming equilibrium with non-neutral attitudes toward uncertainty. Recall that we restrict agents to choose pure strategies, so that "mixed" strategies arise only as distributions of pure strategies within populations of agents.

Definition 4 *A profile of strategy distributions* $\sigma^* = (\sigma^*_i)_{i \in I}$ *is a* smooth selfconfirming equilibrium (SSCE) *of a game with feedback and ambiguity attitudes* (G, ϕ) *if, for each* $i \in I$ *and each* $s^*_i \in supp \sigma^*_i$*, there is a prior* $p_{s^*_i} \in \Delta(\hat{\Sigma}_{-i}(s^*_i, \sigma^*_{-i}))$ *such that*

$$V_i\left(s^*_i, p_{s^*_i}; \phi_i\right) \geq V_i\left(s_i, p_{s^*_i}; \phi_i\right) \qquad \forall s_i \in S_i. \tag{5}$$

This "confirmed rationality" condition can be written as

$$s^*_i \in \arg\max_{s_i \in S_i} \int_{\hat{\Sigma}_{-i}(s^*_i, \sigma^*_{-i})} \phi_i\left(U_i\left(s_i, \sigma_{-i}\right)\right) p_{s^*_i}(\mathrm{d}\sigma_{-i}).$$

It requires that every pure strategy s^*_i that a positive fraction $\sigma^*_i(s^*_i)$ of agents keeps playing must be a best response within S_i to the "evidence," that is, the statistical distribution of messages $\hat{F}_{s^*_i}(\sigma^*_{-i}) \in \Delta(M_i)$ generated by playing s^*_i against the strategy distribution σ^*_{-i}.

Looking at the limit case of the Maxmin criterion and the special case of the SEU criterion, we get two ancillary definitions of selfconfirming equilibrium for a game with feedback. A profile of distributions $\sigma^* = (\sigma^*_i)_{i \in I}$ is a:

(i) *Maxmin selfconfirming equilibrium (MSCE)* if, for each $i \in I$ and each

$s_i^* \in supp\sigma_i^*$,

$$s_i^* \in \arg\max_{s_i \in S_i} \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i\left(s_i, \sigma_{-i}\right); \tag{6}$$

(ii) *Bayesian selfconfirming equilibrium (BSCE)* if, for each $i \in I$ and each $s_i^* \in supp\sigma_i^*$, there is a prior $p_{s_i^*} \in \Delta(\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$ such that

$$s_i^* \in \arg\max_{s_i \in S_i} \int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i\left(s_i, \sigma_{-i}\right) p_{s_i^*}(\mathrm{d}\sigma_{-i}); \tag{7}$$

In section 4, we illustrate these definitions with a detailed analysis of a generalized version of the game of Figure 1. Here we consider a more symmetric example.

$\mathbf{1 \backslash 2}$	$\boldsymbol{O}_2$	$\boldsymbol{H}_2$	$\boldsymbol{T}_2$
$\boldsymbol{O}_1$	$1,1$	$1,2$	$1,2$
$\boldsymbol{H}_1$	$2,1$	$4,0$	$0,4$
$\boldsymbol{T}_1$	$2,1$	$0,4$	$4,0$

Figure 3

Example 5 *Figure 3 gives the reduced strategic form of a sequential game where players unilaterally and simultaneously decide either to stop and get out (O_i) or continue. If they both stop, they get 1 util each; if only one of them does, the player who stops gets 1 util, the other player gets 2 utils; if they both continue, next they play a Matching Pennies subgame. Suppose that each i* only observes his own payoff, *that is,* $F_i(\cdot) = U_i(\cdot)$. *Then, an agent who stops cannot observe anything, while an agent who plays Head or Tails identifies the strategy distribution of the population of opponents:*

$$\hat{\Sigma}_{-i}(O_i, \sigma_{-i}) = \Delta(S_{-i}) \text{ and } \hat{\Sigma}_{-i}(H_i, \sigma_{-i}) = \hat{\Sigma}_{-i}(T_i, \sigma_{-i}) = \{\sigma_{-i}\}$$

for every $i \in \{1, 2\}$ *and* $\sigma_{-i} \in \Delta(S_{-i})$. *A necessary condition for* σ^* *to be an*

SCE is

$$\sigma_i^*(O_i\ ) < 1 \Longrightarrow \sigma_{-i}^*(H_{-i}) = \sigma_{-i}^*(T_{-i}), \quad \forall i \in \{1, 2\},$$

because agents who do not stop identify the opponents' distribution and have to be indifferent between Head and Tail. Next note that stopping is never a best response for an ambiguity neutral agent.

With this, it is easy to check that BSCE and NE coincide: nobody stops and the two populations split evenly between Heads and Tails. But the set of SSCE's is much larger if agents are sufficiently ambiguity averse. Specifically, it can be shown that the belief that minimizes the incentive for an ambiguity averse agent to deviate from O_i is $p_i = \frac{1}{2}\delta_{H_{-i}} + \frac{1}{2}\delta_{T_{-i}}$. That is, agents with such belief think that either all agents in population $-i$ play Head, or all of them play Tail, and that these two extreme distributions are equally likely (see Lemma 15 in the Appendix). Let $\phi_i(U) = U^{1/\alpha}$ with $\alpha > 0$ for each i. Then,

$$V_i(H_i, p_i; \phi_i) = V_i(T_i, p_i; \phi_i) = \left(\frac{1}{2}4^{1/\alpha} + \frac{1}{2}0^{1/\alpha}\right)^\alpha \leq 1 \Longleftrightarrow \alpha \geq 2.$$

Therefore, if $\alpha < 2$ then O_i cannot be a best reply to any prior, and so $SSCE = BSCE = NE$; if $\alpha \geq 2$ then O_i is a best reply to p_i, which is trivially confirmed, and the necessary condition for SCE is also sufficient:

$$SSCE = \{\sigma^* : \forall i \in \{1, 2\}, \sigma_i^*(O_i\ ) < 1 \Longrightarrow \sigma_{-i}^*(H_{-i}) = \sigma_{-i}^*(T_{-i})\}$$

$$= \{\sigma^* : \forall i \in \{1, 2\}, \sigma_{-i}^*(H_{-i})(1 - \sigma_i^*(O_i\ )) = \sigma_{-i}^*(T_{-i})(1 - \sigma_i^*(O_i\ ))\}.$$

We conclude that if agents are sufficiently ambiguity averse, i.e. $\alpha \geq 2$, then they may stop in an SSCE. ▲

As pointed out in the discussion (Section 5), our definition of Bayesian SCE subsumes earlier definitions of conjectural and selfconfirming equilibrium as special cases. Like these earlier notions of SCE, our more general notion is motivated by a partial identification problem: the mapping from strategy distributions to the distributions of observations available to an agent is not one to one. In fact, if for each agent i identification is full – i.e., $\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) = \{\sigma_{-i}\}$

for all s_i and all σ_{-i} – condition (5) is easily seen to reduce to the standard Nash equilibrium condition $U_i\left(s_i^*, \sigma_{-i}^*\right) \geq U_i\left(s_i, \sigma_{-i}^*\right)$. In other words, if none of the agents features a partial identification problem, we are back to the Nash equilibrium notion (in its mass action interpretation).

3 Comparative statics and relationships

In this section we compare the equilibria of games with different ambiguity attitudes. This allows us to nest the different notions of SCE defined above. We also identify a special case where they all collapse to mixed Nash equilibrium.

3.1 Main result

Under the smooth ambiguity criterion ambiguity attitudes are characterized by the weighting function ϕ_i. We say that ϕ_i is *more ambiguity averse* than $\bar{\phi}_i$ if there is a concave and strictly increasing function $\varphi_i : Im\phi_i \to \mathbb{R}$ such that $\bar{\phi}_i = \varphi_i \circ \phi_i$ (see KMM). Game (G, ϕ) is *more ambiguity averse* than $(G, \bar{\phi})$ if, for each i, ϕ_i is more ambiguity averse than $\bar{\phi}_i$. Game (G, ϕ) is *ambiguity averse* if it is more ambiguity averse than $(G, \mathrm{Id}_{\mathbb{U}_1}, ..., \mathrm{Id}_{\mathbb{U}_n})$, i.e., if each function ϕ_i is concave.

Note that the comparison of ambiguity attitudes does not require that the profiles themselves be ambiguity averse. It only matters that one profile be comparatively more ambiguity averse than the other (something that can happen even if both are ambiguity loving).

We can now turn to the main result of this paper, which relies on the status quo bias of Lemma 3: the set of equilibria expands when ambiguity aversion increases.

Theorem 6 *Suppose that the game with feedback G has observable payoffs. If (G, ϕ) is more ambiguity averse than $\left(G, \bar{\phi}\right)$, then the SSCE's of $\left(G, \bar{\phi}\right)$ are also SSCE's of (G, ϕ). Similarly, the SSCE's of any game with feedback and ambiguity attitudes (G, ϕ) are also MSCE's of G.*

We provided intuition for this result in the Introduction. Now we can be more precise: let σ^* be an SSCE of $(G, \bar{\phi})$, the *less* ambiguity averse game, and pick any strategy played by a positive fraction of agents, $s_i^* \in \mathrm{supp}\sigma_i^*$; then, there is a justifying confirmed belief $p_{s_i^*} \in \Delta(\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$ such that s_i^* is a best reply to $p_{s_i^*}$ given $\bar{\phi}_i$, i.e., $V_i(s_i^*, p_{s_i^*}; \bar{\phi}_i) \geq V_i(s_i, p_{s_i^*}; \bar{\phi}_i)$ for each s_i. We interpret $p_{s_i^*}$ as the belief held in the long run by an agent who keeps playing the "status-quo" strategy s_i^* in the i.i.d. environment determined by σ_{-i}^*. Such agent eventually learns the long-run frequencies of the (observable) payoffs of s_i^*; therefore, the value of s_i^* for this agent converges to its objective expected utility, $U(s_i^*, \sigma_{-i}^*)$, which independent of his ambiguity attitudes (cf. Lemma 3). But the value of an untested strategy $s_i \neq s_i^*$ typically depends on ambiguity attitudes and, keeping beliefs fixed, it is higher when ambiguity aversion is lower, that is, $V_i(s_i, p_{s_i^*}; \bar{\phi}_i) \geq V_i(s_i, p_{s_i^*}; \phi_i)$. Therefore $V_i(s_i^*, p_{s_i^*}; \phi_i) = U(s_i^*, \sigma_{-i}^*) \geq V_i(s_i, p_{s_i^*}; \phi_i)$ for all s_i. This means that it is possible to support σ^* as an SSCE of the more ambiguity averse game (G, ϕ) using the same profile of beliefs supporting σ^* as an SSCE of $(G, \bar{\phi})$.

3.2 Relationships

Theorem 6 implies that under observable payoffs:

(i) the set of BSCE's of G is contained in the set of SSCE's of every (G, ϕ) with ambiguity averse players;

(ii) the set of SSCE's of every (G, ϕ) is contained in the set of MSCE's of G.

In other words, under observable payoffs and ambiguity aversion, it holds

$$BSCE \subseteq SSCE \subseteq MSCE. \tag{8}$$

The degree of ambiguity aversion determines the size of the set of selfconfirming equilibria, with the sets of Bayesian and Maxmin selfconfirming equilibria being, respectively, the smallest and largest one.

It is well known that every Nash equilibrium σ^* is also a Bayesian SCE. The same relationship holds more generally for Nash and smooth selfconfirming equilibrium (recall that BSCE is the special case of SSCE neutral attitudes toward model uncertainty). Intuitively, a Nash equilibrium is an SSCE with correct (hence confirmed) beliefs about strategy distributions; since correct beliefs cannot exhibit any model uncertainty, they satisfy the equilibrium conditions independently of ambiguity attitudes.

Lemma 7 *If a profile of distributions σ^* is a Nash equilibrium of G, then it is an SSCE of any game with feedback and ambiguity attitudes (G, ϕ).*

Since the set NE of Nash equilibria is nonempty, we automatically obtain existence of SSCE for any f and ϕ. In particular, we can enrich (8) as follows:

$$\emptyset \neq NE \subseteq BSCE \subseteq SSCE \subseteq MSCE.$$

This shows that, under observable payoffs, every game with feedback G has some MSCE and that every Nash equilibrium is not only an SSCE, but also an MSCE.[19]

The next simple, but instructive result establishes a partial converse. Recall that G has own-strategy independent feedback if what each player can infer ex post about the strategies of other players is independent of his own choice. The following proposition illustrates the strength of this assumption.

Proposition 8 *In every game with observable payoffs and own-strategy independent feedback, every type of SCE is equivalent to Nash equilibrium:*

$$NE = BSCE = SSCE = MSCE.$$

The intuition for this result is quite simple: the strategic-form payoff function $U_i(s_i, \cdot) : S_{-i} \to \mathbb{R}$ is constant on each cell $F_{s_i}^{-1}(m_i)$ of the partition

[19] Battigalli *et al.* (2012) prove existence of MSCE under alternative assumptions about information feedback.

$\mathcal{F}_{s_i} = \{F_{s_i}^{-1}(m_i)\}_{m_i \in M_i}$ (observability of payoffs), but this partition is independent of s_i (own-strategy independence of feedback). This means that, in the long run, an agent does not only learn the objective probabilities of the payoffs associated with his "status quo" strategy, but also the objective probabilities of the payoffs associated with every other strategy. Hence, model uncertainty is irrelevant and he learns to play the best response to the true strategy distributions of the other players/roles even if he does not exactly learn these distributions.[20]

Further results about the relationship between equilibrium concepts can be obtained when G is derived from a game in extensive form under specific assumptions about the information structure (see Battigalli *et al.*, 2011).

We conclude emphasizing the key role played by payoff observability in establishing the inclusions (8). The following example shows that, indeed, such inclusions need not hold when payoffs are not observable.

Example 9 *Consider the zero-sum game of Figure 1 of the Introduction, but now suppose that player* 1 *cannot observe his payoff ex post (he only remembers his actions). For example, the utility values in Figure 1 could be a negative affine transformation of the consumption of player* 2*, reflecting a psychological preference of player* 1 *for decreasing the consumption of player* 2 *(not observed by* 1*) even if the consumption of* 1 *is independent of the actions taken in this game. Then, even if* 1 *plays one of the Matching Pennies subgames for a long time, he gets no feedback: under this violation of the observable payoff assumption* $\hat{\Sigma}_2(s_1, \sigma_2) = \Delta(S_2)$ *for all* (s_1, σ_2)*. Since* $u_1(O) = 1 + \varepsilon$ *is larger than the pure maxmin payoff of each subgame, the outside option* O *is the only MSCE choice of player* 1 *at the root. If* ϕ_1 *is sufficiently concave,* O *is also an SSCE choice. But, as already explained,* O *cannot be an ambiguity neutral best reply. Furthermore, it can be verified that every strategy* s_1 *is an SSCE strategy. Therefore,*

$$BSCE \cap MSCE = \emptyset \quad \textit{and} \quad SSCE \not\subseteq MSCE$$

[20] Related results are part of the folklore on SCE. See, for example, Battigalli (1999) and Fudenberg and Kamada (2011).

and so the inclusions (8) here do not hold. ▲

4 A parametrized example

In this section we analyze the SCE's of a zero-sum example parametrized by the number of strategies. The zero-sum assumption is inessential, but it simplifies the structure of the equilibrium set. The game is related to the Matching Pennies example of the Introduction. We show how the SSCE set gradually expands from the BSCE set to the MSCE set as the degree of ambiguity aversion increases.

To help intuition, we first consider a generalization of the game of Figure 1: player 1 chooses between an outside option O and n Matching-Pennies subgames against player 2. Subgames with a higher index k have "higher stakes," that is, a higher (mixed) maxmin value, but a lower minimum payoff (see Figure 5). The game of Figure 1 obtains for $n = 2$.

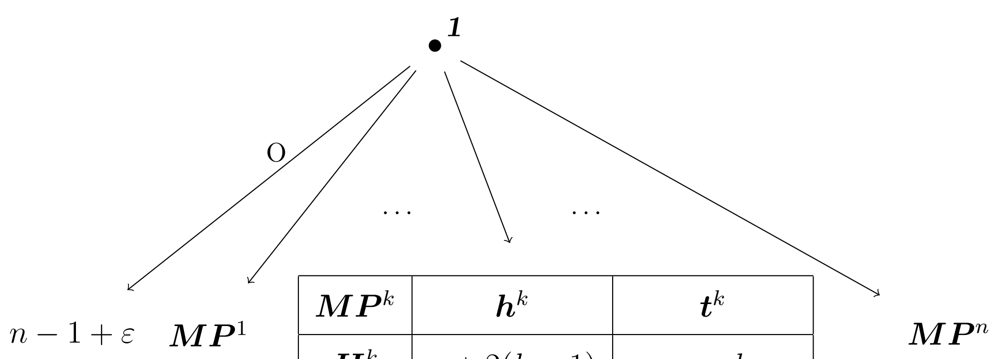

$\boldsymbol{MP}^k$	$\boldsymbol{h}^k$	$\boldsymbol{t}^k$
$\boldsymbol{H}^k$	$n+2(k-1)$	$n-k$
$\boldsymbol{T}^k$	$n-k$	$n+2(k-1)$

Figure 4: Fragment of zero-sum game

In this game, player 1 has $(n+1)\times 2^n$ strategies and player 2 has 2^n strategies. To simplify the notation, we instead analyze an equivalent extensive-form game Γ_n obtained by two transformations. First, player 2 is replaced by a team of opponents $2.1, ..., 2.n$, one for each subgame k, each one with the same payoff function $u_{2.k} = -u_1$. Second, the sequence of moves (k, H^k) of player 1 (go to subgame k then choose Head) – which is common to 2^{n-1} realization-equivalent strategies – is coalesced into the single strategy H^k. Similarly (k, T^k) becomes T^k. The new strategy set of player 1 has $2n+1$ strategies: $S_1 = \{O, H^1, T^1, ..., H^n, T^n\}$. If player 1 chooses H^k or T^k, player $2.k$ moves at information set $\{H^k, T^k\}$ (i.e., without knowing which of the two actions was chosen by player 1) and chooses between h^k and t^k; hence $S_{2.k} = \{h^k, t^k\}$. See Figure 5.

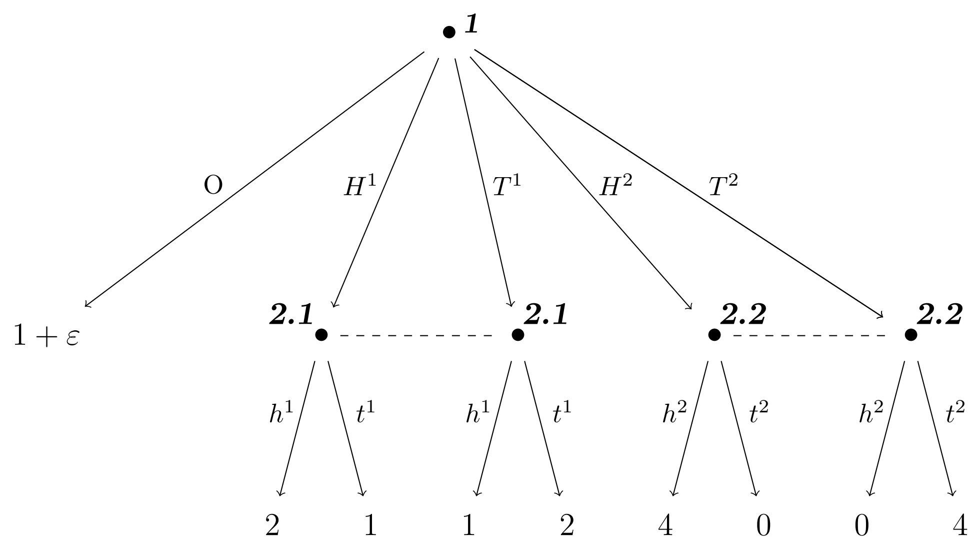

Figure 5: The case $n = 2$

We assume that players observe the terminal node, or – equivalently – that the game has observable payoffs (cf. Example 2).

Although there are no proper subgames in Γ_n, we slightly abuse language and informally refer to "subgame k" when player 1 chooses H^k or T^k, giving the move to opponent $2.k$. The game Γ_n and the previously described game have isomorphic sets of terminal nodes (with cardinality $4n + 1$) and the same reduced normal form (once players $2.1, ..., 2.n$ of the second game are coalesced into a unique player 2). By standard arguments, these two games have equivalent sets of Nash equilibria, equivalent BSCE and MSCE sets, and equivalent

SSCE sets for every ϕ.[21]

That said, consider the game with feedback G_n derived from extensive-form game Γ_n (Figure 5 illustrates the special case Γ_2) under the assumption that the terminal node reached is observed ex post (or that payoffs are observable). It is easily seen that, for every profile of strategy distributions $\sigma_2^* = (\sigma_{2.k}^*)_{k=1}^n$, it holds[22]

$$\hat{\Sigma}_2(O, \sigma_2^*) = \otimes_{k=1}^n \Delta(S_{2.k}), \tag{9}$$

and

$$\hat{\Sigma}_2(H^k, \sigma_2^*) = \hat{\Sigma}_2(T^k, \sigma_2^*) = \{\sigma_2 : \sigma_{2.k} = \sigma_{2.k}^*\}. \tag{10}$$

As a result, next we provide necessary SCE conditions that partially characterize the equilibrium strategy distribution for player/role 1 and fully characterize the equilibrium strategy distributions for the opponents.

Lemma 10 *For every (Bayesian, Smooth, Maxmin) SCE σ^* and every $k = 1, ..., n$,*

$$\sigma_1^*(H^k) + \sigma_1^*(T^k) > 0 \Rightarrow \frac{\sigma_1^*(H^k)}{\sigma_1^*(H^k) + \sigma_1^*(T^k)} = \frac{1}{2} = \sigma_{2.k}^*(h^k). \tag{11}$$

Furthermore, for every σ^* and $\bar{\sigma}^*$, if σ^* is a (Bayesian, Smooth, Maxmin) SCE, and suppσ_1^* = supp$\bar{\sigma}_1^*$, then also $\bar{\sigma}^*$ is a (Bayesian, Smooth, Maxmin) SCE.

Note that these necessary conditions do not restrict at all the set of strate-

[21] Each profile $\sigma = (\sigma_1, (\sigma_{2.k})_{i=1}^n)$ of the new n-person game can be mapped to an equivalent profile $(\bar{\sigma}_1, \bar{\sigma}_2)$ of the old two-person game and viceversa while preserving the equilibrium properties. Specifically, $(\sigma_{2.k})_{k=1}^n$ is also a behavioral strategy of player 2 in the two-person game, which corresponds to a realization-equivalent strategy distribution $\bar{\sigma}_2$ for player 2. Similarly, any such distribution $\bar{\sigma}_2$ can be mapped to a realization-equivalent profile $(\sigma_{2.k})_{k=1}^n$. As for σ_1, for each s_1 in the new game, the probability mass $\sigma_1(s_1)$ can be distributed arbitrarily among the pure strategies of the old two-person game that select the corresponding sequence of moves (that is, either (O), or (k, H^k) or (k, T^k)), thus obtaining a realization-equivalent distribution $\bar{\sigma}_1$. In the opposite direction, every $\bar{\sigma}_1$ of the old game yields a unique realization-equivalent σ_1 in the new game, where $\sigma_1(s_1)$ is the $\bar{\sigma}_1$-probability of the set of (realization-equivalent) strategies that select the same sequence of moves as s_1.

[22] For ease of notation, in this section we denote $\hat{\Sigma}_{-1}$ by $\hat{\Sigma}_2$.

gies that can be played in equilibrium: for every $s_1 \in \{O, H^1, T^1, ..., H^n, T^n\}$ there is some distribution profile σ^* such that $\sigma_1^*(s_1) > 0$ and (11) holds. The formal proof of the lemma is straightforward and left to the reader. Intuitively, if subgame k is played with positive probability, then each agent playing this subgame learns the relative frequencies of Head and Tail in the opponent's population, and the best response conditions imply that an SCE reaching subgame k with positive probability must induce a Nash equilibrium in this Matching-Pennies subgame. Thus, the σ_2^*-value to an agent in population 1 of playing the "status quo" strategy H^k or T^k (with $\sigma_1^*(H^k) + \sigma_1^*(T^k) > 0$) is the mixed maxmin value of subgame k, $n - 1 + k/2$. With this, the value of deviating to another "untested" strategy depends on the exogenous attitudes toward model uncertainty, and on the subjective belief $p_1 \in \Delta(\hat{\Sigma}_2(H^k, \sigma_2^*))$, which is only restricted by $\sigma_{2.k}^*$ (eqs. (9) and (10)). As for the agents in roles $2.1, ..., 2.n$, their attitudes toward uncertainty are irrelevant, because, if they play at all, they learn all that matters to them, that is, the relative frequencies of H^k and T^k.

Suppose that a positive fraction of agents in population 1 play H^k or T^k, with $k < n$. By Lemma 10, in an SCE the value that they assign to their strategy is its vNM expected utility given that opponent $2.k$ mixes fifty-fifty, that is, $n-1+k/2$. But, if they are ambiguity neutral, the subjective value of deviating to subgame n is at least the maxmin value $n-1+n/2 > n-1+k/2$. Furthermore, the outside option O is never an ambiguity neutral best reply.[23] This explains the following:

Proposition 11 *The BSCE set of G_n coincides with the set of Nash equilibria. Specifically,*

$$BSCE = NE = \left\{ \sigma^* \in \Sigma : \sigma_1^*(H^n) = \sigma_1^*(T^n) = \sigma_n^*(h^n) = \frac{1}{2} \right\}.$$

Next we analyze the SSCE's assuming that agents are ambiguity averse in the KMM sense. The following preliminary result, which has some independent interest, specifies the beliefs about opponents' strategy distributions

[23] Indeed, O is strictly dominated by every mixed strategy $\frac{1}{2}H^k + \frac{1}{2}T^k$.

that minimize the subjective value of deviating from a given strategy s_1 to any subgame j.

Lemma 12 *Let ϕ_1 be concave. For all $j = 1, ..., n$, $p_1, q_1 \in \Delta\left(\otimes_{k=1}^{n} \Delta(S_{2.k})\right)$, if*

$$\operatorname{mrg}_{\Delta(S_{2.j})} q_1 = \frac{1}{2}\delta_{h^j} + \frac{1}{2}\delta_{t^j},$$

then

$$\max\{V_1(H^j, p_1; \phi_1), V_1(T^j, p_1; \phi_1)\} \geq V_1(H^j, q_1; \phi_1) = V_1(T^j, q_1; \phi_1).$$

Intuitively, an ambiguity averse agent dislikes deviating to subgame j the most when his subjective prior assigns positive weight only to the highest and lowest among the possible objective expected utility values, i.e., when its marginal on $\Delta(S_j)$ has the form $x\delta_{h^j} + (1-x)\delta_{t^j}$. By symmetry of the 2×2 payoff matrix of subgame k, he would pick within $\{H^k, T^k\}$ the strategy corresponding to the highest subjective weight (H^k if $x > 1/2$). Hence, the subjective value of deviating to subgame j is minimized when the two Dirac measures δ_{h^j} and δ_{r^j} have the same weight $x = 1/2$.

To analyze how the SSCE set changes with the degree of ambiguity aversion of player 1, we consider the one-parameter family of negative exponential weighting functions

$$\phi_1^{\alpha}(U) = -e^{-\alpha U},$$

where $\alpha > 0$ is the coefficient of ambiguity aversion (see KMM p. 1865). Let $SSCE(\alpha)$ denote the set of SSCE's of $(\Gamma_n, \mathrm{Id}_Z, ..., \mathrm{Id}_Z, \phi_1^{\alpha}, \phi_2, ...\phi_n)$. To characterize the equilibrium correspondence $\alpha \mapsto SSCE(\alpha)$, we use the following transformation of $\phi_1^{\alpha}(U)$:

$$M(\alpha, x, y) = (\phi_1^{\alpha})^{-1}\left(\frac{1}{2}\phi_1^{\alpha}(x) + \frac{1}{2}\phi_1^{\alpha}(y)\right).$$

By Lemma 12, this is the minimum value of deviating to a subgame characterized by payoffs x and y. The following known result states that this value is decreasing in the coefficient of ambiguity aversion α, it converges to the mixed maxmin value as $\alpha \to 0$ (approximating the ambiguity neutral case), and it converges to the minimum payoff as $\alpha \to +\infty$.

Lemma 13 *For all $x \neq y$, $M(\cdot; x, y)$ is strictly decreasing, continuous, and satisfies*

$$\lim_{\alpha \to 0} M(\alpha; x, y) = \frac{1}{2}x + \frac{1}{2}y \quad \text{and} \quad \lim_{\alpha \to +\infty} M(\alpha; x, y) = \min\{x, y\}. \tag{12}$$

By Lemma 10, to analyze the $SSCE(\alpha)$ correspondence we only have to determine the strategies s_1 that can be played by a positive fraction of agents in equilibrium, or – conversely – the strategies s_1 that must have measure zero. Let us start from very small values of α, i.e., approximately ambiguity neutral agents. By Lemmas 12 and 13, the subjective value of deviating to the highest-stakes subgame n is approximately bounded below by $n - 1 + n/2 > u_1(O)$. Therefore, the outside option O cannot be a best reply. Furthermore, suppose by way of contradiction that H^k or T^k ($k < n$) are played by a positive fraction of agents. By Lemma 10, the value of playing subgame k is the vNM expected utility $n - 1 + k/2 < n - 1 + n/2$. Hence all agents playing this game would deviate to the highest-stakes subgame n. Thus, for α small $SSCE(\alpha) = BSCE$. By Lemma 13, as α increases, the minimum value of deviating to subgame n decreases, converging to zero for $\alpha \to +\infty$. More generally, the minimum value $M(\alpha, n-j, n+2(j-1))$ of deviating to subgame j converges to $n-j$ for $\alpha \to +\infty$. Since $n-j < u_1(O) < n-1+k/2$, this means that, as α increases, it becomes easier to support an arbitrary strategy s_1 as an SSCE strategy. Therefore there must be thresholds $0 < \alpha_1 < ... < \alpha_n$ such that *only* the higher-stakes subgames $k+1, ...n$ can be played by a positive fraction of agents in equilibrium if $\alpha < \alpha_{n-k}$, and *every* strategy (including the outside option O) can be played by a positive fraction of agents for *some* $\alpha \geq \alpha_{n-k}$. In particular, for α sufficiently large, $SSCE(\alpha)$ coincides with the set of Maxmin SCE's, which is just the set

$$\Sigma^* = \{\sigma^* \in \Sigma : \text{eq. (11) holds}\}$$

of distribution profiles satisfying the necessary conditions of Lemma 10.[24] To summarize, by the properties of function $M(\alpha, x, y)$ stated in Lemma 13, we can define strictly positive thresholds $\alpha_1 < \alpha_2 < ... < \alpha_n$ so that the following indifference conditions hold

$$\max_{j\in\{k+1,...,n\}} M(\alpha_{n-k}, n-j, n+2(j-1)) = n-1+\frac{k}{2}, \qquad k = 1, ..., n-1, \tag{13}$$

$$\max_{j\in\{k+1,...,n\}} M(\alpha_n, n-j, n+2(j-1)) = n-1+\varepsilon, \tag{14}$$

and $SSCE(\alpha)$ expands as α increases, making subgame k playable in equilibrium as soon as α reaches α_{n-k}, expanding to $MSCE$ and making the outside option O playable as soon as α reaches α_n. Formally:

Proposition 14 *Let $\alpha_1 < ... < \alpha_n$ be the strictly positive thresholds defined by (13) and (14). For every α and $k = 1, ...n-1$,*

$$\alpha < \alpha_{n-k} \Longrightarrow SSCE(\alpha) = \left\{\sigma^* \in \Sigma^* : \sigma_1^*(\{O, L^1, T^1, ..., H^k, T^k\}) = 0\right\}$$

and

$$\alpha < \alpha_n \Longrightarrow SSCE(\alpha) = \{\sigma^* \in \Sigma^* : \sigma_1^*(O) = 0\}.$$

Furthermore

$$\mathop{}_{\alpha\geq\alpha_{n-k}} SSCE(\alpha) = \Sigma^* = MSCE,$$

and $SSCE(\alpha) = BSCE = NE$ if $\alpha < \alpha_1$, while $SSCE(\alpha) = MSCE$ if $\alpha \geq \alpha_n$.

5 Concluding remarks and related literature

The SCE concept characterizes stable patterns of behavior in games played recurrently. We analyzed a notion of SCE with agents who have non-neutral

[24] This characterization holds for every parametrized family of distributions that satisfies, at every expected utility value $\bar{U}$, properties analogous to those of Lemma 13, with α replaced by the coefficient of ambiguity aversion $-\phi_1''(\bar{U})/\phi'(\bar{U})$.

attitudes toward uncertainty about the true steady-state data generating process. We argued that this uncertainty comes from a partial identification problem: the mapping from strategy distributions to the distributions of observations available to an agent is not one to one. We used as our workhorse the KMM smooth-ambiguity model, which separates endogenous beliefs from exogenous ambiguity attitudes. This makes our setup particularly well suited to connect with the previous literature on SCE and to analyze how the set of equilibria changes with the degree of ambiguity aversion. Assuming observability of payoffs, we show that the set of smooth SCE's expands when agents become more ambiguity averse. The reason is that agents learn the expected utility values of the strategies played in equilibrium, but not of the strategies they can deviate to, which are thus penalized by higher ambiguity aversion. This allows us to derive intuitive relationships between different versions of SCE. Nash equilibrium is a refinement of all of them, which guarantees existence. All notions of SCE collapse to Nash equilibrium under the additional assumption of own-strategy independence of feedback.

We developed our theoretical insights in the framework of population games played recurrently, but similar intuitions apply to different strategic contexts, such as dynamic games with a stationary Markov structure. Our insights are likely to have consequences for more applied work. For example, the SCE and ambiguity aversion ideas have been applied in macroeconomics to analyze, respectively, learning in policy making (see Sargent, 1999, and the references in Cho and Sargent, 2008) and robust control (Hansen and Sargent, 2008). Our analysis suggests that these two approaches can be fruitfully merged. Fershtman and Pakes (2012) put forward a concept of "experienced based equilibrium" akin to SCE to provide a framework for the theoretical and empirical analysis of dynamic oligopolies. They argue that equilibrium conditions are, in principle, testable when agents beliefs are determined (if only partially) by empirical frequencies, as in their equilibrium concept and SCE. Their model features observable payoffs because firms observe profits, therefore a version of our main result applies: ambiguity aversion expands the set of equilibria.

In the remainder of this section we consider some limitations and possible extensions of our analysis and we briefly discuss the related literature. We refer the reader to the working paper version (Battigalli *et al.* 2011) for a more detailed discussion.

Dynamic consistency To avoid dynamic consistency issues, we assumed that agents play the strategic form of the recurrent game. But when agents really play a game with sequential moves, not its strategic form, they cannot commit to any contingent plan. A strategy for an agent is just a plan that allows him to evaluate the likely consequences of taking actions at any information set. The plan is credible and can be implemented only if it prescribes, at each possible information set, an action that has the highest value, given the agent's beliefs and planned continuation. Plans with this *unimprovability* property can be obtained by means of a "folding back" procedure on the subjective decision tree implied by the agent's beliefs. We can make this precise in the context of the smooth-ambiguity model, and thus provide notions of SSCE assuming unimprovability. For brevity, we write SSCE^{δ} to refer to such smooth selfconfirming equilibria with dynamic consistency. What difference does dynamic consistency make?

It is well known that the strategies that an ambiguity averse agent would commit to, if he could, need not be unimprovable.[25] Therefore, SSCE^{δ} is not equivalent to SSCE. Furthermore, it is not obvious whether our main comparative statics result is valid with this definition. We conjecture that the following version of the comparative statics result holds: if $\bar{\sigma}^*$ is an SSCE^{δ} of $(G, \bar{\phi})$ and (G, ϕ) is more ambiguity averse than $(G, \bar{\phi})$, then there is an SSCE^{δ} σ^* of (G, ϕ) that induces the same distribution of observable consequences, that is, $\hat{F}_i(\bar{\sigma}^*) = \hat{F}_i(\sigma^*)$ for each player i.[26]

There is another problem related to dynamic consistency: as in Fudenberg and Levine (1993a), the notion of best reply used in this paper allows for suboptimal behavior at unexpected information sets. To deal with

[25] Siniscalchi (2011) reports examples and provides an in-depth analysis of dynamic consistency under ambiguity aversion. Battigalli *et al.* (2011) provide a game theoretic example.

[26] In Battigalli *et al.* (2011) we show by example that the analog of Theorem 6 cannot hold exactly for SSCE^{δ}.

this problem we can represent players' beliefs as *conditional probability systems.* An agent in role i has an initial joint belief about pure strategies and strategy distributions, $\pi_i \in \Delta(S_{-i} \times \Delta(S_{-i}))$, such that $(\mathrm{mrg}_{S_{-i}}\pi_i)(s_{-i}) = \int_{\Delta(S_{-i})} \sigma_{-i}(s_{-i})\pi_i(\{s_{-i}\} \times \mathrm{d}\sigma_{-i})$. Furthermore, for each information set h_i there is a corresponding conditional belief $\pi_i(\cdot|S_{-i}(h_i) \times \Delta(S_{-i}))$, where $S_{-i}(h_i)$ is the set of s_{-i} consistent with h_i, and all these beliefs are related to each other *via* Bayes rule whenever possible (see Battigalli and Siniscalchi, 1999). With this, we can give stronger versions of unimprovability and consistent planning to obtain a *refined* notion of SSCE$^\delta$. It can be shown that this refinement does not change the set of equilibrium outcomes. The reason is simple: agents are not assumed to know the preferences of others and may have incorrect beliefs about the choices of others at off-equilibrium-path information sets. The refinement has bite when mutual or common knowledge of the game is assumed, as discussed in the following subsection on rationalizable SSCE.

Mixed strategies In our analysis agents' choice is restricted to pure strategies. This means that we do not allow them to commit to arbitrary objective randomization devices. The equilibrium concept obtained by allowing mixed strategies is not a generalization of SSCE (or MSCE). This can be easily seen in the game of Figure 1: if player 1 delegates his choice to an objective randomization device that selects the high-stakes subgame MP^2 with probability one and splits evenly the probability mass on Head and Tail, he guarantees at least 2 utils in expectation. If this randomized choice were available no agent in population 1 would choose the outside option O or the low-stakes subgame MP^1, and the unique SCE outcome would be the Nash outcome. In general, we can define notions of smooth and Maxmin SCE whereby arbitrary randomizations are allowed, and show that the set of Maxmin SCE's is contained in the set of Bayesian SCE's. On the other hand, our result that under observable payoffs $MSCE \subseteq SSCE \subseteq BSCE$ holds also when agents choose mixed strategies. We conclude that if payoffs are observable and agents can commit to delegate their choice of strategy to arbitrary randomization devices, then ambiguity aversion does not affect the set of selfconfirming equilibrium distributions (though, of course, their rationales can be very different).[27]

The reason why we restrict choice to pure strategies is that credible randomization requires a richer commitment technology than assumed so far. This can be seen by focussing on simultaneous-moves games, where playing a pure strategy simply means that an action is irreversibly chosen. But there is a commitment issue in playing mixed strategies. Suppose that an agent in population i believes that mixed strategy σ_i^* is optimal. If this is true for an ambiguity neutral (SEU) agent, then also each action in the support of σ_i^* is optimal, therefore σ_i^* can be implemented by mapping each action in $\mathrm{supp}\sigma_i$ to the realization of an appropriate roulette spin and then choosing the action associated with the observed realization. On the other hand, an ambiguity averse agent who finds σ_i^* optimal, need not find all the actions in $\mathrm{supp}\sigma_i$ optimal (within the simplex $\Delta(S_i)$). Therefore, unlike an ambiguity neutral agent, an ambiguity averse one has to be able to irreversibly delegate his choice to the random device. At the interpretive level, we are not really assuming that agents are prevented from using randomization devices: it may be the case that agents in population i have a set $\hat{S}_i \subset S_i$ of "truly pure" strategies and that S_i also includes a finite set of choices that are realization-equivalent to randomizations over $\hat{S}_i$.[28] If this is the case, such commitment technology should be explicitly allowed by the rules of the game and represented in the game form.

Rationalizable selfconfirming equilibrium In a selfconfirming equilibrium agents are rational and their beliefs are confirmed. If the game is common knowledge, it is interesting to explore the implications of assuming, on top of this, that there is common (probability-one) belief of rationality and confirmation of beliefs. Interestingly, the set of *rationalizable SCE*s thus obtained may be a strict subset of the set of SCE's consistent with common certainty of rationality, which in turn may be a strict subset of the set of

[27] See Section 6 in the working paper version (Battigalli *et al.*, 2011).

[28] Of course, the definition of F_i has to be adapted accordingly, because $F_i(s_i, s_{-i})$ is a random message when s_i is a randomization device.

SCE's.[29] The separation between ambiguity attitudes and beliefs in the KMM smooth-ambiguity model allows a relatively straightforward extension of this idea to obtain a notion of rationalizable SSCE.

Learning and steady states Fudenberg and Levine (1993b) analyze agents' learning in an overlapping generations model of a population game with stationary aggregate distributions. They show that steady-state strategy distributions approach a selfconfirming equilibrium as agents' life-span increases. The intuition is that agents learn and experiment only when they are young; when the life-span is very long, the vast majority of agents has approximately settled beliefs and choose stage-game best responses to such beliefs. The stationarity assumption is a clever trick that allows to use consistency and convergence results in Bayesian statistics about sampling from a "fixed urn" of unknown distribution.

The separation between ambiguity attitudes and beliefs in the KMM model allows to analyze updating in a Bayesian fashion and attempt an extension of this result to SSCE. Our conjecture is that, as the life-span increases, steady-state strategy distributions should approximate a smooth SCE even faster, because ambiguity averse agents stop experimenting sooner than ambiguity neutral ones. This can be more easily understood if agents observe *only* their own payoffs. In this case, choices that are perceived to give raise to uncertain posterior beliefs coincide with those that are perceived as ambiguous, i.e., those that yield uncertain distributions of payoffs. Therefore the choices that are worth experimenting with are exactly those that an ambiguity averse agent tend to avoid.

Related literature As we mentioned, our notion of SCE subsumes earlier definitions due to Battigalli (1987)[30] and Fudenberg and Levine (1993a) as special cases. These earlier definitions assume SEU maximization and apply to games in extensive form with feedback functions $f_i : Z \to M_i$ defined on the

[29]See Rubinstein and Wolinsky (1994), Dekel *et al.* (1999). See also the references to rationalizable SCE in Battigalli *et al.* (2011).

[30]See also Battigalli and Guaitoli (1988).

set of terminal nodes Z. We can fit this in our strategic-form framework letting $F_i(s) = f_i(\zeta(s))$, where $\zeta : S \to Z$ is the outcome function associating strategy profiles with terminal nodes. Battigalli (1987) allows for general feedback functions f_i with observable payoffs, but he considers only equilibria where all agents playing in a given role have the same independent belief about co-players. Fudenberg and Levine (1993a) assume that players observe the terminal node reached (each f_i is one-to-one). Since payoffs are determined by endnodes, this implies that payoffs are observable.

We are not going to thoroughly review the vast literature on uncertainty and ambiguity aversion, which is covered by a comprehensive recent survey (Gilboa and Marinacci, 2013). We only mention that in the paper we rely on the decision theoretic framework of Cerreia-Vioglio *et al.* (2013a,b) that makes formally explicit the DM's uncertainty about the true probabilistic model, or data generating process.

To the best of our knowledge, the paper most related to our idea of combining SCE with non-neutral attitudes toward uncertainty is Lehrer (2012). In this paper a decision maker is endowed with a "partially specified probability" (PSP), that is, a list of random variables defined on a probability space. The decision maker knows only the expected values of the random variables, hence he is uncertain about the true underlying probability measure within the set of all measures that give rise to such values. Lehrer (2012) axiomatizes a decision criterion equivalent to the maximization of the minimal expected utility with respect to the set of probability measures consistent with the PSP. Then he defines a notion of equilibrium with partially specified probabilities for a game played in strategic form. Lehrer's equilibrium is similar the one we obtain in the "Maxmin" case,[31] but his assumptions on information feedback eliminate the "status-quo advantage" of equilibrium strategies. To better compare our approach to Lehrer's first note that, for each i and s_i, we have a PSP: the probability space is (S_{-i}, σ_{-i}), the random variables are the indicator functions of

[31] Lehrer considers mixed strategy equilibria and does not assume a population game scenario. His equilibrium concept should be compared to the version of WSCE where any mixed strategy is allowed, but all agents in a given role play the same strategy (see Battigalli *et al.* 2011, Section 6).

the different messages (ex post observations), and their expectations are the objective probabilities of the messages given by distribution $\hat{F}_{s_i}(\sigma_{-i}) \in \Delta(M)$. However, in our paper this PSP may depend on the chosen strategy s_i. Lehrer instead assumes that the PSP depends only on σ_{-i} not on s_i, that is, he assumes own-strategy independence of feedback (in n-person games he relies on an even stronger assumption of separability of feedback across co-players). As we noticed, when this strong assumption is coupled with the rather natural assumption of observable payoffs, Nash equilibrium obtains. In other words, once the two frameworks are made comparable, our Proposition 8 shows that the intersection between the class of equilibria considered in the present paper (where observability of payoffs is maintained) and those considered by Lehrer (2012) *only consists of Nash equilibria.* Battigalli *et al.* (2012) characterizes MSCE in greater detail according to the properties of information feedback and provides a rigorous analysis of the relationship between MSCE and Lehrer's equilibrium concept.

6 Appendix

6.1 Proof of Lemma 3

Since payoffs are observable, eq. (1) holds and $U_i = u_i \circ F_i$, where $u_i : M_i \to \mathbb{R}$. Fix s_i and σ^*_{-i}. For each $\sigma_{-i} \in \Delta(S_{-i})$,

$$\begin{aligned} U_i(s_i, \sigma_{-i}) &= \sum_{s_i \in S_{-i}} (u_i \circ F_i)(s_i, s_{-i})\, \sigma_{-i}(s_{-i}) \\ &= \sum_{s_{-i} \in S_{-i}} (u_i \circ F_{s_i})(s_{-i})\, \sigma_{-i}(s_{-i}) \\ &= \sum_{m_i \in M_i} u_i(m_i) \sum_{s_{-i} \in F_{s_i}^{-1}(m_i)} \sigma_{-i}(s_{-i}) = \sum_{m_i \in M_i} u_i(m_i)\, \hat{F}_{s_i}(\sigma_{-i})(m_i). \end{aligned}$$

This implies $U_i(s_i, \sigma_{-i}) = U_i\left(s_i, \sigma^*_{-i}\right)$ if $\hat{F}_{s_i}(\sigma) = \hat{F}_{s_i}\left(\sigma^*_{-i}\right)$, that is, if $\sigma_{-i} \in \hat{\Sigma}_{-i}\left(s_i, \sigma^*_{-i}\right)$. ■

6.2 Proofs for Section 3

Proof of Theorem 6 By definition, there exists a profile $(\varphi_i)_{i\in I}$ of strictly increasing and concave functions such that $\phi_i = \varphi_i \circ \bar{\phi}_i$ for each i. Let σ^* be an SSCE of the less ambiguity averse game $\left(G, \bar{\phi}\right)$. Fix $i \in I$, and pick $s_i^* \in supp\sigma_i^*$, $p_{s_i^*} \in \Delta(\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$ such that $s_i^* \in \arg\max_{s_i \in S_i} V_i(s_i, p_{s_i^*}; \bar{\phi}_i)$. We want to show that $s_i^* \in \arg\max_{s_i \in S_i} V_i(s_i, p_{s_i^*}; \phi_i)$, which implies the first claim. Since payoffs are observable, by Lemma 3 $U_i(s_i^*, \sigma_{-i}) = U_i(s_i^*, \sigma_{-i}^*)$ for each $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$. Thus

$$V_i(s_i^*, p_{s_i^*}; \phi_i) = U_i(s_i^*, \sigma_{-i}^*) = V_i(s_i^*, p_{s_i^*}; \bar{\phi}_i). \tag{15}$$

Next observe that, for any $s_i \in S_i$,

$$\begin{aligned}
V_i(s_i, p_{s_i^*}; \bar{\phi}_i) &= \bar{\phi}_i^{-1}\left(\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} \bar{\phi}_i(U_i(s_i, \sigma_{-i}))p_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&= (\bar{\phi}_i^{-1} \circ \varphi_i^{-1}) \circ \varphi_i\left(\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} \bar{\phi}_i(U_i(s_i, \sigma_{-i}))p_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&\geq (\bar{\phi}_i^{-1} \circ \varphi_i^{-1})\left(\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} (\varphi_i \circ \bar{\phi}_i)(U_i(s_i, \sigma_{-i}))p_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&= \phi_i^{-1}\left(\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} \phi_i(U_i(s_i, \sigma_{-i}))p_{s_i^*}(\mathrm{d}\sigma_{-i})\right) = V_i(s_i, p_{s_i^*}; \phi_i)
\end{aligned}$$

where we used Jensen's inequality and $\phi_i = \varphi_i \circ \bar{\phi}_i$. Hence,

$$V_i(s_i, p_{s_i^*}; \phi_i) \leq V_i(s_i, p_{s_i^*}; \bar{\phi}_i) \leq V_i(s_i^*, p_{s_i^*}; \bar{\phi}_i) = V_i(s_i^*, p_{s_i^*}; \phi_i)$$

for each $s_i \in S_i$, which shows that $s_i^* \in \arg\max_{s_i \in S_i} V_i(s_i, p_{s_i^*}; \phi_i)$.

We now prove that SSCE's are MSCE's. Let σ^* be an SSCE of a game (G, ϕ). Fix $i \in I$, and pick $s_i^* \in supp\sigma_i^*$, $p_{s_i^*} \in \Delta(\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$ such that $s_i^* \in \arg\max_{s_i \in S_i} V_i(s_i, p_{s_i^*}; \phi_i)$. We show that $s_i^* \in \arg\max_{s_i \in S_i} W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$

where, by definition,

$$W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U\left(s_i, \sigma_{-i}\right)$$

Since payoffs are observable, Lemma 3 and (15) imply

$$W_i(s_i^*, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) = U_i(s_i^*, \sigma_{-i}^*) = V_i(s_i^*, p_{s_i^*}; \phi_i) \geq V_i(s_i, p_{s_i^*}; \phi_i) \tag{16}$$

for each $s_i \in S_i$. Next observe that for, each $\sigma'_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$ and each $s_i \in S_i$,

$$U_i(s_i, \sigma'_{-i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U\left(s_i, \sigma_{-i}\right) = W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))$$

hence

$$\phi_i\left(U_i(s_i, \sigma'_{-i})\right) \geq \phi_i\left(W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))\right)$$

because ϕ_i is strictly increasing. This implies that, for each $s_i \in S_i$,

$$\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} \phi_i\left(U_i(s_i, \sigma'_{-i})\right) p_{s_i^*}(\mathrm{d}\sigma'_{-i}) \geq \phi_i\left(W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))\right)$$

which in turn implies

$$\begin{aligned} V_i(s_i, p_{s_i^*}; \phi_i) &= \phi_i^{-1}\left(\int_{\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} \phi_i\left(U_i(s_i, \sigma'_{-i})\right) p_{s_i^*}(\mathrm{d}\sigma'_{-i})\right) \\ &\geq \phi_i^{-1}\left(\phi_i\left(W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*))\right)\right) \\ &= W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)). \end{aligned}$$

This latter inequality paired with (16) delivers that

$$W_i(s_i^*, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) \geq W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) \qquad \forall s_i \in S_i$$

proving the statement. ■

Proof of Lemma 7 Fix a mixed strategy Nash equilibrium σ^* of G. Pick any

i and pure strategy $s_i^* \in supp\sigma_i^*$. Then $U_i\left(s_i^*, \sigma_{-i}^*\right) \geq U_i\left(s_i, \sigma_{-i}^*\right)$ for each $s_i \in S_i$. By definition, it holds $\sigma_{-i}^* \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$. Hence, $\delta_{\sigma_{-i}^*} \in \Delta\left(\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)\right)$. Since $V_i\left(s_i, \delta_{\sigma_{-i}^*}, \phi_i\right) = U_i\left(s_i, \sigma_i^*\right)$ for every weighting function ϕ_i and $s_i \in S_i$, it follows that σ^* is an SSCE of (G, ϕ). ■

Proof of Proposition 8 Given the previous results, we only have to show that every MSCE is a Nash equilibrium. Fix an MSCE σ^*, any player i and any $s_i^* \in supp\sigma_i^*$. Then, for each s_i

$$W_i(s_i^*, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) \geq W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)).$$

By Lemma 3, observability of payoffs implies $U_i(s_i, \sigma_{-i}) = U_i(s_i, \sigma_{-i}^*)$ for each s_i and $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma_{-i}^*)$. Own-strategy independence of feedback implies that, for each s_i,

$$\{F_{s_i}^{-1}(m_i) : m_i \in F_{s_i}(S_{-i})\} = \{F_{s_i^*}^{-1}(m_i) : m_i \in F_{s_i^*}(S_{-i})\}$$

hence

$$\begin{aligned} \hat{\Sigma}_{-i}(s_i, \sigma_{-i}^*) &= \left\{\sigma_{-i} : \forall m_i, \sigma_{-i}\left(F_{s_i}^{-1}(m_i)\right) = \sigma_{-i}^*\left(F_{s_i}^{-1}(m_i)\right)\right\} \\ &= \left\{\sigma_{-i} : \forall m_i, \sigma_{-i}\left(F_{s_i^*}^{-1}(m_i)\right) = \sigma_{-i}^*\left(F_{s_i^*}^{-1}(m_i)\right)\right\} = \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*). \end{aligned}$$

From the above equalities and inequalities we obtain, for each s_i,

$$U_i(s_i^*, \sigma_{-i}^*) = W_i(s_i^*, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) \geq W_i(s_i, \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)) = W_i(s_i, \hat{\Sigma}_{-i}(s_i, \sigma_{-i}^*)) = U_i(s_i, \sigma_i^*).$$

This shows that σ^* is a Nash equilibrium. ■

6.3 Proofs for Section 4

Proof of Proposition 11. For any prior p_1, the ambiguity neutral subjective value of playing any Matching Pennies subgame k is

$$\begin{aligned} & \max\{V_1(H^k, p_1), V_1(T^k, p_1)\} \\ = & \max\left\{ \begin{array}{l} \bar{p}_1^k(h^k)(n+2(k-1)) + (1-\bar{p}_1^k(h^k))(n-k), \\ \bar{p}_1^k(h^k))(n-k) + (1-\bar{p}_1^k(h^k))(n+2(k-1)) \end{array} \right\} \\ \geq & \; n-1+\frac{k}{2} > n-1+\varepsilon = u_1(O), \end{aligned}$$

where $n-1+k/2$ is the mixed maxmin value of subgame k, $\bar{p}_1^k = \mathrm{mrg}_{S_{2.k}}\bar{p}_1$ and $\bar{p}_1$ is the predictive belief. Therefore O cannot be played by a positive fraction of agents in a BSCE because it cannot be a best response to any predictive belief $\bar{p}_1$. Furthermore, no strategy H^k or T^k with $k < n$ can have positive measure in a BSCE. Indeed, by (11), if $s_1^k \in \{H^k, T^k\}$ has positive probability in an equilibrium σ^*, then for every belief $p_1 \in \Delta(\hat{\Sigma}_{-1}(s_1^k, \sigma_2^*))$, the value of s_1^k is

$$V_1(s_1^k, p_1) = U_1\left(s_1^k, \sigma^*_{-\{1,2.k\}} \times \left(\frac{1}{2}h^k + \frac{1}{2}t^k\right)\right) = n-1+\frac{k}{2},$$

while the ambiguity neutral value of deviating to subgame n is

$$\max\{V_1(H^n, p_1), V_1(T^n, p_1)\} \geq n-1+\frac{n}{2}.$$

Therefore, eq. (11) implies $\sigma_1^*(H^n) = \sigma_1^*(T^n) = \sigma_n^*(h^n) = \frac{1}{2}$ in each BSCE σ^*. It is routine to verify that every such σ^* is also a Nash equilibrium. Therefore $BSCE = NE$. ■

The proof of Lemma 12 is based on the following lemma, where $\mathbf{I}$ is the unit interval $[0,1]$ endowed with the Borel σ-algebra.

Lemma 15 *Let $\varphi : \mathbf{I} \to \mathbb{R}$ be increasing and concave. For each Borel probability measure p on $\mathbf{I}$*

$$\max\left\{\int_{\mathbf{I}} \varphi(x)\, p(\mathrm{d}x), \int_{\mathbf{I}} \varphi(1-x)\, p(\mathrm{d}x)\right\} \geq \frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0). \qquad (17)$$

Proof. Let

$$\begin{array}{rccc} \tau: & \mathbf{I} & \to & \mathbf{I} \\ & x & \mapsto & 1-x \end{array}.$$

Then

$$\int_{\mathbf{I}} \varphi(1-x)\, p(\mathrm{d}x) = \int_{\mathbf{I}} \varphi(\tau(x))\, p(\mathrm{d}x) = \int_{\mathbf{I}} \varphi(y)\, p_\tau(\mathrm{d}y)$$

where $p_\tau = p \circ \tau^{-1}$. In particular, for $\varphi = \mathrm{id}_{\mathbf{I}}$ it follows that $1 - \int_{\mathbf{I}} xp(\mathrm{d}x) = \int_{\mathbf{I}} yp_\tau(\mathrm{d}y)$. Thus (17) becomes

$$\max\left\{\int_{\mathbf{I}} \varphi(x)\, p(\mathrm{d}x), \int_{\mathbf{I}} \varphi(x)\, p_\tau(\mathrm{d}x)\right\} \geq \frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0)$$

and either $\int_{\mathbf{I}} xp(\mathrm{d}x) \geq 1/2$ or $\int_{\mathbf{I}} yp_\tau(\mathrm{d}y) \geq 1/2$. Next we show that for each Borel probability measure q on $\mathbf{I}$ such that $\int_{\mathbf{I}} xq(\mathrm{d}x) \geq 1/2$

$$\int_{\mathbf{I}} \varphi(x)\, q(\mathrm{d}x) \geq \frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0). \tag{18}$$

Denote by $F(x) = q([0,x])$ and by $G(x) = \left(\frac{1}{2}\delta_0 + \frac{1}{2}\delta_1\right)([0,x])$. In particular, F and G are increasing, right continuous, and such that $F(1) = G(1) = 1$, moreover $G(x) = 1/2$ for all $x \in [0,1)$. Moreover, there exists $\bar{x} \in (0,1)$ such that $F(\bar{x}) \leq 1/2$. By contradiction, assume $F(x) > 1/2$ for all $x \in (0,1)$, then

$$\frac{1}{2} \leq \int_{\mathbf{I}} xq(\mathrm{d}x) = \int_0^1 (1 - F(x))\,\mathrm{d}x < \frac{1}{2}.$$

if $F(0) = q(\{0\}) > 1/2$, then

$$\int_{\mathbf{I}} xq(\mathrm{d}x) \leq 0q(0) + \int_{(0,1]} xq(\mathrm{d}x) \leq 0q(0) + \int_{(0,1]} 1q(\mathrm{d}x) \leq 0q(\{0\}) + (1 - q(\{0\})) < \frac{1}{2}$$

contradicting $\int_{\mathbf{I}} xq(\mathrm{d}x) \geq 1/2$. Let $x^* = \inf\{x \in \mathbf{I} : F(x) > 1/2\}$, then $0 < \bar{x} \leq x^* \leq 1$.

Therefore $F(1) = G(1) = 1$ and for each $y \in (x^*, 1)$, $F(y) \geq F(x^*) \geq 1/2 \geq G(y)$. For each $y \in [0, x^*)$, $F(y) \leq 1/2 \leq G(y)$. Finally, by the classic Karlin-Novikoff (1963) result F second-order stochastically dominates G, that is (18) holds for all increasing and concave φ. ■

Proof of Lemma 12 Let $x = \sigma_{2.k}(h^k)$. Clearly $U_1(H^k, \sigma_2)$ depends only on x and we can write $U_1(H^k, x)$, and similarly for T^k. Let $\varphi(x) = \phi_1(U_1(H^k, x))$. By symmetry of the payoff matrix, $\varphi(1-x) = \phi_1(U_1(T^k, x))$. Note that φ is strictly increasing and concave. Let $p \in \Delta(\mathbf{I})$ be the marginal belief about $x = \sigma_{2.k}(h^k)$ derived from p_1. Recall that q_1 is a prior such that $\mathrm{mrg}_{\Delta(S_{2.j})} q_1 = \frac{1}{2}\delta_{h^j} + \frac{1}{2}\delta_{t^j}$. With this,

$$
\begin{aligned}
& \max\{V_1(H^j, p_1; \phi_1), V_1(T^j, p_1; \phi_1)\} \\
= \; & \max\left\{\phi_1^{-1}\left(\int_{\mathbf{I}} \varphi(x)\, p(\mathrm{d}x)\right), \phi_1^{-1}\left(\int_{\mathbf{I}} \varphi(1-x)\, p(\mathrm{d}x)\right)\right\} \\
= \; & \phi_1^{-1}\left(\max\left\{\int_{\mathbf{I}} \varphi(x)\, p(\mathrm{d}x), \int_{\mathbf{I}} \varphi(1-x)\, p(\mathrm{d}x)\right\}\right)
\end{aligned}
$$

and

$$
V_1(H^j, q_1; \phi_1) = V_1(T^j, q_1; \phi_1) = \phi_1^{-1}\left(\frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0)\right).
$$

Hence, the thesis is implied by Lemma 15. ■

Proof of Proposition 14 By Lemma 10, $SSCE(\alpha)$ is determined by the set of pure strategies of player 1 that can be played by a positive fraction of agents in equilibrium. Fix $\sigma^* \in \Sigma^*$, i.e., a distribution profile that satisfies the necessary SCE conditions, and a strategy s_1; $\sigma_1^*(s_1) > 0$ is possible in equilibrium if and only if there no incentives to deviate to any subgame j. We rely on Lemma 12 to specify a belief $p_1^{s_1} \in \Delta(\hat{\Sigma}_2(s_1, \sigma_2^*))$ that minimizes the incentive to deviate. Thus, s_1 can be played in equilibrium if and only if it is a best reply to $p_1^{s_1}$. Specifically,

$$
p_1^O = \times_{j=1}^n \left(\frac{1}{2}\delta_{h^j} + \frac{1}{2}\delta_{t^j}\right) \in \Delta(\hat{\Sigma}_2(O, \sigma_2^*)) = \Delta\left(\otimes_{j=1}^n \Delta(S_{j.k})\right),
$$

for each $k = 1, ..., n-1$ and $s_1^k \in \{H^k, T^k\}$,

$$
p_1^k = \delta_{\frac{1}{2}h^k + \frac{1}{2}t^k} \times \left(\times_{j \neq k}\left(\frac{1}{2}\delta_{h^j} + \frac{1}{2}\delta_{t^j}\right)\right)
$$

belongs to $\Delta(\hat{\Sigma}_2(s_1^k, \sigma_2^*)) = \Delta\left(\left\{\sigma_2 : \sigma_{2,k} = \frac{1}{2}h^k + \frac{1}{2}t^k\right\}\right)$. Given such beliefs, the value of deviating from s_1 to subgame j is $M(\alpha, n-j, n+2(j-1)$. Therefore, O is a best reply to p_1^O, and can have positive measure in equilibrium, if and only if

$$n - 1 + \varepsilon \geq \max_{j \in \{1,\ldots,n\}} M(\alpha, n-j, n+2(j-1)). \tag{19}$$

By Lemma 13 there is a unique threshold $\alpha_n > 0$ that satisfies (19) as an equality so that (19) holds if and only if $\alpha \geq \alpha_n$. Similarly, $s_1^k \in \{H^k, L^k\}$ $(k = 1, ..., n-1)$ is a best reply to p_1^k, and can have positive measure in equilibrium, if and only if

$$n - 1 + \frac{k}{2} \geq \max_{j \in \{1,\ldots,n\}} M(\alpha, n-j, n+2(j-1)), \tag{20}$$

where

$$\max_{j \in \{1,\ldots,n\}} M(\alpha, n-j, n+2(j-1)) = \max_{j \in \{k+1,\ldots,n\}} M(\alpha, n-j, n+2(j-1))$$

because, for all α and $j \leq k$

$$M(\alpha, n-j, n+2(j-1)) \leq n - 1 + \frac{j}{2} < n - 1 + \frac{k}{2}.$$

By Lemma 13 there is a unique threshold $\alpha_{n-k} > 0$ that satisfies (20) as an equality so that (20) holds if and only if $\alpha \geq \alpha_{n-k}$. Since $M(\cdot, x, y)$ is strictly decreasing if $x \neq y$, the thresholds are strictly ordered: $\alpha_1 < \alpha_2 < ... < \alpha_n$. It follows that, for each $k = 1, ..., n-1$, $\sigma^*(\{O, H^1, T^1, ..., H^k, T^k\}) = 0$ for every $\sigma^* \in SSCE(\alpha)$ if and only if $\alpha < \alpha_{n-k}$, and every strategy has positive measure in some SSCE if α is large enough (in particular if $\alpha \geq \alpha_n$). Since the equilibrium set in this case is Σ^*, which is defined by necessary SCE conditions, this must also be the MSCE set. If $\alpha < \alpha_1$, then $\sigma^*(\{O, H^1, T^1, ..., H^{n-1}, T^{n-1}\}) = 0$ for each $\sigma^* \in SSCE(\alpha)$; by Proposition 11, $SSCE(\alpha) = BSCE = NE$ in this case. ■

References

[1] Anscombe, F.J. and R. Aumann (1963): "A definition of subjective probability," *Annals of Mathematical Statistics*, 34, 199-205.

[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*, unpublished thesis, Università Bocconi.

[4] Battigalli P. (1999): "A comment on non-Nash equilibria," mimeo, European University Institute.

[5] 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).

[6] Battigalli P. and M. Siniscalchi (1999): "Hierarchies of conditional beliefs and interactive epistemology in dynamic games," *Journal of Economic Theory*, 88, 188-230.

[7] Battigalli P., M. Gilli and M.C. Molinari (1992): "Learning and convergence to equilibrium in repeated strategic interaction," *Research in Economics,* 46, 335-378.

[8] Battigalli P., S. Cerreia-Vioglio, F. Maccheroni and M. Marinacci (2011): "Self-confirming equilibrium and model uncertainty," IGIER W.P. 428, Università Bocconi.

[9] Battigalli P., S. Cerreia-Vioglio, F. Maccheroni and M. Marinacci (2012): "Analysis of information feedback and selfconfirming equilibrium," IGIER W.P. 459, Università Bocconi.

[10] Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci and L. Montrucchio (2013a): "Ambiguity and robust statistics," *Journal of Economic Theory*, 148, 974-1049.

[11] Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci and L. Montrucchio (2013b): "Classical subjective expected utility," *Proceedings of the National Academy of Sciences*, 110, 6754-6759.

[12] Cho, I.K. and T.J. Sargent (2008): "Self-confirming equilibrium," in S. N. Durlauf and L. E. Blume (eds.) *The New Palgrave Dictionary of Economics, 2*nd ed., London: Palgrave Macmillan.

[13] Dekel, E., D. Fudenberg and D.K. Levine (1999): "Payoff information and self-confirming equilibrium," *Journal of Economic Theory*, 89, 165-185.

[14] Fershtman, C. and A. Pakes (2012): "Dynamic games with asymmetric information: a framework for empirical work," *Quarterly Journal of Economics*, 127, 1611-1661.

[15] Fudenberg, D. and Y. Kamada (2011): "Rationalizable partition-confirmed equilibrium," mimeo, Harvard University.

[16] Fudenberg, D. and D. Kreps (1995): "Learning in extensive games, I: self-confirming equilibria," *Games and Economic Behavior,* 8, 20-55.

[17] Fudenberg, D. and D.K. Levine (1993a): "Self-confirming equilibrium," *Econometrica*, 61, 523-545.

[18] Fudenberg, D. and D.K. Levine (1993b): "Steady state learning and Nash equilibrium," *Econometrica*, 61, 547-573.

[19] 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*, Cambridge: Cambridge University Press.

[20] Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with a non-unique prior," *Journal of Mathematical Economics*, 18, 141-153.

[21] Hansen, L.P. and T.J. Sargent (2008): *Robustness.* Princeton: Princeton University Press.

[22] Karlin, S. and A. Novikoff (1963): "Generalized convex inequalities," *Pacific Journal of Mathematics*, 13, 1251-1279.

[23] Klibanoff P., M. Marinacci and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849-1892.

[24] Lehrer, E. (2012): "Partially-specified probabilities: decisions and games," *American Economic Journals, Microeconomics*, 4, 70-100.

[25] Rubinstein, A. and A. Wolinsky (1994): "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability," *Games and Economic Behavior,* 6, 299-311.

[26] Sargent, T.J. (1999): *The Conquest of American Inflation.* Princeton: Princeton University Press.

[27] Siniscalchi, M. (2011): "Dynamic choice under ambiguity," *Theoretical Economics*, 6, 379-421.

[28] von Neumann J. and O. Morgenstern (1947): *The Theory of Games and Economic Behavior* (Second ed.). Princeton: Princeton University Press.

[29] Wald, A. (1950): *Statistical Decision Functions.* New York: Wiley.