Selfconfirming Equilibrium and Model Uncertainty*

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

IGIER and Department of Decision Sciences — Università Bocconi

July 29, 2014

*Battigalli: Università Bocconi and IGIER, via Sarfatti 25, 20136 Milano, ITALIA, pierpaolo.battigalli@unibocconi.it
Cerreia-Vioglio: Università Bocconi and IGIER, via Sarfatti 25, 20136 Milano, ITALIA, simone.cerreia@unibocconi.it
Maccheroni: Università Bocconi and IGIER, via Sarfatti 25, 20136 Milano, ITALIA, fabio.maccheroni@unibocconi.it
Marinacci: Università Bocconi and IGIER, via Sarfatti 25, 20136 Milano, ITALIA, massimo.marinacci@unibocconi.it
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 two anonymous referees for suggestions that lead to a significant improvement of the paper. We also thank Veronica Cappelli, Nicodemo De Vito, Ignacio Esponda, Eduardo Feingold, Alejandro Francetich, 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, Lugano, 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. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

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, D81.

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

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 in 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 benchmark scenario is just a repeated game with a fixed set of players in which there are no intertemporal strategic links between the plays. That is, the individuals who play the game many times are concerned only with their instantaneous payoff, and ignore the effects of their current actions on the other players' future behavior; they simply best respond to their updated beliefs about the current period strategies of the opponents. Although all our results apply to this situation, our presentation is framed into the so called large populations (or Nash's mass action) scenario: There is a large society of individuals

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

who play recurrently a given game G, possibly a sequential game with chance moves: for each player/role 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 opponents, 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, he can observe at most the terminal node reached, but often he will observe even less, e.g., only his monetary payoffs (and not those of his opponents). Each agent is interested in the distribution of strategy profiles adopted by the opponents with whom he is matched, because it determines the expected payoffs of his alternative strategies. Typically, this distribution is not uniquely identified by the long-run frequencies of the agent's observations. This defines the fundamental inference problem he faces, and explains why model uncertainty is pervasive also in steady states. The key difference between SCE and Nash equilibrium is that, in a 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).

Partial identification of the true distribution and awareness of the possible incorrectness of beliefs form the natural domain for *ambiguity aversion*. Yet, according to the traditional SCE concept, agents are Bayesian subjective expected utility maximizers and hence *ambiguity neutral*. Here we modify the notion of SCE to allow for non-neutral attitudes toward model uncertainty (see Gilboa and Marinacci, 2013, for a recent review on the topic). 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 that the agent posits.[2] Each model

[2]In this context, we call "objective probabilities" the possible probability models (distributions) over a state space S. These are not to be confused with the objective probabilities

$\sigma \in \Sigma$ specifies the objective probabilities of states and, for each action a of the decision maker, it determines a von Neumann-Morgenstern expected utility evaluation $U(a, \sigma)$; the decision maker 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 the 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."

A now classical description of ambiguity aversion is the maxmin criterion of Gilboa and Schmeidler (1989), where actions are chosen to solve the problem $\max_a \min_\sigma U(a, \sigma)$. 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 subjective expected utility criterion as a special case.

In a "Smooth" SCE, agents in each role best respond to beliefs consistent with their database, choosing actions with the highest smooth-ambiguity value, and their database is the one that obtains under the true data generating process corresponding to the actual strategy distribution. The following example shows how our notion of SCE differs from the traditional, or Bayesian, SCE.

---

stemming from an Anscombe and Aumann setting. For a discussion, see Cerreia-Vioglio *et al.* (2013b).

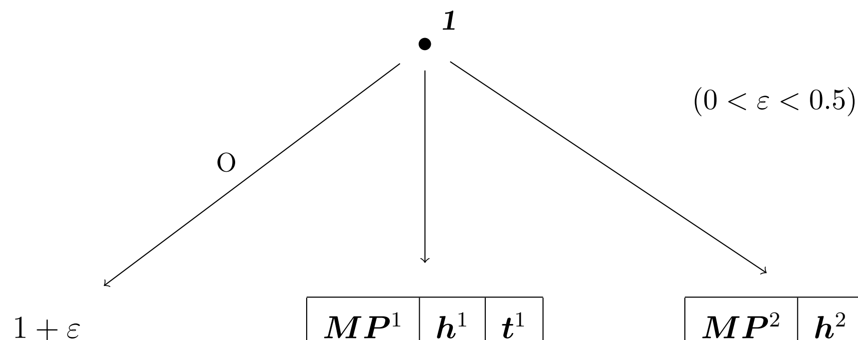

$\boldsymbol{MP}^1$	$\boldsymbol{h}^1$	$\boldsymbol{t}^1$
$\boldsymbol{H}^1$	2	1
$\boldsymbol{T}^1$	1	2

$\boldsymbol{MP}^2$	$\boldsymbol{h}^2$	$\boldsymbol{t}^2$
$\boldsymbol{H}^2$	4	0
$\boldsymbol{T}^2$	0	4

Figure 1: Matching Pennies with increasing stakes

In the zero-sum game[3] 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,[4] 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.[5] Specifically, let $\bar{\mu}^k$ denote the subjective probability assigned by an ambiguity neutral agent in role 1 to h^k, with $k = 1, 2$. Going to the low-stakes subgame MP^1 has subjective value $\max\{\bar{\mu}^1 + 1, 2 - \bar{\mu}^1)\} \geq 1.5$ and going to the high-stakes subgame MP^2 has subjective value $\max\{4\bar{\mu}^2, 4(1 - \bar{\mu}^2)\} \geq 2$. 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,

[3]The zero-sum feature simplifies the example, but it is inessential.
[4]We call "outcome" a distribution on terminal nodes.
[5]See Section III for a rigorous analysis.

and all agents in population 2, would learn the relative frequencies of Head and Tail. Since in a 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[6] 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 a SCE if they are extremely averse to model uncertainty, as represented by the maxmin criterion.[7] 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 sepa-

[6] 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).

[7] As anticipated, the maxmin criterion is a limit case of the smooth one, therefore the same result holds for very high degrees of ambiguity aversion.

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

rates the endogenous subjective beliefs about the true strategy distribution 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. Finally, we show how our results for Smooth SCE extend to other robust decision criteria.

The rest of the paper is structured as follows. Section I gives the setup and our definition of SCE. In Section II, the core of the paper, we present a comparative statics result and analyze the relationships between equilibrium concepts. Section III illustrates our concepts and results with a detailed analysis of a generalized version of the game of Figure 1. Section IV 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.

I. Recurrent games and selfconfirming equilibrium

A. Games with feedback

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; that is, they simultaneously and irreversibly choose a pure strategy, which is then mechanically implemented by some device.

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 $S = \Pi_{i \in I} S_i$ and $S_{-i} = \Pi_{j \neq i} S_j$ denote the set of all strategy profiles and of i's opponents' strategy profiles, respectively;
- M_i is a set of messages that player i may receive *ex post* (at the end of the game);
- $F_i : S \to M_i$ is a 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 typically includes, but needs not be limited to, the material consequences of interaction observed by i, such as his consumption. 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.

Example 1 *Three natural special cases are: For every $i \in I$ and $s \in S$,*

- $F_i(s) = \zeta(s)$*, each player observes the terminal node (reached under the realized strategy profile), that is, f_i is the identity on Z;*
- $F_i(s) = g(\zeta(s))$*, each player observes everybody's material consequences at the terminal node, that is, f_i is the consequence function g;*[9]

[9]The consequence function $g : Z \to \Pi_{i \in I} C_i$ associates profiles of consequences with terminal nodes, where C_i denotes the set of all material consequences that player i may face at the end of the game.

- $F_i(s) = g_i(\zeta(s))$, *each player observes his own material consequences at the terminal node, that is,* f_i *is the* i*-th projection of* g. ▲

Note that, while in the first two cases all agents obtain the same feedback, in the third one feedback is personal. We implicitly assume that player i knows the feedback function F_i and remembers the strategy s_i he just played. Hence, upon playing s_i and receiving message m_i, he infers that the strategy profile played by his opponents must belong to the set

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

where $F_{i,s_i} : S_{-i} \to M_i$ denotes the section at s_i of the feedback function F_i.[10] To streamline notation, and inspired by the important special cases in which $F_i = F$ does not depend on i, we write F_{s_i} instead of F_{i,s_i}. With this, every strategy s_i gives rise to an ex post information partition of the set of opponents' strategy profiles:

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

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:*[11]

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

where $a^1.a^2$ *denotes the strategy of player 2 that chooses action* $a^1 \in \{h^1, t^1\}$ *(resp.* $a^2 \in \{h^2, t^2\}$*) in subgame* MP^1 *(resp.* MP^2*). Summing up,* $\mathcal{F}_{s_1}$ *depends on* s_1 *and it never fully reveals the strategy played by the opponent.* ▲

A game form with feedback $(I, (S_i, M_i, F_i)_{i \in I})$ satisfies *own-strategy independence of feedback* if the ex post information partition $\mathcal{F}_{s_i}$ is independent

[10] That is, $F_{i,s_i}(s_{-i}) = F_i(s_i, s_{-i})$ for every $s_{-i} \in S_{-i}$.

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

of s_i for every $i \in I$.[12] 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 game discussed above.

B. Players' preferences

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, as usual, by a von Neumann-Morgenstern payoff function

$$U_i : S \to \mathbb{R}.$$

We say that game G has *observable payoffs* whenever the payoff of every player only depends on his ex post information about play. Our main results rely on this maintained assumption, which can be formalized as follows: For each $i \in I$, each $s_i \in S_i$, and every $s'_{-i}, s''_{-i} \in S_{-i}$ such that $F_{s_i}\left(s'_{-i}\right) = F_{s_i}\left(s''_{-i}\right)$, we have

$$U_i(s_i, s'_{-i}) = U_i(s_i, s''_{-i}).$$

Contrapositively, this means that, upon playing a fixed strategy and obtaining different utilities, the agent would detect a difference in his opponents' counter strategies (receive a different feedback).[13]

We call *game with feedback* the tuple $G = (I, (S_i, M_i, F_i, U_i)_{i \in I})$ where agents' payoffs are specified.

For each $i \in I$, the attitudes toward uncertainty, or *ambiguity attitudes*, of agents in population i are represented by a strictly increasing and continuous function $\phi_i : \mathbb{U}_i \to \mathbb{R}$, where $\mathbb{U}_i = [\min_{s \in S} U_i(s), \max_{s \in S} U_i(s)]$. Suppose that player i is uncertain about the true distribution $\sigma_{-i} \in \Delta(S_{-i})$ of strategies in

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

[13] Mathematically, this amounts to $\mathcal{F}_{s_i}$-measurability of each section $U_{s_i} = U_{i,s_i}$ of U_i.

the population of potential opponents,[14] and that his uncertainty is expressed by some prior belief μ_i with support on a posited subset Σ_{-i} of $\Delta(S_{-i})$.[15] Then, the value to player i of playing strategy $s_i \in S_i$ is given by the KMM smooth ambiguity criterion:

$$V_i^{\phi_i}(s_i, \mu_i) = \phi_i^{-1}\left(\int_{\mathrm{supp}\mu_i} \phi_i(U_i(s_i, \sigma_{-i}))\mu_i(\mathrm{d}\sigma_{-i})\right), \tag{1}$$

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 von Neumann-Morgenstern expected utility of s_i under σ_{-i}, so that (1) is a certainty equivalent expressed in utils. The standard Bayesian SEU criterion

$$V_i^{\mathrm{id}}(s_i, \mu_i) = \int_{\mathrm{supp}\mu_i} U_i(s_i, \sigma_{-i})\mu_i(\mathrm{d}\sigma_{-i}), \tag{2}$$

corresponds to an affine ϕ_i;[16] while a robust criterion *à la* Gilboa and Schmeidler

$$V_i^{\omega}(s_i, \mu_i) = \min_{\sigma_{-i} \in \mathrm{supp}\mu_i} U_i(s_i, \sigma_{-i}), \tag{3}$$

can be obtained as a limit of (1) when the measure of ambiguity aversion $-\phi_i''/\phi_i'$ converges pointwise to infinity (see KMM for details). Alternative robust preferences are discussed in Section IV.

We call *game with feedback and ambiguity attitudes* a pair (G, ϕ), where G is a game with feedback and $\phi = (\phi_i)_{i \in I}$ is a profile of ambiguity attitudes. We adopt the conventional equality $\phi_i = \omega$ for some or all i in order to encompass preferences represented as in (3). Note that the latter preferences are fully characterized by U_i and the set $\mathrm{supp}\mu_i$ of opponents' strategy distributions

[14] In games with three or more players, i is facing a profile of strategy distributions $(\sigma_j)_{j \neq i} \in \Pi_{j \neq i}\Delta(S_j)$. 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 actually a product distribution.

[15] The simplex $\Delta(S_{-i})$ in $\mathbb{R}^{S_{-i}}$ is endowed with the Borel sigma-algebra.

[16] Since this is by far the most well known functional form, the superscript "id"(which stands for "identity function") will sometimes be omitted.

that agent i deems plausible.

C. Partial identification

Next we describe how an agent who keeps playing a fixed strategy in a stationary environment can partially identify the opponents' strategy distributions, and, if payoffs are observable, he can learn – in the long run – the expected payoff of the fixed strategy itself.

The probability of observing a given message m_i for a player that chooses s_i and faces populations of opponents described by σ_{-i} is

$$\sigma_{-i}\left(\{s_{-i} \in S_{-i} : F_i\left(s_i, s_{-i}\right) = m_i\}\right) = \sigma_{-i}(F_{s_i}^{-1}(m_i)).$$

The corresponding distribution of messages $\sigma_{-i} \circ F_{s_i}^{-1} \in \Delta\left(M_i\right)$ is denoted $\hat{F}_{s_i}\left(\sigma_{-i}\right)$. Therefore, if i plays the pure strategy s_i and observes the long-run frequency distribution of messages $\nu_i \in \Delta(M_i)$, then he can infer that the set of (product) strategy distributions of the opponents that may have generated ν_i is[17]

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

If $\sigma_{-i}^* = \Pi_{j\neq i}\sigma_j^*$ is the true strategy distribution of his opponents, the long-run frequency distribution of messages observed by i when playing s_i is (almost certainly) the one induced by the objective distribution σ_{-i}^*, that is, $\nu_i^* = \hat{F}_{s_i}\left(\sigma_{-i}^*\right)$.[18] The set of possible distributions from i's (long-run empiricist) perspective is thus

$$\begin{aligned}\hat{\Sigma}_{-i}(s_i, \sigma_{-i}^*) &= \left\{\sigma_{-i} \in \Pi_{j\neq i}\Delta(S_j) : \hat{F}_{s_i}\left(\sigma_{-i}\right) = \hat{F}_{s_i}\left(\sigma_{-i}^*\right)\right\} \\ &= \left\{\sigma_{-i} \in \Pi_{j\neq i}\Delta(S_j) : \sigma_{-i|\mathcal{F}_{s_i}} = \sigma_{-i|\mathcal{F}_{s_i}}^*\right\}.\end{aligned}$$

This is, the set of all product probability measures on S_{-i} that coincide with

[17] With a slight abuse of notation we are identifying the product set $\Pi_{j\neq i}\Delta(S_j)$ with the corresponding set of product distributions on S_{-i}.

[18] As common in steady state analysis, we are heuristically relying on a law-of-large-numbers argument.

σ^*_{-i} on the information partition $\mathcal{F}_{s_i}$: Although σ^*_{-i} remains unknown, its restriction to $\mathcal{F}_{s_i}$ is learned in the limit.

The identification correspondence $\hat{\Sigma}_{-i}(s_i, \cdot)$ is nonempty (since $\sigma^*_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$) 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. Hence, he knows 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.[19]

If payoffs are observable, then i can learn their time average and, in the long run, their expectation.[20]

Lemma 1 *If payoffs are observable in the game with feedback G, then, for every i, s_i, and σ^*_{-i},*

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

In contrast, if a different strategy $s'_i \neq s_i$ is considered, the value of $U_i(s'_i, \sigma_{-i})$ as σ_{-i} ranges in $\hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$ remains uncertain: The set $\{U_i(s'_i, \sigma_{-i}) : \sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})\}$ is not, in general, a singleton.[21] This is the feature that, under ambiguity aversion, will generate a kind of status-quo bias in favor of the strategy s_i that has been played for a long time.

As a matter of interpretation, we assume that each agent in population i knows I, S, M_i, F_i, U_i, and ϕ_i, but he may not know F_{-i}, U_{-i}, and ϕ_{-i}. In Section IV we comment extensively on the limitations and possible extensions of our framework.

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

[20] Again, by a law-of-large-numbers heuristic.

[21] Because $U_{s'_i} : S_{-i} \to \mathbb{R}$ is $\mathcal{F}_{s'_i}$-measurable and not, in general, $\mathcal{F}_{s_i}$-measurable.

D. 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 1 *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* $s_i^* \in \mathrm{supp}\sigma_i^*$*, there is a prior* $\mu_{s_i^*}$ *with support contained in* $\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$ *such that*

$$V_i^{\phi_i}\left(s_i^*, \mu_{s_i^*}\right) \geq V_i^{\phi_i}\left(s_i, \mu_{s_i^*}\right) \qquad \forall s_i \in S_i. \tag{4}$$

The "confirmed rationality" condition (4) requires that every pure strategy s_i^* that a positive fraction $\sigma_i^*(s_i^*)$ of agents keep 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}^*.

If all ϕ_i's are affine, we obtain a definition of *Bayesian selfconfirming equilibrium (BSCE)* that subsumes the earlier definitions of conjectural and selfconfirming equilibrium. Finally, we also consider the corresponding classical (as opposed to Bayesian) case of maxmin selfconfirming equilibrium.

Definition 2 *A profile of strategy distributions* $\sigma^* = (\sigma_i^*)_{i \in I}$ *is a* maxmin selfconfirming equilibrium (MSCE) *of a game with feedback* G *if, for each* $i \in I$ *and each* $s_i^* \in \mathrm{supp}\sigma_i^*$*,*

$$\min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i\left(s_i^*, \sigma_{-i}\right) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i\left(s_i, \sigma_{-i}\right) \qquad \forall s_i \in S_i. \tag{5}$$

Formally, this definition is a special case of the previous one. In fact, an MSCE is a SSCE of a game (G, ϕ) with $\phi \equiv \omega$ under the additional assumption that, for each s_i^* played by a positive fraction of agents, the justifying prior $\mu_{s_i^*}$ has full support on $\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$. However, we state it separately since this maxmin notion also admits a conceptually different, classical, statistical interpretation in which priors are absent and so agents are empirical frequentists.

In Section III, 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 2

Example 3 *Figure 2 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 a 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 $\mu_i = \frac{1}{2}\delta_{H_{-i}} + \frac{1}{2}\delta_{T_{-i}}$. 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. Let $\phi_i(U) = U^{1/\alpha}$ with $\alpha > 0$ for each i. Then,*

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

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 μ_i, which is trivially confirmed, and the necessary condition for a 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 a SSCE. ▲

As anticipated above and discussed in Section IV, 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 – that is, $\hat{\Sigma}_{-i}(s_i, \sigma_{-i}) = \{\sigma_{-i}\}$ for all s_i and all σ_{-i} – condition (4) 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).

II. 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 Nash equilibrium.

A. Main result

Ambiguity attitudes are characterized by the weighting functions' profile $\phi = (\phi_i)_{i\in I}$. We say that ϕ_i is *more ambiguity averse* than ψ_i if there is a concave and strictly increasing function $\varphi_i : \psi_i(\mathbb{U}_i) \to \mathbb{R}$ such that $\phi_i = \varphi_i \circ \psi_i$ (see KMM).[22] Game (G, ϕ) is *more ambiguity averse* than (G, ψ) if, for each i, ϕ_i is more ambiguity averse than ψ_i. Game (G, ϕ) is *ambiguity averse* if it is more ambiguity averse than $(G, \mathrm{id}_{\mathbb{U}_1}, ..., \mathrm{id}_{\mathbb{U}_n})$, that is, if each function ϕ_i is concave.

Observe that we do not assume that the ϕ_i's are concave. Therefore, our comparison of ambiguity attitudes does not hinge on this assumption. In other words, for the relation of being more ambiguity averse, it only matters that profile ϕ be comparatively more ambiguity averse than profile ψ, something that can happen even if both are ambiguity loving.

Building on the non ambiguity of the expected payoff of the long-run strategy, established in Lemma 1, we can now turn to the main result of this paper: The set of equilibria expands as ambiguity aversion increases.

Theorem 1 *If (G, ϕ) is more ambiguity averse than (G, ψ), then the SSCE's of (G, ψ) 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 provide intuition for this result in the Introduction. Now we can be more precise: Let σ^* be an SSCE of (G, ψ), 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 $\mu_{s_i^*}$ such that s_i^* is a best reply to

[22] With the convention that $\phi_i = \omega$ is more ambiguity averse than any ψ_i, and that if ϕ_i is more ambiguity averse than ω then $\phi_i = \omega$.

$\mu_{s_i^*}$ given ψ_i, that is, $V_i^{\psi_i}(s_i^*, \mu_{s_i^*}) \geq V_i^{\psi_i}(s_i, \mu_{s_i^*})$ for all s_i. We interpret $\mu_{s_i^*}$ as the belief held in the long-run by an agent who keeps playing the long-run strategy s_i^* in the stationary 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 $U(s_i^*, \sigma_{-i}^*)$, independently of his ambiguity attitudes (cf. Lemma 1). 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, $V_i^{\psi_i}(s_i, \mu_{s_i^*}) \geq V_i^{\phi_i}(s_i, \mu_{s_i^*})$. Therefore

$$V_i^{\phi_i}(s_i^*, \mu_{s_i^*}) = U(s_i^*, \sigma_{-i}^*) = V_i^{\psi_i}(s_i^*, \mu_{s_i^*}) \geq V_i^{\psi_i}(s_i, \mu_{s_i^*}) \geq V_i^{\phi_i}(s_i, \mu_{s_i^*})$$

for all s_i. This means that it is possible to justify σ^* as an SSCE of the more ambiguity averse game (G, ϕ) using the same profile of beliefs justifying σ^* as an SSCE of (G, ψ).

B. Relationships

Theorem 1 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 that

$$BSCE \subseteq SSCE \subseteq MSCE. \tag{6}$$

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.[23]

[23]But note that the inclusions $BSCE \subseteq MSCE$ and $SSCE \subseteq MSCE$ do not require ambiguity aversion. Furthermore, one can show that, in two-person games, $BSCE \subseteq SSCE$ independently of the ambiguity attitudes ϕ, due to the convex-valuedness of $\hat{\Sigma}_{-i}(s_i, \cdot)$ in this case (see Battigalli *et al.*, 2011).

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 equilibria (also when agents are ambiguity loving). 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 2 *If a profile of distributions σ^* is a Nash equilibrium of a game with feedback G, then it is a 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 ϕ.[24] In particular, we can enrich the chain of inclusions in (6) as follows:

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

under observable payoffs and ambiguity aversion.

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

[24]Hence, we also obtain existence of MSCE, by Theorem 1.

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.[25]

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 by emphasizing the key role played by payoff observability in establishing the inclusions in (6). The following example shows that, indeed, these inclusions need not hold when payoffs are not observable.

Example 4 *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 minimum 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 (justified by a suitable prior). 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,*

$$BSCE \cap MSCE = \emptyset \quad \textit{and} \quad SSCE \nsubseteq MSCE$$

and so the inclusions of (6) here do not hold. ▲

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

III. 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 that yields $n-1+\varepsilon$ utils ($0 < \varepsilon < 1/2$) and $n \geq 2$ 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 3). The game of Figure 1 obtains for $n = 2$.

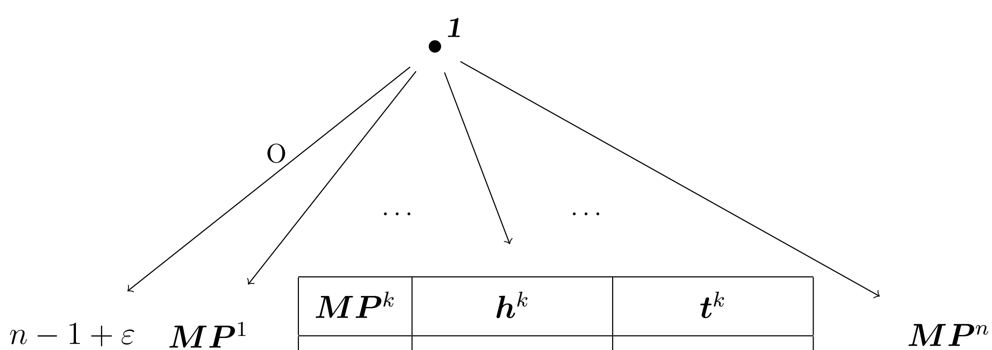

$\boldsymbol{MP}^k$	$\boldsymbol{h}^k$	$\boldsymbol{t}^k$
$\boldsymbol{H}^k$	$n+2(k-1)$	$n-k$
$\boldsymbol{T}^k$	$n-k$	$n+2(k-1)$

Figure 3: 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 (zero-sum) subgame k. 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 4.

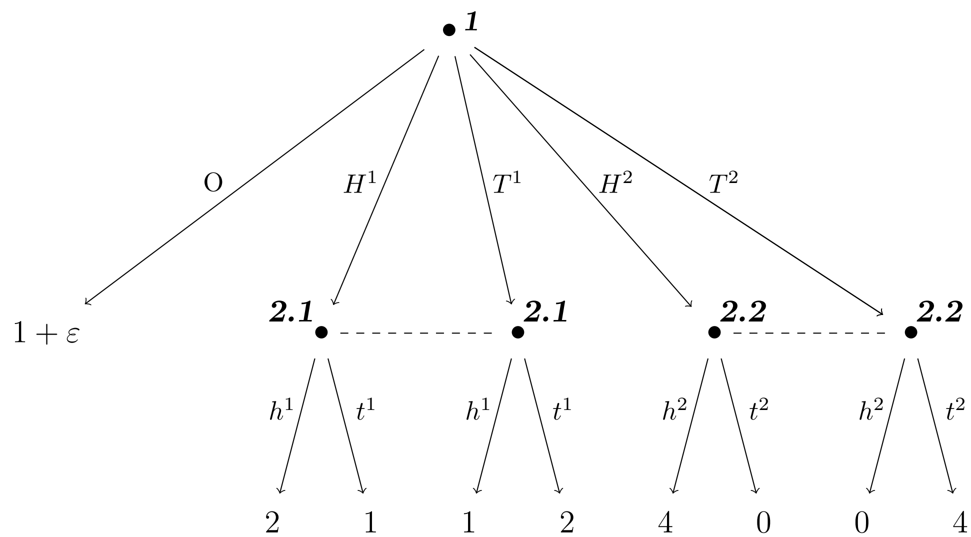

Figure 4: The case $n = 2$

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.[26]

[26]Each profile $\sigma = (\sigma_1, (\sigma_{2.k})_{k=1}^n)$ of the new n-person game can be mapped to an equiv-

That said, consider the game with feedback G_n derived from extensive-form game Γ_n 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 that[27]

$$\hat{\Sigma}_2(O, \sigma_2^*) = \Pi_{k=1}^n \Delta(S_{2.k}), \tag{7}$$

and

$$\hat{\Sigma}_2(H^k, \sigma_2^*) = \hat{\Sigma}_2(T^k, \sigma_2^*) = \{\sigma_2 : \sigma_{2.k} = \sigma_{2.k}^*\}. \tag{8}$$

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

$$\sigma_1^*(H^k) + \sigma_1^*(T^k) > 0 \Rightarrow \frac{\sigma_1^*(H^k)}{\sigma_1^*(H^k) + \sigma_1^*(T^k)} = \frac{1}{2} = \sigma_{2.k}^*(h^k). \tag{9}$$

Furthermore, for every (σ_1^*, σ_2^*) and $\bar{\sigma}_1^*$, if (σ_1^*, σ_2^*) is a (Bayesian, Smooth, Maxmin) SCE, and $\mathrm{supp}\sigma_1^* = \mathrm{supp}\bar{\sigma}_1^*$, then $(\bar{\sigma}_1^*, \sigma_2^*)$ is also a (Bayesian, Smooth, Maxmin) SCE.

Note that these necessary conditions do not restrict at all the set of strategies 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 (9) holds. The

alent 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.

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

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 a 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 $\mu_1 \in \Delta(\hat{\Sigma}_2(H^k, \sigma_2^*))$, which is only restricted by $\sigma_{2.k}^*$ (eqs. (7) and (8)). 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 3, in a SCE, the value that they assign to their strategy is its von Neumann-Morgenstern 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 mixed maxmin value $n - 1 + n/2 > n - 1 + k/2$. Furthermore, the outside option O is never an ambiguity-neutral best reply.[28] This explains the following:

Proposition 2 *The BSCE set of G_n coincides with the set of Nash equilibria. Specifically,*

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

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

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

Lemma 4 *Let ϕ_1 be concave. For all $j = 1, ..., n$, $\mu_1, \nu_1 \in \Delta\left(\Pi_{k=1}^n \Delta(S_{2.k})\right)$, if*

$$\mathrm{mrg}_{\Delta(S_{2.j})}\nu_1 = \frac{1}{2}\delta_{h^j} + \frac{1}{2}\delta_{t^j},$$

then

$$\max\{V_1^{\phi_1}(H^j, \mu_1), V_1^{\phi_1}(T^j, \mu_1)\} \geq V_1^{\phi_1}(H^j, \nu_1) = V_1^{\phi_1}(T^j, \nu_1).$$

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 $\Delta(S_{2.j})$ 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_{t^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 $(G_n, \phi_1^\alpha, \phi_2, ...\phi_n)$. 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 4, 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 5 *For all $x \neq y$, $M(\cdot, x, y)$ is strictly decreasing, continuous, and satisfies*

$$\lim_{\alpha \to 0} M(\alpha, x, y) = \frac{1}{2}x + \frac{1}{2}y \quad \text{and} \quad \lim_{\alpha \to +\infty} M(\alpha, x, y) = \min\{x, y\}. \qquad (10)$$

By Lemma 3, 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 4 and 5, 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 3, 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 5, 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

$$\Sigma^* = \{\sigma^* \in \Sigma : \text{eq. (9) holds}\}$$

of distribution profiles satisfying the necessary conditions of Lemma 3.[29] To

[29] This characterization holds for every parametrized family of distributions that satisfies, at every expected utility value $\bar{U}$, properties analogous to those of Lemma 5, with α replaced

summarize, by the properties of the function $M(\alpha, x, y)$ stated in Lemma 5, we can define strictly positive thresholds $\alpha_1 < \alpha_2 < ... < \alpha_n$ so that the following indifference conditions hold

$$\max_{j\in\{k+1,\dots,n\}} M(\alpha_{n-k}, n-j, n+2(j-1)) = n-1+\frac{k}{2}, \qquad k = 1, ..., n-1, \quad (11)$$

$$\max_{j\in\{1,\dots,n\}} M(\alpha_n, n-j, n+2(j-1)) = n-1+\varepsilon, \qquad (12)$$

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 3 *Let $\alpha_1 < ... < \alpha_n$ be the strictly positive thresholds defined by (11) and (12). 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

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

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

IV. Concluding remarks and related literature

The SCE concept characterizes stable patterns of behavior in games played recurrently. We analyze a notion of SCE with agents who have non-neutral attitudes toward uncertainty about the true steady-state data generating process. We showed that this uncertainty comes from a partial identification

by the coefficient of ambiguity aversion $-\phi_1''(\bar{U})/\phi'(\bar{U})$.

problem: The mapping from strategy distributions to the distributions of observations available to an agent is not one to one. We use 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 those of the strategies they can deviate to, which are thus penalized by higher ambiguity aversion. This allows us to derive intuitive relationships between different notions 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 develop our theoretical insights in the framework of population games played recurrently, but similar intuitions apply to different strategic contexts, such as repeated games, or 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 "experience 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 in 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) and to Battigalli *et al.* (2014) for a more detailed discussion.

More on robust preferences It is well known that the smooth ambiguity criterion corresponding to $\phi_i(t) = -e^{-\frac{t}{\alpha}}$ for all $t \in \mathbb{R}$, with constant absolute ambiguity aversion coefficient $\alpha > 0$, can be written as

$$V_i^{\phi_i}(s_i, \mu_i) = \inf_{\nu_i \ll \mu_i} \left(\int_{\mathrm{supp}\nu_i} U_i(s_i, \sigma_{-i}) \, \nu_i(\mathrm{d}\sigma_{-i}) + \alpha H(\nu_i || \mu_i) \right).$$

Here H is the Kullback-Leibler divergence; thus the corresponding smooth criterion is akin to the multiplier criterion of Hansen and Sargent (2001). This suggests considering robust preferences of the form

$$V_i^{\Phi_i}(s_i, \mu_i) = \inf_{\nu_i \ll \mu_i} \left(\int_{\mathrm{supp}\nu_i} U_i(s_i, \sigma_{-i}) \, \nu_i(\mathrm{d}\sigma_{-i}) + \Phi_i(\nu_i || \mu_i) \right), \tag{13}$$

where Φ_i is a generic divergence between priors, that is, a function

$$\Phi_i : \Delta(\Delta(S_{-i})) \times \Delta(\Delta(S_{-i})) \to [0, \infty]$$

such that $\Phi_i(\cdot || \mu_i)$ is convex and $\Phi_i(\mu_i || \mu_i) = 0$ for every μ_i. Maccheroni, Marinacci and Rustichini (2006) and Cerreia-Vioglio, Maccheroni, Marinacci and Montrucchio (2013a) show how $\Phi_i(\cdot || \cdot)$ captures ambiguity attitudes in a simple way: Φ_i is *more ambiguity averse* than Ψ_i if $\Phi_i(\cdot || \mu_i) \leq \Psi_i(\cdot || \mu_i)$ for every $\mu_i \in \Delta(\Delta(S_{-i}))$. It can be shown that all results in Section II hold when the smooth criterion (1) is replaced with the robust criterion (13).

Dynamic consistency and conditional beliefs To avoid dynamic consistency issues, we assume that agents play the strategic form of the recurrent game, i.e., an essentially simultaneous stage 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 conditional beliefs and planned continuation. The plans with this unimprovability property are obtained by means of a "folding back" procedure on the subjective decision tree implied by the agent's beliefs. We sketch how we can make this precise in the context of the smooth-ambiguity model, and thus provide a notion of dynamically consistent SSCE. Next we discuss the properties of this concept. We assume that agents' feedback functions satisfy *ex post perfect recall*, that is, after playing the game agents remember the information sets they crossed and the actions chosen at such information sets. For an in-depth analysis with proofs of claims, see Battigalli *et al.* (2014).

Each agent in role i has a system of beliefs $\mu_i(\cdot|\cdot)$ about distributions σ_{-i} given by a prior $\mu_i \in \Delta(\Delta(S_{-i}))$ and a posterior $\mu_i(\cdot|h_i)$ at each information set h_i of i. The predictive probability of reaching information set h_i – given that the agent chooses the actions leading to h_i – is $\mathbb{P}_{\mu_i}(h_i) = \int_{\Delta(S_{-i})} \sigma_{-i}(S_{-i}(h_i))\, \mu_i(\mathrm{d}\sigma_{-i})$, where $S_{-i}(h_i)$ denotes the set of strategy profiles s_{-i} consistent with h_i. If $\mathbb{P}_{\mu_i}(h_i) > 0$, the posterior belief $\mu_i(\cdot|h_i)$ is derived from the prior using Bayes rule;[30] otherwise, $\mu_i(\cdot|h_i)$ is derived from $\mu_i(\cdot|\bar{h}_i)$, where $\bar{h}_i$ is the information set closest to the root such that $\mathbb{P}_{\mu_i}(h_i|\bar{h}_i) > 0$ (note, it may be $\bar{h}_i = h_i$). Such system of beliefs yields a conditional probability system on $S_{-i} \times \Delta(S_{-i})$ given the collection of conditioning cylindrical events $S_{-i}(h) \times \Delta(S_{-i})$ (cf. Battigalli and Siniscalchi, 1999). A plan s_i is a *sequential best reply* to $\mu_i(\cdot|\cdot)$ if, at each information set h_i of i, it selects an action maximizing the KMM value $V_i^{\phi_i}$, given $\mu_i(\cdot|h_i)$ and the s_i-continuation after h_i. A profile of distributions σ^* is a *dynamically consistent SSCE*, for brevity SSCE^{DC}, if each s_i with $\sigma_i^*(s_i) > 0$ is a sequential best reply to some $\mu_{s_i}(\cdot|\cdot)$ such that the prior μ_{s_i} satisfies the confirmation condition $\mathrm{supp}\mu_{s_i} \subseteq \hat{\Sigma}_{-i}(s_i, \sigma_{-i}^*)$.

By the dynamic consistency of SEU maximization, SSCE^{DC} is realization-

[30] That is,

$$\mu_i(E_{-i}|h_i) = \frac{1}{\mathbb{P}_{\mu_i}(h_i)} \int_{E_{-i}} \sigma_{-i}(S_{-i}(h_i))\, \mu_i(\mathrm{d}\sigma_{-i})$$

for every Borel set $E_{-i} \subseteq \Delta(S_{-i})$.

equivalent to SSCE if agents are ambiguity neutral. The reason is that an ex ante SEU-optimal strategy can prescribe suboptimal actions only at information sets that the agent subjectively deems impossible; given ex post perfect recall, in a selfconfirming equilibrium, such information sets must be off the equilibrium path.[31] But, in general, SSCE^{DC} outcomes may differ from SSCE outcomes, because – as is well known[32] – an ex ante optimal strategy of an ambiguity averse agent may prescribe ex post suboptimal actions even at information sets the agent deems possible. If agents truly play the game in a sequential fashion, SSCE^{DC} is the relevant concept. Does a version of our comparative statics result (Theorem 1) hold for SSCE^{DC}? We can prove that, in games where each player moves at most once on any path (e.g., the game of Figure 4), if (G, ϕ) is more ambiguity averse than (G, ψ), then every SSCE^{DC} of (G, ψ) is also an SSCE^{DC} of (G, ϕ). The main intuition for the result is again a kind of status quo bias: By ex post perfect recall and observability of payoffs, actions chosen at information sets on the equilibrium path are unambiguous, whereas deviations may be perceived as ambiguous; hence, higher ambiguity aversion penalizes deviations.[33] We can show by example that the comparative statics statement cannot be generalized *as is* to all games, but we conjecture that a version of the result holds for outcome distributions.

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 – 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 SCE's.[34] The separation between ambiguity attitudes and beliefs in the KMM smooth-

[31] As we discuss below, dynamic consistency makes a difference if beliefs are also assumed to be rationalizable.

[32] See, e.g., Siniscalchi (2011).

[33] The proof, however, requires a non-trivial ancillary result.

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

ambiguity model allows, a relatively straightforward extension of this idea, which yields a notion of rationalizable SSCE, and a notion of rationalizable SSCE^{DC}. For truly dynamic games, rationalizable SSCE^{DC} excludes outcome distributions that are allowed by either rationalizable SSCE, or mere SSCE^{DC}, as pointed out in the early literature on selfconfirming equilibrium for the case of ambiguity-neutral agents (e.g., Dekel, Fudenberg and Levine, 1999). The reason is that rationalizable beliefs must assign probability zero to opponents' strategies that prescribe suboptimal actions at information sets off the equilibrium path. Hence, some "threats," or "promises" that support SSCE^{DC} outcomes may be deemed non-credible according to rationalizable SSCE^{DC}. We can easily extend our comparative statics result to rationalizable SSCE. As for rationalizable SSCE^{DC}, we can prove a version of the result for games where each player moves at most once (see Battigalli *et al.*, 2014).

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 $BSCE \subseteq SSCE \subseteq MSCE$ 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).[35]

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 focusing 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 each action in the support of σ_i^* is also optimal, therefore σ_i^* can be implemented by mapping each action in suppσ_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 suppσ_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 $\hat{S}_i$.[36] But, if this is the case, such commitment technology should be explicitly allowed by the rules of the game and represented in the game form.

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 using consistency and convergence

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

[36]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.

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 us 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 ambiguity averse agents tend to avoid.

Related literature As we mentioned, our notion of SCE subsumes earlier definitions due to Battigalli (1987) and Fudenberg and Levine (1993a) as special cases.[37] 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 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 opponents. 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 in 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), which makes formally explicit the decision maker's uncertainty about the true prob-

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

abilistic 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,[38] 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 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 assumes instead 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 opponents). 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 1 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

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

feedback, and provides a rigorous analysis of the relationship between MSCE and Lehrer's equilibrium concept.

Appendix

Proof of Lemma 1

Fix $i \in I$, $s_i \in S_i$, and $\sigma^*_{-i} \in \Delta(S_{-i})$. Since payoffs are observable, the payoff function $U_{s_i} : S_{-i} \to \mathbb{R}$ is $\mathcal{F}_{s_i}$-measurable, and therefore, for every $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i, \sigma^*_{-i})$, we have

$$\begin{aligned} U_i(s_i, \sigma_{-i}) &= \sum_{s_{-i} \in S_{-i}} U_i(s_i, s_{-i}) \sigma_{-i}(s_{-i}) = \int_{S_{-i}} U_{s_i} \sigma_{-i} \\ &= \int_{S_{-i}} U_{s_i} \sigma_{-i|\mathcal{F}_{s_i}} = \int_{S_{-i}} U_{s_i} \sigma^*_{-i|\mathcal{F}_{s_i}} = \int_{S_{-i}} U_{s_i} \sigma^*_{-i} \\ &= \sum_{s_{-i} \in S_{-i}} U_i(s_i, s_{-i}) \sigma^*_{-i}(s_{-i}) = U_i(s_i, \sigma^*_{-i}), \end{aligned}$$

as wanted. ∎

Proofs for Section II

Proof of Theorem 1 For every $i \in I$, $s_i \in S_i$, $(s^*_i, \sigma^*_{-i}) \in S_i \times \Delta(S_{-i})$, $\mu_{s^*_i}$ with support in $\hat{\Sigma}_{-i}(s^*_i, \sigma^*_{-i})$, and every ϕ_i more ambiguity averse than ψ_i

$$V_i^{\psi_i}(s_i, \mu_{s^*_i}) \geq V_i^{\phi_i}(s_i, \mu_{s^*_i}) \geq V_i^{\omega}(s_i, \mu_{s^*_i}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s^*_i, \sigma^*_{-i})} U_i(s_i, \sigma_{-i}). \quad (14)$$

The last inequality is obvious because $\mathrm{supp}\mu_{s^*_i} \subseteq \hat{\Sigma}_{-i}(s^*_i, \sigma^*_{-i})$. The central inequality is also obvious if $\phi_i = \omega$, otherwise choose $\sigma'_{-i} \in \mathrm{supp}\mu_{s^*_i}$ such that

$$U_i(s_i, \sigma'_{-i}) = \min_{\sigma_{-i} \in \mathrm{supp}\mu_{s^*_i}} U_i(s_i, \sigma_{-i}) = V_i^{\omega}(s_i, \mu_{s^*_i})$$

now

$$\begin{aligned}
V_i^{\phi_i}(s_i,\mu_{s_i^*}) &= \phi_i^{-1}\left(\int_{\mathrm{supp}\mu_{s_i^*}} \phi_i(U_i(s_i,\sigma_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&\geq \phi_i^{-1}\left(\int_{\mathrm{supp}\mu_{s_i^*}} \phi_i(U_i(s_i,\sigma'_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&= \phi_i^{-1}\left(\phi_i(U_i(s_i,\sigma'_{-i}))\right) = V_i^{\omega}(s_i,\mu_{s_i^*}),
\end{aligned}$$

as desired. As for the first inequality of (14):

- if $\psi_i = \omega$, then ϕ_i, being more ambiguity averse than ω, coincides with ω, and the inequality is an equality;
- if $\psi_i \neq \omega$ and $\phi_i = \omega$, the inequality follows from the previous argument;
- if $\psi_i \neq \omega$ and $\phi_i \neq \omega$, then there exists a continuous concave and strictly increasing function φ_i such that $\phi_i = \varphi_i \circ \psi_i$,[39] by the Jensen's inequality we have

$$\begin{aligned}
V_i^{\psi_i}(s_i,\mu_{s_i^*}) &= \psi_i^{-1}\left(\int_{\mathrm{supp}\mu_{s_i^*}} \psi_i(U_i(s_i,\sigma_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&= (\psi_i^{-1}\circ\varphi_i^{-1})\circ\varphi_i\left(\int_{\mathrm{supp}\mu_{s_i^*}} \psi_i(U_i(s_i,\sigma_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&\geq (\psi_i^{-1}\circ\varphi_i^{-1})\left(\int_{\mathrm{supp}\mu_{s_i^*}} (\varphi_i\circ\psi_i)(U_i(s_i,\sigma_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) \\
&= \phi_i^{-1}\left(\int_{\mathrm{supp}\mu_{s_i^*}} \phi_i(U_i(s_i,\sigma_{-i}))\mu_{s_i^*}(\mathrm{d}\sigma_{-i})\right) = V_i^{\phi_i}(s_i,\mu_{s_i^*}).
\end{aligned}$$

Now let σ^* be a SSCE of the less ambiguity averse game (G,ψ). Fix $i \in I$, and pick $s_i^* \in \mathrm{supp}\sigma_i^*$, $\mu_{s_i^*}$ with support in $\hat{\Sigma}_{-i}(s_i^*,\sigma_{-i}^*)$ such that

$$s_i^* \in \arg\max_{s_i\in S_i} V_i^{\psi_i}(s_i,\mu_{s_i^*}). \tag{15}$$

[39] Note that $\phi_i, \psi_i : \mathbb{U}_i \to \mathbb{R}$ are continuous and $\mathbb{U}_i$ is connected. Moreover, $\varphi_i : \psi_i(\mathbb{U}_i) \to \mathbb{R}$ is increasing and such that $\phi_i = \varphi_i \circ \psi_i$. Therefore φ_i is continuous too.

We want to show that

$$s_i^* \in \arg\max_{s_i \in S_i} V_i^{\phi_i}(s_i, \mu_{s_i^*}), \tag{16}$$

which implies the first claim. Since payoffs are observable, by Lemma 1, $U_i(s_i^*, \sigma_{-i}) = U_i(s_i^*, \sigma_{-i}^*)$ for every $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*) \supseteq \text{supp}\mu_{s_i^*}$. Thus

$$U_i(s_i^*, \sigma_{-i}^*) = V_i^{\phi_i}(s_i^*, \mu_{s_i^*}) = V_i^{\psi_i}(s_i^*, \mu_{s_i^*}) = V_i^{\omega}(s_i^*, \mu_{s_i^*}) = \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i(s_i^*, \sigma_{-i}) \tag{17}$$

concluding the proof of the first part of the statement since, together with (15) and (14), it delivers (16), in fact

$$V_i^{\phi_i}(s_i^*, \mu_{s_i^*}) = V_i^{\psi_i}(s_i^*, \mu_{s_i^*}) \geq V_i^{\psi_i}(s_i, \mu_{s_i^*}) \geq V_i^{\phi_i}(s_i, \mu_{s_i^*}) \quad \forall s_i \in S_i.$$

We now prove that all SSCE's are MSCE's. Let σ^* be an SSCE of a game (G, ψ). Fix $i \in I$, and pick $s_i^* \in \text{supp}\sigma_i^*$, $\mu_{s_i^*}$ with support in $\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$ such that (15) holds. By (17), we have

$$\min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i(s_i^*, \sigma_{-i}) = U_i(s_i^*, \sigma_{-i}^*) = V_i^{\psi_i}(s_i^*, \mu_{s_i^*}),$$

thus (15) and (14) deliver

$$\min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i(s_i^*, \sigma_{-i}) = V_i^{\psi_i}(s_i^*, \mu_{s_i^*}) \geq V_i^{\psi_i}(s_i, \mu_{s_i^*}) \geq \min_{\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)} U_i(s_i, \sigma_{-i})$$

for all $s_i \in S_i$, as wanted. ■

Proof of Lemma 2 Fix a mixed strategy Nash equilibrium σ^* of G. Pick any i and pure strategy $s_i^* \in$suppσ_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, $\delta_{\sigma_{-i}^*}$ is with support in $\hat{\Sigma}_{-i}(s_i^*, \sigma_{-i}^*)$. Since $V_i^{\phi_i}\left(s_i, \delta_{\sigma_{-i}^*}\right) = U_i\left(s_i, \sigma_i^*\right)$ 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 1 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 $s_i^* \in$suppσ_i^*, then,

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

By Lemma 1, observability of payoffs implies $U_i(s_i,\sigma_{-i}) = U_i(s_i,\sigma_{-i}^*)$ for every s_i and $\sigma_{-i} \in \hat{\Sigma}_{-i}(s_i,\sigma_{-i}^*)$. Own-strategy independence of feedback implies that, for each s_i, $\mathcal{F}_{s_i} = \mathcal{F}_{s_i^*}$ hence

$$\begin{aligned}
\hat{\Sigma}_{-i}(s_i,\sigma_{-i}^*) &= \left\{\sigma_{-i} \in \Pi_{j\neq i}\Delta(S_j) : \sigma_{-i|\mathcal{F}_{s_i}} = \sigma_{-i|\mathcal{F}_{s_i}}^*\right\} \\
&= \left\{\sigma_{-i} \in \Pi_{j\neq i}\Delta(S_j) : \sigma_{-i|\mathcal{F}_{s_i^*}} = \sigma_{-i|\mathcal{F}_{s_i^*}}^*\right\} = \hat{\Sigma}_{-i}(s_i^*,\sigma_{-i}^*).
\end{aligned}$$

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

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

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

Proofs for Section III

Proof of Proposition 2. For any prior μ_1, the ambiguity-neutral subjective value of playing any Matching Pennies subgame k is

$$\begin{aligned}
& \max\{V_1(H^k,\mu_1), V_1(T^k,\mu_1)\} \\
= \; & \max\left\{\begin{array}{l} \bar{\mu}_1^k(h^k)(n+2(k-1)) + (1-\bar{\mu}_1^k(h^k))(n-k), \\ \bar{\mu}_1^k(h^k))(n-k) + (1-\bar{\mu}_1^k(h^k))(n+2(k-1)) \end{array}\right\} \\
\geq \; & n-1+\frac{k}{2} > n-1+\varepsilon = u_1(O),
\end{aligned}$$

where $n-1+k/2$ is the mixed maxmin value of subgame k, $\bar{\mu}_1^k = \mathrm{mrg}_{S_{2.k}}\bar{\mu}_1$ and $\bar{\mu}_1$ 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 $\bar{\mu}_1$.[40] Furthermore, no strategy H^k or T^k with $k < n$ can have positive measure in a BSCE. Indeed, by (9), if $s_1^k \in \{H^k, T^k\}$ has positive probability in an equilibrium σ^*, then for every belief $\mu_1 \in \Delta(\hat{\Sigma}_2(s_1^k, \sigma_2^*))$, the value of s_1^k is

$$V_1(s_1^k, \mu_1) = U_1\left(s_1^k, \sigma^*_{-\{1,2.k\}} \times \left(\frac{1}{2}h^k + \frac{1}{2}t^k\right)\right) = n - 1 + \frac{k}{2},$$

while the ambiguity-neutral value of deviating to subgame n is

$$\max\{V_1(H^n, \mu_1), V_1(T^n, \mu_1)\} \geq n - 1 + \frac{n}{2}.$$

Therefore, eq. (9) implies $\sigma_1^*(H^n) = \sigma_1^*(T^n) = \sigma_{2.n}^*(h^n) = \frac{1}{2}$ 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 4 is based on the following lemma, where $\mathbf{I}$ is the unit interval $[0, 1]$ endowed with the Borel σ-algebra.

Lemma 6 *Let $\varphi : \mathbf{I} \to \mathbb{R}$ be increasing and concave. For each Borel probability measure μ on $\mathbf{I}$*

$$\max\left\{\int_{\mathbf{I}} \varphi(x)\,\mu(\mathrm{d}x), \int_{\mathbf{I}} \varphi(1-x)\,\mu(\mathrm{d}x)\right\} \geq \frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0). \qquad (18)$$

Proof. Let

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

Then

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

[40] Recall that given a prior μ_i on a Borel subset Σ_{-i} of $\Delta(S_{-i})$, its predictive $\bar{\mu}_i$ is defined by

$$\bar{\mu}_i(s_{-i}) = \int_{\Sigma_{-i}} \sigma_{-i}(s_{-i})\,\mu_i(\mathrm{d}\sigma_{-i}).$$

where $\mu_\tau = \mu \circ \tau^{-1}$. In particular, for $\varphi = \mathrm{id}_{\mathbf{I}}$ it follows that $1 - \int_{\mathbf{I}} x\mu(\mathrm{d}x) = \int_{\mathbf{I}} y\mu_\tau(\mathrm{d}y)$. Thus (18) becomes

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

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

$$\int_{\mathbf{I}} \varphi(x)\,\nu(\mathrm{d}x) \geq \frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0)\,. \tag{19}$$

Denote by $F(x) = \nu([0,x])$ 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 $x \in [0,1)$. Note that 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

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

a contradiction. 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 (19) holds for all increasing and concave φ. ■

Proof of Lemma 4 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 $\mu \in \Delta(\mathbf{I})$ be the marginal belief about $x = \sigma_{2.k}(h^k)$ derived from μ_1. Recall that ν_1 is a prior such that $\mathrm{mrg}_{\Delta(S_{2.j})}\nu_1 =$

$\frac{1}{2}\delta_{h^j}+\frac{1}{2}\delta_{t^j}$. With this,

$$\begin{aligned}
& \max\{V_1^{\phi_1}(H^j,\mu_1),V_1^{\phi_1}(T^j,\mu_1)\} \\
= & \max\left\{\phi_1^{-1}\left(\int_{\mathbf{I}}\varphi(x)\,\mu(\mathrm{d}x)\right),\phi_1^{-1}\left(\int_{\mathbf{I}}\varphi(1-x)\,\mu(\mathrm{d}x)\right)\right\} \\
= & \phi_1^{-1}\left(\max\left\{\int_{\mathbf{I}}\varphi(x)\,\mu(\mathrm{d}x),\int_{\mathbf{I}}\varphi(1-x)\,\mu(\mathrm{d}x)\right\}\right)
\end{aligned}$$

and

$$V_1^{\phi_1}(H^j,\nu_1)=V_1^{\phi_1}(T^j,\nu_1)=\phi_1^{-1}\left(\frac{1}{2}\varphi(1)+\frac{1}{2}\varphi(0)\right).$$

Hence, the thesis is implied by Lemma 6. ■

Proof of Proposition 3 By Lemma 3, $SSCE(\alpha)$ 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 are no incentives to deviate to any subgame j. We rely on Lemma 4 to specify a belief $\mu_1^{s_1}\in\Delta(\hat{\Sigma}_2(s_1,\sigma_2^*))$ that minimizes the incentive to deviate. Thus, s_1 can be played in equilibrium if and only if it is a best reply to $\mu_1^{s_1}$. Specifically,

$$\mu_1^O=\times_{j=1}^n\left(\frac{1}{2}\delta_{h^j}+\frac{1}{2}\delta_{t^j}\right)\in\Delta(\hat{\Sigma}_2(O,\sigma_2^*))=\Delta\left(\Pi_{j=1}^n\Delta(S_{j.k})\right),$$

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

$$\mu_1^k=\delta_{\frac{1}{2}h^k+\frac{1}{2}t^k}\times\left(\times_{j\neq k}\left(\frac{1}{2}\delta_{h^j}+\frac{1}{2}\delta_{t^j}\right)\right)$$

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 μ_1^O, and can have positive measure in equilibrium, if and only if

$$n-1+\varepsilon\geq\max_{j\in\{1,\dots,n\}}M(\alpha,n-j,n+2(j-1)). \tag{20}$$

By Lemma 5 there is a unique threshold $\alpha_n > 0$ that satisfies (20) as an equality so that (20) 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 μ_1^k, and can have positive measure in equilibrium, if and only if

$$n - 1 + \frac{k}{2} \geq \max_{j \in \{1,...,n\}} M(\alpha, n-j, n+2(j-1)), \tag{21}$$

where

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

because, for all $\alpha > 0$ and $j \leq k$

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

By Lemma 5 there is a unique threshold $\alpha_{n-k} > 0$ that satisfies (21) as an equality so that (21) 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 2, $SSCE(\alpha) = BSCE = NE$ in this case. ■

References

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

[2] Azrieli, Yaron (2009): "On pure conjectural equilibrium with non-manipulable information," *International Journal of Game Theory*, 38,

209-219.

[3] Battigalli, Pierpaolo (1987): *Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali*, unpublished thesis, Università Bocconi.

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

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

[7] Battigalli, Pierpaolo, Mario Gilli and M. Cristina 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] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni and M. Marinacci (2014): "Selfconfirming equilibrium and sequential rationality with ambiguity averse agents," mimeo, Università Bocconi.

[11] Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci and Luigi Montrucchio (2013a): "Ambiguity and robust statistics," *Journal of Economic Theory*, 148, 974-1049.

[12] Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci and Luigi Montrucchio (2013b): "Classical subjective expected utility," *Proceedings of the National Academy of Sciences*, 110, 6754-6759.

[13] Cho, In-Koo and Thomas J. Sargent (2008): "Self-confirming equilibrium," in S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics, 2*$^{\text{nd}}$ ed., London: Palgrave Macmillan.

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

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

[16] Fudenberg, Drew and Yuichiro Kamada (2011): "Rationalizable partition-confirmed equilibrium," mimeo, Harvard University.

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

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

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

[20] Gilboa, Itzhak and Massimo 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.

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

[22] Hansen, Lars Peter and Thomas J. Sargent (2001): "Robust control and model uncertainty," *American Economic Review*, 91, 60-66.

[23] Hansen, Lars Peter and Thomas J. Sargent (2008): *Robustness.* Princeton: Princeton University Press.

[24] Karlin, Samuel and Albert Boris J. Novikoff (1963): "Generalized convex inequalities," *Pacific Journal of Mathematics*, 13, 1251-1279.

[25] Klibanoff Peter, Massimo Marinacci and Sujoy Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849-1892.

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

[27] Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006): "Ambiguity aversion, robustness, and the variational representation of preferences," *Econometrica*, 74, 1447-1498.

[28] Rubinstein, Ariel and Asher Wolinsky (1994): "Rationalizable conjectural equilibrium: Between Nash and rationalizability," *Games and Economic Behavior,* 6, 299-311.

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

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

[31] von Neumann, John and Oskar Morgenstern (1947): *The Theory of Games and Economic Behavior* (Second ed.). Princeton: Princeton University Press.

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