CERTAINTY CRISES, AMBIGUITY AVERSION AND SELF-CONFIRMING EQUILIBRIUM

Matteo Pozzi †

DOI: 10.82029/2025014

> "Ignoranti quem portum petat nullus suus ventus est.[1]„
>
> Lucius Annaeus Seneca,
> *Epistula ad Lucilium LXXI*

Abstract. I explore how doubt affects economic outcomes in strategic settings. Doubt is defined as a revision of subjective beliefs over payoff-relevant statistical models consistent with objective evidence, leading to consideration of more extreme possibilities. This work builds upon research on Self-Confirming Equilibrium (SCE) with ambiguity-averse players developed by Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2015), extending it to games with payoff uncertainty and exploring the consequences of differences in the perception of ambiguity. The smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) is applied, for it allows clear separation of attitudes toward ambiguity from its perception. The key finding is that players with heightened awareness of uncertainty present in their decision environment are more likely to exhibit *status quo* bias, favoring past behavior and leading to objectively worse outcomes. This occurs even when ambiguity aversion is mild. This result is demonstrated through a comparative statics exercise on the equilibrium set, where players' beliefs are adjusted to reflect varying degrees of ambiguity perception. An application to a simple contracting game illustrates how ambiguity aversion and perception can lead to market shutdowns, such as chronic worker shortages or underinvestment, in spite of mechanisms designed to mitigate these issues.

JEL classification: C72, C73

Keywords: Ambiguity Aversion, Uncertainty Perception, Self-Confirming Equilibrium.

1. INTRODUCTION

Doubt is a common occurrence in decision-making. Worldly wisdom suggests doubt is desirable, for it helps to carefully evaluate all available information, and act upon considerate judgment thereof. Second-guessing oneself is thus arguably reasonable within environments characterized by persistent uncertainty to the extent that this may aid the deliberation process. I clarify the game-theoretic consequences of doubt, which can be informally defined as the revision of subjective beliefs over plausible statistical models relevant to a decision problem, to consider a larger set of possibilities consistent with objective evidence. Such a revision is not put into formal relation with the standard rules of Bayesian updating. It serves exercises of model equilibrium comparative statics in the event of believing that more state distributions are possible than was previously thought. The formal model developed in the thesis combines elements from Decision Theory and Game Theory, building on the existing literature of equilibrium analysis with ambiguity averse-players.

†Pozzi: m.pozzi@studbocconi.it; matteo.pozzi@ecb.europa.eu. I thank Prof. Massimo Marinacci for his excellent supervision throughout my thesis work, and Prof. Pierpaolo Battigalli for insightful comments and feedback in view of my defense. I also thank the editor and two anonymous referees for useful comments and feedback.

[1]Translation: "No wind is favorable to the sailor who knows not the harbor he seeks."

The relevance of this topic can be understood from a historical perspective. The enthusiastic intellectual attitudes and ingenuity that spread at the time of the Industrial Revolution fostered the drive for innovation that led to marvelous inventions. In turn, this brought to unforeseen increases in the marginal efficiency of capital. Almelhem et al. (2023) recently provided robust evidence in support of this explanation for modern economic growth, attributable to Mokyr (2017). Burn-Murdoch (2024) instead presented evidence that Western cultures are currently abandoning their culture of progress, turning "toward one of caution, worry and risk-aversion", which echoes the notion of doubt described above. In parallel with Mokyr's argument, he observes that this shift has been accompanied by increasingly sluggish growth, suggesting a connection between these phenomena. In what follows, I offer a solid framework to rationalize this hypothesis. Before, however, we must make a series of nontrivial modeling decisions on which tools to use to represent these behavioral traits. The proposed notion of doubt does not easily lend itself to a proper formalization.

A relatively straight-forward approach is to adapt the framework of Battigalli et al. (2015) to study equilibrium sets constrained by alternative sets of conjectures, carefully constructed to reflect varying degrees of ambiguity perception. Note that beliefs cannot truly abstract from the decision environment, as the probabilistic weight assigned to alternative hypotheses necessarily depends on what the latter are. It may then seem unappealing to consider preferences whose subjective probabilistic component is constrained *a priori*. However, the analysis reveals that the extent of ambiguity perception indeed has normative implications. Since this is the approach we take, we will need to combine contemporary findings from Game Theory, particularly Self-Confirming Equilibrium analysis, with a solid Decision-theoretic model for representing uncertainty averse preferences.

Following the abstract axiomatization of Expected Utility by von Neumann and Morgenstern (1944), a staple within the theory of Choice Under Risk, Decision Theory has long been concerned with modeling Choice Under Uncertainty, where the latter is interpreted in the sense of Knight (1921). The ground-breaking leap of Savage's (1954) Subjective Expected Utility (SEU) representation was challenged by field and thought experiments in the spirit of Ellsberg (1961). The first proposals within the Rational Choice Theory approach to represent preferences with alternative attitudes toward this type of uncertainty appeared starting with Schmeidler (1986) and Gilboa and Schmeidler (1989). Then, a vast literature blossomed, whose works all share the underlying common motivation to accommodate for ambiguity aversion. The ensuing representation Theorems account for this feature but are often limited operationally in that they do not cleanly separate attitudes toward ambiguity from its perception. The representation of Klibanoff et al. (2005) is a notable exception, which allows the modeler to ask questions about the effects of varying degrees of both uncertainty aversion and perceived uncertainty. While risk and uncertainty attitudes shall be regarded as stable personal traits, elicited through consistent revealed preferences, perception can be elicited but may vary upon the context of the decision problem.

At the same time and in the same opus, von Neumann and Morgenstern (1944) put forth the fundamental elements for the development of a new framework to deal with economic problems more rigorously and flexibly: Game Theory. This field of study proved extremely prolific, and allowed the formalization of a rich set of solution concepts. While the most widespread and appreciated notions in the theoretical and applied literature may be Rationalizability (Bernheim, 1984; Pearce, 1984) and Nash (1950) Equilibrium, growing attention has been given to Self-Confirming Equilibrium, learning dynamics, and their interplay (Battigalli et al., 1992; Fudenberg and Levine, 1993, 1998). Steady-states of a learning process are characterized by stationary strategy profile distributions which are optimal under the equilibrium private information and beliefs of each player. Self-Confirming Equilibrium is the solution concept we apply for it forms a natural domain to explore the relationship between decision-makers, their knowledge, and behavior under uncertainty (Battigalli et al., 2015).

The reader may question the appropriateness of such theoretical approach for the problem at hand. Within the rational choice theory, decision models shall capture tolerable economic behavior. Starting with a minimal set of assumptions that appeal to some "normative" intuition, a utilitarian representation of preferences is derived through logical inference. Such assumptions are always held to a standard of rationality, albeit the latter term may *prima facie* seem quite vague. For instance, the transitivity assumption present in almost all economic models has been predicated upon Money Pump arguments (Gustafsson, 2022). Hence, it would be a maintained normative assumption because of its desirability and not due to its supposed veracity. Indeed, it has been pointed out that such arguments are often not very compelling in motivating the normative desirability of transitivity.[2] Conducting ourselves a formal investigation of such issues is beyond the scope of the discussion. Nevertheless, Gilboa et al. (2010) provide us with an intuitive motivation for the axiomatic approach. Imagine a decision-maker is held accountable of her decisions; would she necessarily feel embarrassed if asked to explain her reasoning?

This approach differs from Behavioral Economics, where traditional economic assumptions are typically traded off in favor of reliable psychological and neurological foundations to decision-making, thus introducing notions of biases or entertaining that individuals' rationality is bounded.[3] Consider, for instance, a traditional consumer problem under constraint. Traditionally, enlarging the choice set of available alternatives shall weakly improve the consumer's welfare. Following a behavioral approach, however, we may take into account the neurological and psychological evidence that the so-called Set-Size Effect is nonmonotonic.[4] While behavioral models offer valuable insights, this is not the approach followed here; the aim is to reflect on the implications of normative models of ambiguity aversion and perception.

The representation of Klibanoff et al. (2005) does not impose restrictions on the elicited subjective second-order probability assigned to models compatible with the evidence. Since the precise shape of the representation depends on preferences themselves, the analysis confronts decision-makers who ultimately differ in their "taste" in the processing of information. Hence, it may not be concluded that more naïve judgments are irrational. They may be attributed to mood, intuition; perhaps *animal spirits* have an influence on beliefs (Keynes, 1936; Akerlof and Shiller, 2009). Still, it is sensible to confront decision-makers who differ in their naïveté when confronting uncertain environments. Results help highlight the role of human attitudes in economic problems even within standard economic frameworks. Simple formal conditions on beliefs needed to uphold some equilibria rather than others may even lead, *inter alia*, to a better understanding of the effects of potentially ambiguous communication between institutions and the public.

Therefore, I explore how markets can plausibly get stuck into "certainty traps" without departures from (subjective) rationality. The Smooth Self-Confirming Equilibrium solution concept introduced by Battigalli et al. (2015) is applied to perform a comparative statics exercise on the equilibrium set of a game with feedback. Formally, I compare alternative sets of beliefs that differ in their perceived uncertainty. I leverage an example showing how more awareness of uncertainty (a doubt, or more intuitively a *certainty crisis*) can bring about objectively worse equilibrium strategies. The analysis is generalized to encompass games where own payoff-relevant information is not fully observable to every player.

In the rest of this work, I adopt whenever possible the formal notation used by Battigalli et al. (2023) and specify when is done otherwise. In Section 2, I provide a brief discussion of decision models used to deal with ambiguity attitudes, focusing on the representation of Klibanoff et al. (2005), its positioning within the literature, and the motivation for applying this criterion. I present a general definition of ambiguity-neutral (or Bayesian) Self-Confirming Equilibrium for games

[2]For instance, they "require a non-sequitur that calls for the additional hypothesis that a [decision-maker] willingly participates in a process that leaves him worse off". For a nuanced discussion of the normative appeal of transitivity and the treatment of intransitive preferences, see Fishburn (1991).

[3]These models are thus said to be "descriptive": they translate empirical regularities into assumptions, whose reliability rests on the modeler's trust that they are natural, i.e. do not depend on contextual factors.

[4]See for instance Iyengar and Lepper (2000) and Seuanez-Salgado (2006).

with feedback and payoff uncertainty, and its generalization to games with ambiguity attitudes introduced by Battigalli et al. (2015). In Section 3, I define ancillary games characterized by the sets of conjectures players may be assumed to hold when entering the game; these are used to produce results concerning equilibrium inclusion that complement the Theorems on comparative ambiguity aversion of Battigalli et al. (2015). In Section 4, I present an application of this comparative statics exercise to a simple contracting game that resonates with the literature on adverse selection. In Section 5, I draw conclusions and discuss limitations and possible extensions to my work. All proofs are contained in the Appendix.

2. AMBIGUITY ATTITUDES AND SELF-CONFIRMING EQUILIBRIUM

Throughout the following, I analyze the strategic form of a finite multistage game with payoff uncertainty, feedback and ambiguity attitudes within the formal framework developed by Battigalli et al. (2015) (hereon, BCMM). Section 2.1 details the structure of the game and introduces notation. It is convenient to consider the strategic form of the dynamic game and assume that, upon entering the game, players choose a fully specified strategy (*covert commitment*) that is then automatically implemented as the play unfolds.

It must be clear that this analysis is not equivalent to the extensive-form analysis of the same game whenever players are not ambiguity-neutral decision-makers. Dynamic consistency underpins the notion that "we can analyze the essential aspects of Self-Confirming Equilibrium for multistage games with feedback by looking at the strategic (or normal) form [since] any fixed strategy profile is a Self-Confirming Equilibrium of [the extensive-form game] if and only if there is a Self-Confirming Equilibrium [of the strategic form] inducing the same terminal history".[5] (Battigalli et al., 2023). A game with non-Bayesian decision-makers may not satisfy this property. If the intrapersonal version of the One-Deviation Principle fails, as is the general case for players averse to ambiguity, *ex ante* optimal strategies do not correspond to unimprovable strategies in the sense of Battigalli et al., 2019a, meaning that sophisticated players in extensive-form games who cannot fully commit to carrying out the strategy established ex-ante may have a strict preference for deviating somewhere along the path (provided that at least one player moves more than once). At the same time, this may not rule out altogether the existence of equilibrium profiles, which are also sequentially optimal.

Nevertheless, the strategic form analysis is both convenient from an analytic standpoint, and meaningful in classes of games where covert commitment is credible and in the special case of static games (degenerate multistage games where all players are active only at the root). More broadly, it is relevant in light of the relationship between the Self-Confirming Equilibrium solution concept and learning dynamics: monotonicity results regarding the strategic form stem from the influence ambiguity aversion has on learning *regardless* of dynamic inconsistency, as more uncertainty averse agents tend to stifle their experimentation after a few repetitions of a game (Battigalli et al., 2019b).

2.1. Game description

I adopt the *Nash mass action* (or anonymous interaction) view detailed in BCMM, and generalize their definitions of Smooth and Maxmin Self-Confirming Equilibrium to a setting with payoff uncertainty. According to the anonymous interaction interpretation, each player *role* $i \in I$ represents a large population, from which an agent is "drawn at random and matched to play the game [with agents drawn at random from all other roles $j \neq i$], then separated and rematched with (almost certainly) different opponents. After each play in which he was involved, an agent obtains some evidence on how the game was played." This conveniently allows to assume that any one player always only focuses on the present stage and disregards future payoffs ("myopic"

[5] The two equilibrium strategy profiles are said to be "realization-equivalent" in this case.

agents)—because she assigns negligible probability to the event of rematching with the same person twice. In addition, we assume uncertainty stems both from coplayers' strategies $(S_i)_{j\neq i}$, and their private information types $(\Theta_i)_{j\neq i}$ (along with nature Θ_0). The uncertainty space is thus $S_{-i} \times \Theta_{-i} \times \Theta_0$, where $S_{-i} := \times_{j\neq i} S_j$ and $\Theta_{-i} := \times_{j\neq i}\Theta_j$, and we only consider finite sets for simplicity. As described above, I restrict the analysis to the strategic-form of the game. Thus analogously to BCMM, with the addition of information types, the rules of the game make up game form with feedback[6]

$$\langle I, \Theta_0, q_0, (\Theta_i, q_i, S_i, M_i, F_i)\rangle$$

- I is the (finite) set of player roles;
- $\Theta := \Theta_0 \times \bigtimes_{i\in I} \Theta_i$ is the set of information type profiles —which is relevant to the game form to the extent that information determines the (outcome associated with each) terminal history;
- $q := (q_i)_{i\in I\cup\{0\}}$ is the profile of (exogenous) information type distributions;
- $S_i := \bigtimes_{h\in H} \mathcal{A}_i(h)$ is the individual finite strategy set of a player $i \in I$;
- M_i is the set of messages player $i \in I$ can observe *ex post* (whenever a play of the game is terminated);
- $F_i := f_i \circ \zeta : \Theta \times S \to M_i$ is the private strategic form feedback function of $i \in I$.

The extensive form of the game is not very relevant to us; nonetheless, I offer a brief description to foster intuition and clarify notation. For each t in a finite number of periods T, there is a non-empty subset $J_t \subseteq I$ of players who are *active* in the sense that they have to choose an action during that turn. Every sequence of profiles of actions taken from the beginning of the game draws a map, or *history*, through the nodes crossed during the game. The root of the game is the empty sequence. As players choose their actions at t, they reach history $h = (a^1, a^2 ..., a^t) \in \overline{H}$. H is the set of non-terminal histories, i.e., all decision nodes where there are still actions to be chosen by at least some player, and the game has not terminated yet. By contrast, Z denotes the set of terminal histories. Clearly, $H \cup Z = \overline{H}$ (while their intersection is empty). The action set of each player i is A_i. We may however want to allow for the specific set of available moves at any point in the game to depend on how the game has unfolded so far. Thus we denote by $\mathcal{A}_i : H \rightrightarrows A_i$ the feasibility correspondence of i, so that by definition S_i contains all possible *plans* of i, i.e. prescriptions on which actions she should carry out among those available depending on the contingency. $\zeta : S \to Z$ denotes the *path function*, which associates any profile of strategies to the terminal node it induces.

The feedback structure of the game form is described by the profile of functions $f_i : \Theta \times Z \to M_i$. When these are combined with the path function, we obtain profile $F = (F_i)_{i\in I}$ which gives rise, for each $i \in I, \theta_i \in \Theta_i$ and pure strategy $s_i \in S_i$, to an ex-post information partition of the uncertainty space,

$$\mathcal{F}_{s_i,\theta_i} = \left\{F^{-1}_{s_i,\theta_i}(m_i) \subseteq \Theta_0 \times \Theta_{-i} \times S_{-i} : m_i \in M_i\right\} \tag{1}$$

Battigalli et al. (2015, 2016, 2023) offer thorough discussions of general feedback structures, together with definitions of games with feedback and ambiguity attitudes. Importantly, it is the case that a decision-maker's (partial) identification of the plausible statistical models does not hinge on (is independent of) preferences and ambiguity attitudes, i.e. it can be described relying solely on the game form. On the other hand, information types distributions q are needed to derive the objective distribution of messages for each player given the decision functions mapping types into strategies.

The payoff of a player i, $u_i : \Theta \times Z \to \mathbb{R}$, depends on the true information types of the players $(\theta_i)_{i\in I}$ and the terminal history $z \in Z$ that is reached at the end of game. The payoff function itself

[6]The notation is explained below. More detail can be found in Battigalli et al. (2023).

is the composition $u_i := v_i \circ g$ of an outcome function $g : \Theta \times Z \to Y$ and a Bernoulli (or vN-M) utility $v_i : \Theta_i \times Y \to \mathbb{R}$. The strategic form payoff function of each $i \in I$ is $U_i := u_i \circ \zeta : \Theta \times S \to \mathbb{R}$. We so obtain the strategic form $G = \mathcal{N}(\Gamma)$ of a finite[7] multistage game with payoff uncertainty and exogenous type distributions Γ:

$$(G, F) := \left\langle I, \Theta_0, q_0, (\Theta_i, q_i, S_i, M_i, F_i, U_i)_{i \in I} \right\rangle \tag{2}$$

The ensuing structure has an evident mathematical similarity to (simple) Bayesian games with exogenous type-independent beliefs. In this framework, however, it must be understood that the game with payoff uncertainty is simply endowed with exogenous information-type distributions.[8] Players are assumed to know I, S, M_i, F_i, U_i and, importantly, θ_i. The overall interactive situation is instead *not* assumed to be common knowledge among the players. Additionally, players will not always be assumed to be Bayesian. An important assumption that is maintained[9] in the following is that of Observed Payoffs:

Definition 1 (**Observed Payoffs**)**.** *Feedback function F_i satisfies Observed Payoffs relative to payoff function U_i if payoff only depends on feedback, that is, if*

$$\begin{aligned} &\forall (s_i, \theta_i) \in S_i \times \Theta_i, \forall (s'_{-i}, \theta'_{-i}), (s''_{-i}, \theta''_{-i}) \in S_{-i} \times \Theta_{-i}, \\ &\qquad F_{\theta_i, s_i}(\theta'_{-i}, s'_{-i}) = F_{\theta_i, s_i}(\theta''_{-i}, s''_{-i}) \Rightarrow U_{\theta_i, s_i}(\theta'_{-i}, s'_{-i}) = U_{\theta_i, s_i}(\theta''_{-i}, s''_{-i}) \end{aligned}$$

Finite game with feedback (G, F) satisfies Observed Payoffs if the feedback functions of every $i \in I$ do.

In simpler terms, this property of feedback requires that players are always able to discern their realized utility. As will be discussed below, this fundamental assumption underpins the monotonicity results in BCMM because it ensures measurability of the section of the payoff function at the strategy that has been played in the long-run. This fact extends seamlessly to payoff-uncertain settings. As per their discussion, while this assumption may be reasonable in many games of interest, it is not to be taken for granted *a priori.* For instance, if players had altruistic motivations with regard to others but could not directly observe their internal utility index, this assumption would be quite far-fetched. The description of the rules of the game is complete. Next I introduce some decision criteria used to deal with ambiguity attitudes. Through these, we will revisit the traditional, Bayesian definition of Self-Confirming Equilibrium under payoff uncertainty, allowing the analysis of broader patterns of behavior.

2.2. **Preferences and decision criteria**

The representation of behavior in uncertain environments requires a decision-maker (DM) to be able to rank acts that yield (possibly risky) consequences dependent upon each state of the world, whose probabilities are unknown. In our setting, an individual player of type $\theta_i \in \Theta_i$ from population $i \in I$ is faced with the choice of a strategy $s_i \in S_i$. The outcome of any such strategy is uncertain and depends on the choices and types of coplayers. The uncertainty space of this choice is thus described by the Cartesian Product of the sets of other agents' strategies, private information types, and residual uncertainty: $S_{-i} \times \Theta_{-i} \times \Theta_0$. The DM attaches to each strategy s_i a subjective

[7] All sets that are part of the structure are finite.

[8] The problem of terminology in Game Theory is a tricky one. Harsanyi (1967, 1968a,b) originally introduced the notion of Bayesian Games and Bayesian Equilibrium. His definition is appropriate to analyze game structures that contain a unique profile of prior beliefs over a relevant uncertainty space. In our context, the word "(non-)Bayesian" is used to denote the decision attitudes of these kinds of players. However, we shall not speak of a Bayesian game, because we allow for profiles of sets of priors, and therefore, if players are not ambiguity neutral, i.e. they do not evaluate events according to their predictive probability (cfr. Section 2.2), there is not an evident way to update their beliefs according to Bayes' rule.

[9] It will be stated explicitly whenever we do *not* make this assumption.

value, denoted $V_{\theta_i}(s_i)$, parametrized by her tastes, characteristics and beliefs. To directly relate the player's decision problem to a framework of decision under uncertainty, this outcome uncertainty space may be interpreted as the *grand state space* discussed e.g. in Gajdos et al. (2008) and Gilboa and Marinacci (2013). Beliefs regarding the state space are constrained solely by the objective information an agent possesses.

Savage (1954) axiomatized the well-known SEU criterion, which extended the classical Expected Utility framework used to analyze lotteries with known probabilities to the more general case of Knightian uncertainty, eliciting the DM's subjective probability (a belief over states of the world) by requiring that he be able to (consistently) rank bets. Applications of this model have proved incredibly prolific in Game Theory and Economics. An alternative axiomatization of the same decision criterion in a more flexible setting is due to Anscombe and Aumann (1963). However, the thought experiments proposed by Ellsberg (1961) challenged (the normative desirability of) the SEU criterion, and led to the formulation of decision models whereby, even assuming the DM had a unique probabilistic model (the equivalent of a lottery) in mind for the problem at hand, faint conviction in such model could drive her to avoid uncertainty. This critique conflicted with one of Savage's postulates, the "Sure Thing Principle", and correspondingly with the Independence axiom in the Anscombe-Aumann formulation.

Ellsberg's first experiment involves two urns, I and II, containing 100 balls each. All balls are either white or black; in Urn I there are 50 black and 50 white balls, whereas nothing is known about the composition of Urn II. A DM must choose an urn; then, a ball is drawn from each of the two urns, and the DM is given a prize if the ball extracted from the urn of her choice is black. After this, the balls are put back in their respective urn, and the DM must choose an urn again; the draw is repeated, but the new prize is given to her only if the ball is white. Oftentimes, the known bet of urn I is preferred in both cases, a phenomenon that is incompatible with the evaluation of acts through a unique, additive probability measure. Allowing for the decision patterns that emerged in Ellsberg's experiments necessarily requires abandoning at least some of the properties of the SEU representation. A variety of approaches have been proposed to deal with this type of decision-making. These include considering non-additive probabilities or *capacities*, allowing for multiple subjective probability measures, and incorporating context-dependent beliefs. I begin by introducing some decision models that explore these possibilities. The overview will provide support in favor of the specific decision criterion adopted for the game-theoretic analysis.

The first axiomatization of a non-Bayesian decision model, which does not impose additivity on the DM's elicited subjective probability was the Choquet Expected Utility (CEU) representation of Schmeidler (1986, 1989). By restricting the Independence axiom of the Anscombe-Aumann formulation to comonotonic acts, i.e. acts whose ranking is independent of the state of the world, all acts are evaluated by means of a subjective capacity, through Choquet integration (Choquet, 1953). Nonadditivity of capacities allows a DM to express mild confidence in the relative likelihood of events. Schmeidler also showed that, if a DM has preferences for hedging against uncertainty, the capacity itself is endowed with a peculiar convexity property, which makes Choquet integration over the capacity akin to minimization of the integral over the probability measures in the core. The core is the set of additive probabilities whose lower bound on any measurable event contained in the event algebra is given by the value of the capacity itself. There are, however, situations in which the objective information in possession of the DM lets her entertain a set of plausible prior beliefs that does not constitute the core of any capacity.[10]

Gilboa and Schmeidler (1989) further relaxed the Independence axiom, restricting it to mixing with lottery acts, i.e., acts whose consequences occur with known probabilities. Together with uncertainty aversion, this axiom leads to their Maxmin Expected Utility (MEU) representation. This model ranks acts according to their value in the worst-case scenario envisioned by the DM. In any decision environment, the DM deems possible some set of probability models, and this set is constrained by objective information. Let $\hat{\Sigma}$ represent such objective information. Then, in our

[10]See the example at p. 200 in Gilboa and Marinacci (2013).

game-theoretic setting, the DM can deem possible any $\hat{\pi} \in \hat{\Sigma} \subseteq \Delta(S_{-i} \times \Theta_{-i} \times \Theta_0)$. A frequentist MEU DM, who infers the probability of events strictly from their long-run relative frequency, and considers all possibilities that to her knowledge are compatible with her objective information, would select strategy $s_i \in S_i$ if

$$\forall s_i' \in S_i,\ \min_{\hat{\pi}\in\hat{\Sigma}} U_i(s_i, \theta_i, \hat{\pi}) \geq \min_{\hat{\pi}\in\hat{\Sigma}} U_i(s_i', \theta_i, \hat{\pi}) \tag{3}$$

where

$$U_i(s_i, \theta_i, \hat{\pi}) = \int_{\Theta_0\times\Theta_{-i}\times S_{-i}} U_i(s_i, \theta, s_{-i})\, \hat{\pi}(\mathrm{d}\theta_0, \mathrm{d}\theta_{-i}, \mathrm{d}s_{-i}) \tag{4}$$

is the Bernoulli utility of s_i under a risky but certain probability model $\hat{\pi}$ (a lottery). The above can often be reduced to

$$U_i(s_i, \theta_i, \hat{\pi}) = \sum_{(\theta_0,\theta_{-i},s_{-i})\in\Theta_0\times\Theta_{-i}\times S_{-i}} U_i(s_i, \theta, s_{-i})\, \hat{\pi}(\theta_0, \theta_{-i}, s_{-i}) \tag{5}$$

since typically only models characterized by finitely many consequences are considered. This is our case because game G is finite. When $\hat{\pi}$ is a Dirac measure $\delta(\theta_0, \theta_{-i}, s_{-i})$ on some deterministic profile, U_i is the vN-M utility of the certain outcome induced by that profile, which we have introduced as part of the game description in (2).

Criterion 3 is flexible and intuitive, if extreme, and motivates the MSCE definition of Section 2.3. It has been pointed out in the literature that it is not obvious why such a "paranoid" attitude shall characterize a DM who dislikes uncertainty. As stated by Gilboa and Marinacci (2013), however,

> "An individual who satisfies the axioms [of Gilboa and Schmeidler] can be thought of as if he or she entertained a set C of priors and maximized the minimal expected utility with respect to this set. Yet, this set of priors need not necessarily reflect the individual's knowledge. Rather, information and personal taste jointly determine the set C. Smaller sets may reflect both better information [or, we may add, more naïvete] and a less averse uncertainty attitude. "

The MEU model thus allows the DM to restrict her attention to a subset of the models which are not contradicted by objective information in her possession. If the set of priors entertained by the DM is some $C \subset \hat{\Sigma}$, she chooses $s_i \in S_i$ if

$$\forall s_i' \in S_i,\ \min_{\hat{\pi}\in C} U_i(s_i, \theta_i, \hat{\pi}) \geq \min_{\hat{\pi}\in C} U_i(s_i', \theta_i, \hat{\pi})$$

In the revealed preferences tradition, the MEU model doesn't require frequentist decision-making. As Gilboa et al. (2010) put it, an uncertainty-averse DM can convince others that decisions made through the Unanimity criterion (Bewley, 2002) are rational, while she cannot be convinced that she is wrong in making decisions by completing her preferences through the MEU criterion. Such completion through revealed preferences need not be frequentist in nature. There are, however, ways to more starkly separate the DM's attitude toward uncertainty from her perception of, and information regarding, uncertainty. The approach followed here is to assume information-type θ_i has SEU preferences over "second-order acts," i.e. expected utility profiles induced by the array of plausible objective models. This requires him to hold a second-order subjective probabilistic belief

$$\mu^i_{\theta_i} \in \Delta(\hat{\Sigma}) \subseteq \Delta^2(\Theta_0 \times \Theta_{-i} \times S_{-i})$$

over such plausible models.[11] Then, negative attitudes toward uncertainty may be described by a functional which "penalizes" more uncertain beliefs. Some belief μ is regarded as more uncertain

[11] As a matter of notation *and* interpretation, a conjecture of a player i may be common or differentiated across types and strategies chosen. For now, consider a level of generality whereby at least types may hold different conjectures, because this will be an important general concept for the later equilibrium analysis. That the conjecture may differ for players playing different strategies is something we informally take for granted here, whereas it will be stated explicitly when necessary (e.g. for equilibrium definitions).

(or uncertainty aware) than another ν if μ is a mean-preserving spread of ν (see Section 3). According to the smooth representation introduced by Klibanoff et al. (2005) [hereon, KMM], an uncertainty-averse DM evaluates available act (strategy) $s_i \in S_i$ under the subjective second-order prior according to criterion

$$V_{\theta_i}^{\varphi_i}\left(s_i, \mu_{\theta_i}^i\right) = \varphi_i^{-1}\left(\int_{\mathrm{supp}\mu_{\theta_i}^i} \varphi_i(U_i(s_i, \theta_i, \hat{\pi}))\ \mu_{\theta_i}^i(d\hat{\pi})\right) \tag{6}$$

where φ_i is a continuous, strictly increasing and concave[12] real-valued function. One of the significant advantages of criterion (6) is that it is able to separate a DM's ambiguity attitudes from her ambiguity perception. This tractability allows to analyze the two features separately. Attitudes toward risk and ambiguity are typically regarded as stable personal traits; while perception of uncertainty may depend on stable personal features as well, it does not seem far-fetched to suppose it is often affected by contextual factors.

The shape of function φ_i captures the player's ambiguity attitudes. The support of the second-order prior $\mu_{\theta_i}^i$ may be interpreted "cognitively" (Wald, 1949; Gilboa and Marinacci, 2013): a motivation for a smooth ambiguity representation may be the desire to represent the behavior of a DM who does *not* exclude any plausible model consistent with objective information from her decision problem, and nevertheless does not evaluate acts solely based on the worst-case scenario for each. This is, however, far from necessary for the representation; in fact, supports of beliefs upholding criterion (6) will be conveniently used to describe varying degrees of uncertainty perception in the game theoretic analysis of Section 3. That is, although objective information might *not* rule out some objective statistical model, the DM may unwillingly ignore it. Note that if $\mu_{\theta_i}^i$ is a Dirac measure on some objective statistical model over the uncertainty space, the DM behaves as a naïve Bayesian; that is, she *de facto* entertains a unique first-order prior.

The presence of operator φ_i^{-1} in (6) constitutes an innocuous normalization that will facilitate comparison of subjective values with risky but unambiguous alternatives. "Ambiguity neutrality" corresponds to the case where φ_i is affine: in this case, the player computes a SEU calculation with respect to the *predictive probability* of each profile of strategies and information types, given by

$$\hat{\mu}_{\theta_i}^i(\theta_0, \theta_{-i}, s_{-i}) = \int_{\mathrm{supp}\mu_{\theta_i}^i} \hat{\pi}(\theta_0, \theta_{-i}, s_{-i})\mu_{\theta_i}^i\,(\mathrm{d}\hat{\pi}) \tag{7}$$

This is no more the case for players who are even moderately averse to ambiguity; this key aspect will play a pivotal role in the analysis of Section 3. According to the KMM criterion, an agent i is characterized as *more ambiguity averse* than j if there exists a function h such that $\varphi_i = h \circ \varphi_j$, where h is strictly increasing and concave or, equivalently, when both φ_i and φ_j are twice continuously differentiable, if $\varphi_i''/\varphi_i' \leq \varphi_j''/\varphi_j'$. When the ambiguity aversion coefficient $-\varphi_i''/\varphi_i'$ converges to infinity, we obtain the limit case of *extreme* ambiguity aversion, denoted ω, corresponding to the Maxmin criterion:

$$V_{\theta_i}^{\omega}\left(s_i, \mu_{\theta_i}^i\right) = \min_{\hat{\pi} \in \mathrm{supp}\mu_{\theta_i}^i} U_i(s_i, \theta_i, \hat{\pi}) \tag{8}$$

Note that, in this case, the set $\mathrm{supp}\mu_{\theta_i}^i \subseteq \hat{\Sigma}$ does *not* reflect ambiguity attitudes—it only embodies objective information *and its processing* on the part of the DM. With this, we are now equipped to embark on the game-theoretic analysis. The models presented above constitute fundamental contributions to Decision Theory, as they axiomatize revealed preferences inconsistent with SEU within Rational Choice Theory. The survey in this Section is far from exhaustive; its scope is limited to the purpose of understanding uncertainty averse decision making, how it favors the *status quo*, and the normative interplay that exists among ambiguity attitudes (personal

[12] If the DM is uncertainty-loving instead, φ_i is convex.

tastes), knowledge, and perception. In later sections, only the KMM model will be used; the above exposition is, however, crucial to interpret correctly the results of Section 3.

In economic problems, usage of uncertainty-averse (or loving) decision criteria may heavily affect predictions, as they characterize individual incentives differently and, therefore, also equilibrium sets according to any solution concept that does not assume perfect monitoring. In other words, departure from the *Bayesian paradigm* allows to encompass broader choice patterns, which address in a descriptively consistent while normatively desirable way Knightian uncertainty. This is especially palatable given that (i) “unmeasurable uncertainty” is pervasive in the evolution of recurrent anonymous interaction phenomena, e.g., in the stock market, as pointed out by Keynes (1936), and (ii) Ellsberg’s Paradox, together with the ensuing academic discussion, has shown that neutrality toward uncertainty may not be necessarily regarded as the most normatively appropriate attitude.

2.3. Self-confirming equilibrium with ambiguity attitudes

The first definitions for Conjectural or Self-Confirming Equilibrium were put forth by Battigalli and Guaitoli (1988) and Fudenberg and Levine (1993). Both assumed the SEU decision criterion, and were suited for the analysis of games in extensive-form. Battigalli and Guaitoli had subsumed the results of their independent undergraduate theses, and their framework allowed for general feedback functions. On the other hand, while Fudenberg and Levine’s work more specifically assumed that players could perfectly observe ex-post the path of play, their analysis would apply to large population games, justifying the assumption that players exclusively focus on their current-stage payoff even if they are not impatient.

A general definition of anonymous-interaction Self-Confirming Equilibrium for games with feedback and payoff uncertainty, which incorporates the contributions above, and in which all players are SEU maximizers, can be adapted from Battigalli et al. (2023). In a Self-Confirming Equilibrium, players respond rationally to confirmed conjectures about the behavior of coplayers. Information feedback for each player i is described by the private feedback function $F_i : \Theta \times S \to M_i$.[13] A player of type θ_i playing strategy s_i who holds conjecture $\mu^i_{s_i,\theta_i}$ expects to receive each message m_i with probability

$$\mathbb{P}^{F_i}_{s_i,\mu^i}(m_i|\theta_i) := \sum_{\theta_0,\theta_{-i},s_{-i}:F_i(\theta,s)=m_i} \mu^i_{s_i,\theta_i}(\theta_0,\theta_{-i},s_{-i}) \tag{9}$$

whereas, if coplayers are collectively enacting mixed strategy profile σ, random matching and a standard law of large numbers argument imply that in a steady state the long-run observed frequency of each message is given by

$$\mathbb{P}^{F_i}_{s_i,\sigma_{-i},q_{-i}}(m_i|\theta_i) := \sum_{\theta_0,\theta_{-i},s_{-i}:F_i(\theta,s)=m_i} q_0(\theta_0)\prod_{j\neq i}\sigma_j(s_j|\theta_j)q_j(\theta_j) \tag{10}$$

To understand the relevance of these formulae, consider a game that has already been repeated several times, such that a stationary state has been reached. That is, we are in “equilibrium,” which I shall define explicitly shortly below. A player of type θ_i from population i has been playing s_i for a long time, and wants to determine whether it is sensible to keep doing so. I implicitly assume that players exhibit inertia, that is, unless they have a *strict* incentive to deviate, they do not do so.[14]

[13] To ease intuition and exposition, the notation is kept consistent throughout with the strategic-form framework introduced at the beginning of Section 2. It should be kept in mind, however, that the scope of the definition of Bayesian Self-Confirming Equilibrium is more general and adapts seamlessly to games in extensive-form.

[14] This assumption is standard in equilibrium analysis, and the reason should be intuitive: if we assumed the opposite, players may keep bouncing from one strategy to another. Even though this does not affect their own payoffs *directly*, it may dynamically alter coplayers’ incentives, leading to chaotic interaction. Although it is in principle totally legitimate to conduct such an analysis, this would be formally much more complex, and equilibria (which would then need to be defined differently) may fail to exist even in relatively simple games.

Through the long-run objective distribution of messages described by Equation 10, θ_i is able to partition the uncertainty space (cfr. Equation 1) and infer the relative frequency of (unobservable) subsets of the uncertainty space, which we also call *events*. At the same time, θ_i enters a play of the game holding a conjecture $\mu^i_{s_i,\theta_i}$ which describes what he believes of the uncertainty space and, in turn, of the relative frequency of the messages he can observe. Is the player free to hold *any* conjecture? In principle, yes, but a key assumption we make in Self-Confirming Equilibrium analysis is that players have sufficient cognitive ability not to hold beliefs that are inconsistent with the stationary distribution of observed messages.

Suppose you are playing Texas Hold 'em against a friend whom you suspect is a *cheater* (this would be his information type). In the beginning, you do not have evidence for or against this hypothesis, so you are free to trust him as you deem fit. This remains the case as long as what you observe does not definitively imply that he is either cheating or not. However, after a long time, you notice that he folds 10% of the time, and *every single time* he does not, he was "lucky" enough to have drawn Double Aces from the deck. Next time, would it not be foolish *not* to assume he has Double Aces in his hand? Assuming the deck is fair, the number of plays does not have to be very large for you to rule out that the observed message frequency (i.e., the outcome of the games he keeps winning) is inconsistent with him being of type *non-cheater*. This naturally leads us to the definition of (Bayesian) Self-Confirming Equilibrium.

A conjecture is confirmed in the long run if it is not disproven, that is, if the subjective probability assigned to the event of observing any given message corresponds to its observed frequency. When a player's conjecture is confirmed, she has no incentive to deviate to another strategy, if she is acting optimally already. These facts provide the two conditions for stationarity that uphold equilibrium.

Definition 2 (SCE)**.** *Consider a finite game G and a profile of feedback functions $F = (F_i)_{i\in I}$. A profile of mixed strategies and conjectures*

$$\left(\sigma_i^*(\cdot|\theta_i), \left(\mu^i_{s_i,\theta_i}\right)_{s_i\in supp\sigma_i(\cdot|\theta_i)}\right)_{i\in I,\theta_i\in\Theta_i} \in \bigtimes_{i\in I}\left(\Delta(S_i)\times\Delta^{S_i}(S_{-i})\right)^{\Theta_i}$$

*is an **anonymous self-confirming equilibrium** (SCE) of the game with feedback (G, F) if, for every role $i \in I$, information type $\theta_i \in \Theta_i$ and strategy $s_i \in supp\sigma_i^*$, the following conditions hold:*

1. *(rationality)* $\sigma_i(s_i|\theta_i) > 0 \Rightarrow \forall \tilde{s}_i \in S_i, U_i\left(s_i, \theta_i, \mu^i_{s_i,\theta_i}\right) \geq U_i\left(\tilde{s}_i, \theta_i, \mu^i_{s_i,\theta_i}\right)$
2. *(confirmation)* $\mathbb{P}^{F_i}_{s_i,\mu^i}(\cdot|\theta_i) = \mathbb{P}^{F_i}_{s_i,\sigma^*_{-i},q_{-i}}(\cdot|\theta_i)$

This definition allows for fractions of agents in the same population to hold different assessments depending on their information type and the strategies they are playing. This is important as the partition of the uncertainty space can differ depending on these parameters. Note that an equilibrium is identified as a profile of strategies *and beliefs*—the former alone only make up an *equilibrium (strategy) profile*. BCMM discuss the generalization of the SCE solution concept to non-neutral attitudes toward uncertainty, comparing sets of equilibria underpinned by increasing degrees of ambiguity aversion, assuming that players are KMM decision-makers. Their definition is easily generalized to a payoff-uncertain setting. Let $\nu^i := \mathbb{P}^{F_i}_{s_i,\sigma^*_{-i},q_{-i}}(\cdot|\theta_i) \in \Delta(M_i)$ denote the long-run distribution of messages observed by θ_i playing s_i within equilibrium strategy profile σ^*; then

$$\hat{\Sigma}_{-i,\theta_i} := \hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i}) = \left\{\hat{\pi} \in \Delta(\Theta_0 \times \Theta_{-i} \times S_{-i}) : \hat{F}_{s_i,\theta_i}(\hat{\pi}) := \hat{\pi} \circ F^{-1}_{s_i,\theta_i} = \nu^i\right\} \tag{11}$$

denotes the *partial identification set*, that is, the set of all distributions of profiles of strategies and types deemed possible from θ_i's long-run empiricist perspective. This is the hard information constraint faced by each DM, which in Section 2.2 we had more informally denoted by $\hat{\Sigma}$. Since the distributions contained in this set are undistinguishable by construction from θ_i's perspective, they form the domain of uncertainty of his or her decision problem. This fact motivates a more general definition of Self-Confirming Equilibrium, which encompasses the possibility of non-Bayesian DMs. Denote a game with feedback and ambiguity attitudes by (G, φ): this structure differs from (G, F) in that we add a profile $\varphi := (\varphi)_{i \in I}$ of ambiguity smoothing functions, in the sense of Equation 6.

Definition 3 (SSCE). *A profile of strategy distributions $\sigma^* = (\sigma^*_{\theta_i})_{\theta_i \in \Theta_i, i \in I}$ is a **smooth** Self-Confirming Equilibrium (SSCE) profile of game with feedback and ambiguity attitudes (G, φ) if, for every $i \in I$, $\theta_i \in \Theta_i$, and $s_i \in supp\sigma^*_{\theta_i}$, there exists μ_{s_i,θ_i} such that*

1. *(rationality)* $\forall \tilde{s}_i \in S_i, V^{\varphi_i}_{\theta_i}\left(s_i, \mu_{s_i,\theta_i}\right) \geq V^{\varphi_i}_{\theta_i}\left(\tilde{s}_i, \mu_{s_i,\theta_i}\right)$
2. *(confirmation)* $\mathrm{supp}\mu_{s_i,\theta_i} \subseteq \hat{\Sigma}_{-i,\theta_i}$

For a strategy profile to constitute a SSCE, there must exist a profile of assessments *consistent with long-run feedback* (that is the role of $\hat{\Sigma}_{-i,\theta_i}$) such that the subjective value of the game as evaluated through φ_i by each player is greatest when choosing that strategy. It is important to note that the existence of the profile of assessments here constitutes a necessary condition for a strategy profile to constitute an equilibrium, and it is not included as a component itself. In other words, the equilibrium profile *is* the equilibrium strategy profile. This feature is formally convenient to describe the monotonicity results of Section 3. For this reason, in the remainder, when I write "SCE" to denote an equilibrium set, I will also be referring to equilibrium strategy profiles only, leaving beliefs in the background.

This definition differs from BCMM in that the partially identified set of distributions $\hat{\Sigma}_{-i,\theta_i}$ is a subset of the larger space $\Delta\left(\Theta_0 \times \Theta_{-i} \times S_{-i}\right)$. The original analysis does not explicitly discuss payoff uncertainty.[15] BCMM refer to games with chance moves; the latter can be alternatively modeled as payoff parameters (see Section 3.5). At the same time, separately modeling payoff uncertainty allows to describe more flexibly the *ex ante* distribution of information in a game. One can more naturally interpret information types as private, stable features of individuals.[16] BCMM's results rely on an important Lemma upheld by the assumption of Observed Payoffs (Definition 1), extended here to enclose payoff uncertainty. First, note that we can define the objective expected payoff of playing strategy s_i for type θ_i when coplayers are behaving according to σ^*_{-i} as[17]

$$U_i(s_i, \theta_i, q_{-i}, \sigma^*_{-i}) := \sum_{(\theta_0,\theta_{-i},s_{-i}) \in \Theta_0 \times \Theta_{-i} \times S_{-i}} U_i\left(s_i, \theta, s_{-i}\right) q_0(\theta_0) \prod_{j \neq i} \sigma^*_j\left(s_j|\theta_j\right) q_j(\theta_j) \qquad (12)$$

Lemma 1. *If payoffs are observable in game G, then for every $i \in I, \theta_i \in \Theta_i, s_i \in S_i$ and $\sigma^*_{-i} \in \Delta(S_{-i})$,*

$$\forall \hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}\left(s_i, q_{-i}, \sigma^*_{-i}\right), \qquad U_i(s_i, \theta_i, \hat{\pi}) = U_i(s_i, \theta_i, q_{-i}, \sigma^*_{-i})$$

[15] If the game has no payoff uncertainty, it does not necessarily feature complete information, which would instead correspond to the informal assumption of the occurrence of event

$$\bigcap_{n \in \mathbb{N}_0} \text{(everybody knows that)}^n \text{ everybody knows all the rules of the game.}$$

[16] Modeling the distribution of information through chance moves would also introduce significant technical complications for SCE analysis. A game representation equivalent to the one above, if carried out through alternative paths of play induced by realizations of chance moves, would hinder clarity and require a fastidious description of what players can observe regarding those chance moves.

[17] Computation of expected payoff according to Equations 4 and 5 requires one-step integration as subjective models $\hat{\pi}$ are joint distributions of strategies and types. These models can be decomposed to obtain marginal beliefs on types, and conditional beliefs on strategies depending on types. Hence the notation abuse in Equation 12, where U_i takes in an extra argument, is moot.

In words, the objective expected payoff of playing strategy s_i is constant across all distributions in the partially identified set. This is quite natural for stationary distributions as, under payoff observability, a DM shall be able to rule out distributions inducing expected payoffs different from the long-run time average observed when playing s_i. BCMM's main result relies on this fact: as long as beliefs are consistent with evidence, their support and shape do not influence the value of the strategy being played—whereas they indeed (potentially) affect the value of untested strategies. The analysis in Section 3 also hinges on this important observation.

As a limit case of KMM, aversion to uncertainty takes the maxmin form of Gilboa and Schmeidler (1989)—although recall from the discussion in Section 2.2 that the set of priors is not to be interpreted in the same way. The following definition is provided separately, because the Maxmin criterion may admit a strictly frequentist interpretation, intended here, where revealed preferences are not the result of a set of priors:

Definition 4 (MSCE)**.** *A profile of strategy distributions $\sigma^* = (\sigma^*_{\theta_i})_{\theta_i \in \Theta_i, i \in I}$ is a **maxmin** self-confirming equilibrium (MSCE) profile of a game with feedback (G, F) if, for every $i \in I$, $\theta_i \in \Theta_i$ and $s^*_i \in supp\sigma^*_i$,*

$$\min_{\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s^*_i, q_{-i}, \sigma^*_{-i})} U_i(s^*_i, \theta_i, \hat{\pi}) \geq \min_{\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s^*_i, q_{-i}, \sigma^*_{-i})} U_i(s_i, \theta_i, \hat{\pi}) \ \forall s_i \in S_i \tag{13}$$

BCMM establish the following important monotonicity result: if a game with ambiguity attitudes is more ambiguity averse than another, in the sense that they share the same underlying structure but every player is weakly more ambiguity averse than the former, the set of smooth equilibria weakly expands. This surely implies that, in general, the set of SSCE is a weak superset of the set of SCE determined by SEU-maximizing players (Definition 2).

Recall that, in the long-run, the value of tested strategies is known objectively, whereas that of untested ones remains ambiguous. BCMM state that "keeping beliefs fixed, [the latter] is higher when ambiguity aversion is lower": the more a player is ambiguity averse, the more he exhibits status quo bias. As will be shown, a related complementary result holds. Keeping ambiguity aversion fixed, the value of untested strategies is in general lower the more uncertain the belief (the more the DM perceives ambiguity, or is troubled by the imprecision in the information she is presented with).

The definition of SSCE assigns an instrumental role to beliefs. The underlying motivation is that, in principle, one may entertain any belief; thus, it is best to consider all of them. Section 2.2 clarifies how, absent additional assumptions, a prior whose support is a strict subset of $\hat{\Sigma}_{-i,\theta_i}$ characterizes personal taste and characteristics in the processing of objective information, given the context. It is thus a stretch to assume that players may not be characterized in terms of their perception, and that we cannot compare agents who differ in this respect. It is sensible to discuss how non-Bayesian DMs may interact differently depending on their naïveté, as embodied by the width of the support of their second-order prior, *for a fixed degree of ambiguity aversion.* Thus in Section 3 an ancillary framework is developed to apply the SSCE solution concept to "doubtful" thinking. That is, I will consider how alternative beliefs entailing varying degrees of ambiguity perception impact steady states. The implications will highlight the effect that doubts early in the game may have on equilibrium.

2.4. **A note on payoff observability**

In the definitions of SSCE and MSCE, the subjective value optimization under some profile of conjectures is sufficient to ascertain that the candidate strategy profile is indeed an equilibrium, due to Lemma 1. Since any statistical model in the support of a conjecture shall yield the same subjective value as the conjecture that corresponds to the true distribution of strategies and types, it does not matter which conjecture is actually held within the set $\hat{\Sigma}_{-i,\theta_i}$ for determining the subjective value of the strategy that is currently being played. As pointed out by BCMM, however, the set

$$\left\{U_i(\tilde{s}_i, \theta_i, \hat{\pi}) : \hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s_i^*, q_{-i}, (\sigma^*_{\theta_k})_{\theta_k \in \Theta_k, k \neq i})\right\}$$

for a generic strategy $\tilde{s}_i$ distinct from s_i^* is in general not a singleton. This is because partitions of the uncertainty space induced by alternative pure strategies need not correspond for any given type—and this is what drives action in favor of the status quo.

Let $\hat{\sigma}$ be some Bayes-Nash strategy profile of the game. If payoffs are not observable, even in the long-run it is not granted that θ_i will learn the expected payoff associated with any $s_i \in \text{supp}\hat{\sigma}_{\theta_i}$. This implies that the role of uncertainty becomes all the more relevant in determining decision-makers' course of action, as the appeal of less ambiguous strategies inherently grows. In this case, the monotonicity result of BCMM does not hold.

For instance, consider the case of constant (no) feedback: $\hat{\Sigma}_{-i,\theta_i}$ necessarily corresponds to the entire uncertainty space. Even though players know their payoff function, they cannot observe their realizations so that, as more distributions become consistent with the evidence (the partition induced by feedback becomes coarser), more paranoid, uncertainty-averse DMs may envision a larger set of possibilities. More conservative courses of actions may then offer higher minimal expected payoff than all $s_i \in \text{supp}\hat{\sigma}_{\theta_i}$ also if coplayers were playing according to $\hat{\sigma}_{-i}$, because the expected payoff induced by $\hat{\sigma}_{-i}$ can never be observed. This is indeed the case of Example 1 below: without feedback, an extremely uncertainty averse DM will never enter the MP subgame. While it is desirable to keep this in mind, in what follows the assumption of Payoff Observability is always maintained.

3. COMPARATIVE AMBIGUITY PERCEPTION

In this section I extend the SCE analysis under uncertainty aversion to consider restrictions on exogenous beliefs, more precisely on profiles of second-order priors. Section 3.1 introduces an ancillary game structure to deal with such restrictions, and shows that priors predicting the same distribution on average may induce nested sets of equilibria if they entail varying degrees of (perceived) ambiguity. Section 3.2 illustrates this reasoning through an example, whereby the threshold of ambiguity attitude needed to justify more conservative courses of action depends on the spread of the prior. The example considers finite support measures for instructive purposes, but it is important to note that the result has a more general scope and, in some games, even very small probabilities assigned to more extreme events by priors with infinite support may affect the decision-making outcome. This important theoretical point is stated and discussed in Section 3.3. The implications of another relevant assumption on feedback, namely "own-strategy independence of feedback", are explored in Section 3.4. The last subsection instead considers how payoff uncertainty affects the formal description of the game.

3.1. Stochastic dominance of beliefs and equilibrium

Definition 5. *Consider a game with feedback and ambiguity attitudes (G, φ). Denote by (G, φ, μ) the μ-restriction of the game, that is, the same game with the additional restriction on exogenous beliefs whereby the belief profile must be described by μ. Denote by (G, φ, C) the C-restriction of the game, where $C := \bigtimes_{i \in I} C_i$ and C_i is the collection of beliefs that player $i \in I$ is allowed to hold in the restricted game.*

Definition 5 yields an ancillary class of games identified by prior profiles.[18] Clearly, a strategy profile σ^* can be a SSCE of $(G, \varphi, \tilde{\mu})$ if and only if it is a SSCE of (G, φ) such that for every i, θ_i

[18]The equilibrium set of a game restricted in this way may be empty even if mixed equilibria exist, because exogenously posed priors may be inconsistent with any Nash strategy profile.

and $s_i^* \in \text{supp}\sigma_i^*$, we have $\text{supp}\tilde{\mu}^i \subseteq \hat{\Sigma}_{-i,\theta_i}$ and

$$V_{\theta_i}^{\varphi_i}(s_i^*, \tilde{\mu}_{\theta_i,s_i^*}) \geq V_{\theta_i}^{\varphi_i}(s_i, \tilde{\mu}_{\theta_i,s_i^*}) \ \forall s_i \in S_i$$

Thus more broadly, σ^* is a SSCE of (G, φ, C) if and only if it is a SSCE of (G, φ) where the confirmation condition on the profile of beliefs μ of Definition 3 ($\text{supp}\mu^i_{s_i,\theta_i} \subseteq \hat{\Sigma}_{-i,\theta_i}$) is replaced with the stronger condition

$$\mu^i_{s_i,\theta_i} \in C_i \cap \Delta(\hat{\Sigma}_{-i,\theta_i})$$

This definition allows us to analyze how the equilibrium set may change depending on the prior belief profile held by players. To this end, the following result equips us with a comparative notion on priors that relies on stochastic dominance. The proof relies on the foundational work of Rothschild and Stiglitz (1970), and allows the derivation of Proposition 1 below.

Lemma 2. *If payoffs are observable, second-order prior $\mu^i \in \Delta(\hat{\Sigma}_{-i,\theta_i})$ is a (weak) mean-preserving spread of $\nu^i \in \Delta(\hat{\Sigma}_{-i,\theta_i})$ if and only if μ^i (weakly) second-order stochastically dominates ν^i.*

Probability distribution μ^i is said to be a mean-preserving spread of ν^i if, formally,[19] $\mu^i \overset{D}{=} \nu^i + z$ where z is random noise with null conditional expectation. In simpler terms, a mean-preserving spread of some distribution is another distribution with the same expected value which, however, assigns more weight to the tails of the original distribution. Note that mean-preserving spreads only make up a partial (incomplete) order. It can be shown that if μ^i is a mean-preserving spread of ν^i, then it has higher variance and identical expected value; the converse is not true in general, which can be seen by simply observing that variance induces a complete ordering. A special case of MPS obtains when the predictive probabilities of two beliefs coincide and the *convex hulls* of their supports are nested in each other.

Lemma 3. *If $\mu^i, \nu^i \in \Delta(\hat{\Sigma}_{-i,\theta_i})$ induce the same predictive probability, and* conv $\text{supp}\nu^i \subset$ conv $\text{supp}\mu^i$*, then μ^i is a mean-preserving spread (MPS) of ν^i.*

The notion of MPS has long been studied in the literature of choice under risk, motivated by the objective of defining what makes some lottery "riskier" than another. In this setting, we are dealing with beliefs rather than lotteries; thus an MPS of some belief shall be interpreted as the belief in the same average outcome which is, at the same time, more concerned with more uncertain outcomes.

Consistently with the interpretation in choice under risk, then, one may say that a belief is a MPS of another if it is "more uncertain", or it *perceives* more (Knightian) uncertainty. Indeed, KMM had already noted that ambiguity aversion is "an aversion to the subjective uncertainty about ex ante evaluations. Analogous to risk aversion, aversion to this uncertainty is taken to be the same as disliking a mean preserving spread" in the distribution of expected utility values induced by a belief and a strategy. They also state that a "useful comparative statics exercise is to hold ambiguity attitudes fixed and ask how the equilibrium is affected if the perceived ambiguity is varied", an exercise which we conduct through the following important result, which complements the analysis of BCMM:

Proposition 1. *Let (G, φ) be a finite game with feedback and ambiguity attitudes, which satisfies Observed Payoffs (Def. 1) and in which all players are weakly ambiguity averse. If μ and ν are profiles of beliefs, and for each i, θ_i and s_i, $\mu^i_{s_i,\theta_i}$ is a (weak) MPS of $\nu^i_{s_i,\theta_i}$, with $\text{supp}\mu^i_{s_i,\theta_i} \subseteq \hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})$ for every $\sigma^* \in SSCE(G, \varphi, \nu)$, then*

$$SSCE(G, \varphi, \nu) \subseteq SSCE(G, \varphi, \mu)$$

[19]Symbol $\overset{D}{=}$ means "is identical in distribution".

Proposition 1 establishes that greater ambiguity aversion need not be the unique determinant of "certainty traps" within the same game form. In fact, it is also possible that players who are equally averse to ambiguity are not all equally aware of the range of models compatible with the evidence, or they disregard some possibilities. To see this point more clearly, recall that the definition of SSCE admits beliefs whose support is strictly contained in $\hat{\Sigma}_{-i,\theta_i}$. Status quo biases, therefore, do not necessarily follow from a more ambiguity averse attitude; they may result from strong ambiguity *perception*.

How shall stochastic dominance be interpreted? It is wrong to conclude that an ambiguity averse decision-maker "prefers" holding beliefs that entail smaller uncertainty. It is behavior that elicits the representation, not the converse. In other words, the second-order prior held by a decision-maker in an interactive situation is symptomatic of the extent to which she would consider all available possibilities before making a choice. The stochastic dominance argument simply embodies the idea that the status quo bias is stronger for agents who are bewildered by the presence of uncertainty, irrespective of *how* averse to it they may be.

We are led to a startling paradox: provided that players slightly dislike uncertainty, outright ignoring it shall grant objectively better payoffs. Conversely, for any given prediction regarding the average outcome, frantically evaluating all possibilities, including the most extreme ones, increases the likelihood of certainty traps, even if relatively little weight is assigned to such extreme possibilities. The derivation of Proposition 1 allows a straight-forward corollary whereby wider ambiguity perception restricts the reachable equilibrium set for ambiguity-loving players.

The assumption of mean preservation means we require that players hold the same prediction regarding the behavior and types of others, while the spread implies that in one case they believe a wider range of outcomes to be possible. This affects their confidence in the prediction, so that the same average prediction may act as basis for action in different ways. Note that so far I have disregarded cases where the predicted outcome is different; this point will be addressed in Section 3.3. The KMM criterion is rather helpful in the derivation of Proposition 1, because it allows to deal with probabilities. Analogous reasoning would result, however, if we were to represent the DM's behavior e.g. through the CEU model (Schmeidler, 1986). Instead of comparing mean-preserving spreads, one may start from a capacity, and reduce the weight assigned to less uncertain objective distributions, i.e. those leading to expected utility values closer to the mean.

Proposition 1 complements the formulation and analysis conducted in BCMM of the status quo bias induced by aversion to uncertainty: when information is insufficient, for ambiguity averse players, the more contemplated possibilities the more equilibria are possible. Restricting the prior beliefs of players confines the comparison to beliefs that can be partially ordered. However, it enables a first evaluation of the consequences of players second-guessing themselves. When many priors are possible, an agent enters the game with a belief determined by her taste, objective information, and how the latter is processed, and acts accordingly. The evidence she accumulates thereafter depends on previous plays. Ex-ante more paranoid DMs, those who incur a certainty crisis, stop learning and get stuck more easily.[20] This is shown in the example of Section 3.2.

3.2. An illustration

Example 1. Consider the Matching Pennies game with an Outside option whose "tree" is depicted in Figure 1. Assume $\varepsilon \in \left(0, \frac{1}{2}\right)$. The game is a simplification of the running example in BCMM.

In this game without payoff uncertainty, player **A** repeatedly decides whether to leave the game (O) or enter and play matching pennies (MP) against Nature, which selects actions h, t according to a deterministic probability law $\sigma_0 \in \Delta(\{h, t\})$. Assume $F_{\mathbf{A}} = U_{\mathbf{A}}$, that is, **A**'s feedback coincides with her payoff function. If **A** plays **H** or **T**, in the long-run she learns her expected payoff and, given the game structure, also the exact distribution σ_0. However, if her initial belief is compatible with choosing strategy O, she may not learn anything and get stuck in the objectively

[20]For a rigorous argument on learning dynamics, see Battigalli et al., 2019b.

FIGURE 1: Matching Pennies with Outside Option

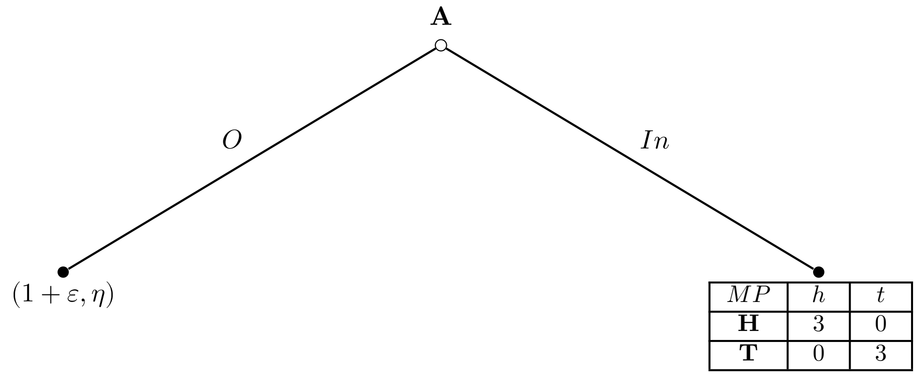

worse outcome. As in BCMM, this can never be the case for a SEU-maximizing player since for any conjecture $\overline{\mu}$ the subgame MP would necessarily have value

$$\max\left\{3\overline{\mu}, 3\left(1-\overline{\mu}\right)\right\} \geq 1.5 > 1+\varepsilon$$

Now consider $\mathbf{A}$'s decision-making for different priors. Let $\varphi_{\mathbf{A}} = U^{1/\alpha}, \alpha > 0$. For $\alpha \in (0,1)$, $\mathbf{A}$ would be ambiguity loving; for $\alpha = 1$, she'd be ambiguity neutral; otherwise, she is ambiguity averse. In general, the higher the value of α, the more $\varphi_{\mathbf{A}}$ is ambiguity averse.[21] Also let

$$\mu_{\mathbf{A},1} := \delta_{\frac{1}{2}\delta_h + \frac{1}{2}\delta_t}$$

$$\mu_{\mathbf{A},2} := \frac{1}{4}\delta_h + \frac{1}{2}\delta_{\frac{1}{2}\delta_h + \frac{1}{2}\delta_t} + \frac{1}{4}\delta_t$$

$$\mu_{\mathbf{A},3} := \frac{1}{2}\delta_h + \frac{1}{2}\delta_t$$

In words, if $\mathbf{A}$ believes $\mu_{\mathbf{A},1}$, she is certain that half the times Nature chooses h, and the other half it chooses t. If she believes $\mu_{\mathbf{A},2}$, she additionally deems it fairly plausible that either it always chooses h or always chooses t. If she believes $\mu_{\mathbf{A},3}$, she only believes these latter possibilities, and she assigns equal probability to them. Strategy O is a best-reply only if $\max\left\{V(H,\mu_{\mathbf{A}}), V(T,\mu_{\mathbf{A}})\right\} \leq 1$. Note that, by construction,

$$\hat{\mu}_{\mathbf{A},1} = \hat{\mu}_{\mathbf{A},2} = \hat{\mu}_{\mathbf{A},3} = \mu_{\mathbf{A},1}$$

where $\hat{\mu}$ is the predictive probability in the sense of Equation 7. Hence the expected utility profile for an ambiguity neutral agent would be the same for all three measures: on average, she expects the same behavior of Nature. For $\mu_{\mathbf{A},1}$, we have that

$$V_{\mathbf{A}}^{\varphi}\left(s_{\mathbf{A}}, \mu_{\mathbf{A},1}\right) = U_{\mathbf{A}}\left(s_{\mathbf{A}}, \mu_{\mathbf{A},1}\right) \;\; \forall s_{\mathbf{A}} \in S_{\mathbf{A}}$$

so that ambiguity attitudes have no bite and O can never be justified as a best-reply, because a prior with singleton-support believes no uncertainty is involved in the choice. On the other hand, observe that

$$V_{\mathbf{A}}^{\varphi}\left(\mathbf{H}, \mu_{\mathbf{A},2}\right) = V_{\mathbf{A}}^{\varphi}\left(\mathbf{T}, \mu_{\mathbf{A},2}\right) = \left(\frac{1}{4}3^{1/\alpha} + \frac{1}{2}1.5^{1/\alpha}\right)^{\alpha} \leq 1 \iff \alpha \geq 2.3 > \log_2(3)$$

[21] Note that this shape for $\varphi_{\mathbf{A}}$ entails a nonconstant degree of ambiguity aversion; see KMM for details.

and

$$V_{\mathbf{A}}^{\varphi}\left(\mathbf{H}, \mu_{\mathbf{A},3}\right) = V_{\mathbf{A}}^{\varphi}\left(\mathbf{T}, \mu_{\mathbf{A},3}\right) = \left(\frac{1}{2}3^{1/\alpha} + \frac{1}{2}0^{1/\alpha}\right)^{\alpha} \leq 1 \iff \alpha \geq \log_2(3)$$

By comparing $\mu_{\mathbf{A},1}$ with the other beliefs, we can observe that while the expected utility profile as expressed by the predictive probability is unchanged (and equal to $\mathbb{E}_{\mu_{\mathbf{A},1}}[U(In)] = 1.5$), for a sufficiently ambiguity averse player **A** these beliefs, which are mean-preserving spreads of the former, increase the likelihood of choosing the outside option. By comparing $\mu_{\mathbf{A},2}$ with $\mu_{\mathbf{A},3}$, we note that the latter *reduces* the support of the prior in terms of cardinality, but it does so while still increasing uncertainty in the payoff, because it assigns zero probability to the certain (in the sense of an objective lottery) outcome. Indeed, a moderately uncertainty averse **A** with $\alpha \in \left(\log_2(3), 2.3\right)$ may select O as a best reply if she held belief $\mu_{\mathbf{A},3}$, while she would not if she held $\mu_{\mathbf{A},2}$. $\triangleleft$

Note that all beliefs considered in the example are compatible with evidence observed by playing O. If **A** considers a wider range of possibilities, as long as she is sufficiently uncertainty averse she will avoid exploring even if she expects the same outcome on average. A greater perception of the threat presented by ambiguity makes the decision-maker more likely to get stuck into a suboptimal decision. In the example, it is especially interesting to confront beliefs whose supports are such that one is a superset of the other, as was the case for $\mu_{\mathbf{A},1}$ and $\mu_{\mathbf{A},2}$. Such a case can be more naturally interpreted as increased awareness of more remote possibilities on the part of the DM, rather than a belief envisioning exclusively different scenarios that just so happens to preserve the mean prediction.

It is also important to note that the objective value of the game given available information has not decreased, and that uncertainty aversion is present in the model in the same degree; it just is possible for it not to show and stay *hidden* as long as assessments are naïve enough, i.e. they are relatively unaware of the uncertainty involved in the choice. On the other hand, one can verify that the status of some strategy profiles as equilibria may not be affected by ambiguity perception. Without loss of generality, let $\sigma_0(h) > 1/2$ in the example and consider strategy $In.\mathbf{H}$; by playing the latter, **A** learns $\sigma_0(h)$ with certainty, and no ambiguity is present in her decision problem anymore, so that she may not hold conjectures that entail varying degrees of uncertainty.

The example clarifies that, if players revise their prior beliefs in a way that does not alter predictions but is more aware of the underlying uncertainty, it may be easier to justify strategies and reach SSCE profiles which would otherwise be avoided. It highlights how the possibility of doubting own beliefs may bring about less ambiguous, though possibly suboptimal, equilibria.

3.3. **Generalization to sets of beliefs**

It is possible to derive a result on equilibrium inclusion which imposes weaker restrictions on exogenous beliefs, that is, where we compare nonsingleton sets of prior beliefs connected through second-order stochastic dominance. This is desirable as the assumptions of Proposition 1 might appear somewhat stringent. Recall that stochastic dominance only induces a partial order; therefore, we cannot compare beliefs which are not in such relation with each other. However, we may leverage pairwise relations among beliefs that belong to different sets.

Recall that for any subjective belief μ^i with nonsingleton support, we may find the associated predictive probability $\hat{\mu}^i$ through Equation 7. If we simply compare nested sets of beliefs, it is relatively trivial to conclude that the induced sets of justifiable strategies shall enlarge irrespectively of ambiguity attitudes. Thus we want to compare beliefs that can be put into the MPS relation with each other, i.e. vary in their perceived uncertainty, without imposing or forbidding other features of beliefs. In this respect, the convexity of $\hat{\Sigma}_{-i,\theta_i}$ is relevant to determine whether we can find in general a valid (in the sense of compatible with evidence) belief with support convex hull different from μ^i that yields the same predictive probability $\hat{\mu}^i$, so that moving from a set to another we are only adding beliefs that are second-order stochastically dominated by the rest. The potential non-convexity of $\hat{\Sigma}_{-i,\theta_i}$ is represented by the dashed square in Figure 2. Again following BCMM, we may grant convexity by assuming (informally) that players are totally ignorant about the matching process; that is, no one is aware that other players are randomly drawn *independently of each other*. This assumption is unimportant for the case of two-person games, as is the case for the game in Section 4.

FIGURE 2: An intuitive graphical representation of the partial identification set and supports of beliefs which belong to R_ε and R_ε^γ. m_k with $k = 1, 2, 3$ denote predictive probabilities.

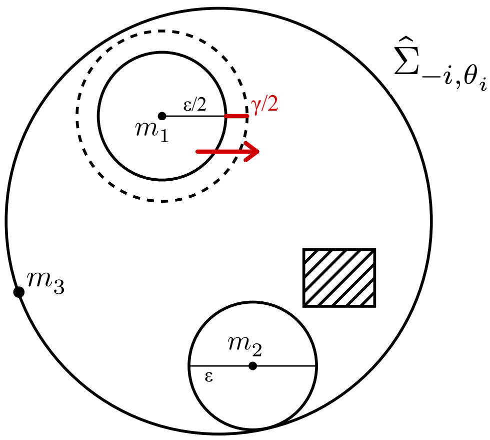

The maximal degree of generality can be obtained by considering sets of beliefs where perceived uncertainty is varied *wherever possible*. In other words, we compare sets where we replace any belief that admits an MPS compatible with objective information with that MPS, and leave unchanged beliefs that do not allow so. This procedure avoids outright excluding some plausible expectations on the game play: note that the strategy profiles in $SCE\,(G, F)$ coincide with the set of $SSCE$ given by all singleton-support beliefs. Consider a belief with singleton support lying on the boundary of $\hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})$ for some $\sigma^*, s_i \in \mathrm{supp}\sigma^*_i$: there is no possibility to construct a MPS whose support lies within $\hat{\Sigma}_{-i,\theta_i}$, as the spread would necessarily involve including models from outside the set.

While dealing with infinite-support beliefs can be demanding, we may rely on Lemma 3, considering the support on the space of objective expected utility values induced by statistical models in $\Delta\,(\Theta_0 \times \Theta_{-i} \times S_{-i})$. Endow this space with the following topological structure: let objective models be subjectively related by the maximal L^1-distance on the space of expected utility values, which is a subset of the real line:

$$\forall p, r \in \Delta(\Theta_0 \times \Theta_{-i} \times S_{-i}), d_{\theta_i}(p, r) = \sup_{s_i \in S_i} \left| \int U_i(s_i, \theta_i, \cdot)\ \mathrm{d}p - \int U_i(s_i, \theta_i, \cdot)\ \mathrm{d}r \right|$$

Then, given $\varepsilon \in \mathbb{R}_+$, we can define a set R_ε as the smallest closed set of beliefs whose supports are composed of objective models with maximum pair-wise distance (diameter) ε.

$$R_\varepsilon := \bigtimes_{i \in I} \left\{ \mu^i \in \Delta^2\,(\Theta_0 \times \Theta_{-i} \times S_{-i}) : \forall \theta_i, \forall p, r \in \mathrm{supp}\mu^i_{\theta_i}, d_{\theta_i}(p, r) \leq \varepsilon \right\}$$

The reader can verify that set R_ε is closed even when the partial identification sets are non-convex, because they are nonetheless compact-valued (BCMM). By increasing ε, we add profiles of beliefs with larger support convex hulls and which are necessarily MPSs of beliefs that are already present, provided that $\hat{\Sigma}_{-i,\theta_i}$ is always convex. As an auxiliary tool, let $R^*_\varepsilon \subseteq R_\varepsilon$ denote the subset whose beliefs already span the maximal convex hull compatible with objective information for every equilibrium. That is, R^*_ε includes each belief profile ν in R_ε such that no equilibrium $\hat{\sigma}$ exists under ν which would allow, for at least one i, θ_i and $s_i \in \text{supp}\hat{\sigma}_i$, constructing an MPS by strictly enlarging to a width larger than ε the convex hull of $\nu^i_{s_i,\theta_i}$ within the information constraint $\hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})$ consistent with equilibrium:

Note that this set is independent of specific values of γ when $\hat{\Sigma}_{-i,\theta_i}$ is convex. Finally, for any $\gamma \geq 0$ let $R^\gamma_\varepsilon := \left(R_{\varepsilon+\gamma} \backslash R_\varepsilon\right) \cup R^*_\varepsilon$. This last set serves the objective described above of varying perceived uncertainty wherever possible, without restricting the range of predictions an agent may hold. We can now state the main result of the paper: the stochastic dominance relation leads to weakly larger equilibrium sets when perceived uncertainty grows.[22]

Theorem 1. *Let (G, φ) be a finite game with feedback and ambiguity attitudes, which satisfies Observed Payoffs and in which all players are ignorant about the matching process and weakly averse to ambiguity. Then,*

$$\forall \varepsilon, \gamma \in \mathbb{R}_+, \qquad SSCE\,(G, \varphi, R_\varepsilon) \subseteq SSCE\,(G, \varphi, R^\gamma_\varepsilon) \subseteq SSCE\,(G, \varphi)$$

3.4. **Own-strategy independence of feedback**

An important, nontrivial assumption often considered when studying self-confirming equilibria is that of Own-Strategy Independence (cfr. Section 2.1 for notation):

Definition 6. *Feedback function F_i satisfies Own-Strategy Independence of Feedback about others (OSI) if, for every θ_i,*[23]

$$\forall s'_i, s''_i \in S_i, \mathcal{F}_{s'_i,\theta_i} = \mathcal{F}_{s''_i,\theta_i} \tag{14}$$

Game with feedback (G, F) satisfies Own-Strategy Independence if condition (14) holds for each player $i \in I$.

In words, i's feedback function satisfies this property if the partition of the uncertainty space it induces is independent of which strategy i actually chooses, for every type of player in that population. The learning patterns do not differ depending on one's own strategy, for one's given own type. In other words, changing from one strategy to another in between plays of the game reveals nothing more and nothing less about the behavior and types of coplayers. If payoffs are not observable, this property raises the incentives for ambiguity averse players to choose ex-ante less ambiguous options. It is natural to ask, then, what happens when it is combined with Observed Payoffs.

Lemma 4. *If (G, φ) satisfies Definitions 1 and 6, then for every $\varepsilon \in \overline{\mathbb{R}}_+$,*

$$BNE = SCE = SSCE(G, \varphi, R_\varepsilon) = MSCE$$

where BNE denotes the set of Bayes-Nash Equilibria.

As in BCMM, if OSI holds together with Observed Payoffs, in the long-run every player learns the expected payoff profile associated with every strategy, not only the one they have been playing. That is, set

$$\left\{U_i(\tilde{s}_i, \theta_i, \hat{\pi}) : \hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s^*_i, q_{-i}, (\sigma^*_{\theta_k})_{\theta_k \in \Theta_k, k \neq i})\right\}$$

[22] One can easily show that Theorem 1 holds with equality if players are not averse to ambiguity.
[23] A less stringent formulation requires condition 14 to hold only for non-dominated strategies of each type θ_i in each role i, since rational players never pick dominated strategies irrespective of their beliefs and information.

is a singleton for every $\tilde{s}_i$, because all strategies induce the same partition of the uncertainty space. As a result, the role of prior beliefs becomes immaterial. The set of (S)SCE collapses to that of BNE. The following example shows this point, i.e. that there are no implications to alternative prior specifications in this class of games.

Example 2. Figure 3 represents the strategic form of a game with feedback and ambiguity attitudes. Consider only pure strategy profiles for simplicity. The unique pure BNE is $(\mathbf{M}, \ell)$. Let again $F_i = U_i$; the game thus satisfies both Observed Payoffs and OSI, as can be checked by the payoff structure. Given that Colin's feedback is perfect and independent of his strategy, $\mu_c = \delta_{s_{\mathbf{R}}}$ is necessary: no uncertainty is involved for him. Whenever Colin (c) is playing ℓ or m, Rowena ($\mathbf{R}$) may be uncertain of which; any belief encompassing such strategies yields the same (feedback and) payoff for $\mathbf{M}$, that is 3. Irrespectively of her beliefs, however, she knows that as long as c plays ℓ or m, her expected payoff profile is constant for every action she may play. In other words, whatever strategy she decides to play in the long-run, the expected payoff of alternative strategies involves no more uncertainty. Thus the support of her belief cannot induce any status quo in favor of any strategy that has been played for a long time. ◁

FIGURE 3: A game where stochastic dominance of priors does not matter.

$\mathbf{R}$ \ c	ℓ	m	r
U	$\mathbf{2}, \mathit{0}$	$\mathbf{2}, \mathit{1}$	$\mathbf{1}, \mathit{0}$
M	$\mathbf{3}, \mathit{1}$	$\mathbf{3}, \mathit{0}$	$\mathbf{0}, \mathit{1}$
D	$\mathbf{0}, -\mathit{1}$	$\mathbf{0}, -\mathit{1}$	$-\mathbf{1}, \mathit{2}$

3.5. **The role of payoff uncertainty**

In the examples made so far, payoff uncertainty has not been given a pivotal role, mostly for ease of exposition. In general, however, adding uncertain parameters plays an important role in SCE analysis. For any given feedback structure, the fewer elements known by the players, the more distributions are consistent with the evidence. Thus, uncertainty cannot but grow (though possibly weakly). As the discussion shall have made clear by now, this inherently reinforces the status quo bias implied by ambiguity aversion when players are sufficiently aware of it. This is shown in the following example.

Example 3. In the game of Figure 4, let $F_i = U_i$ for $i = 1, 2$, and $\theta = 0.1$. Suppose first that θ is private knowledge in possession of player 2. The game then satisfies Observed Payoffs and Own-Strategy Independence, thus has a unique smooth SCE, corresponding to the unique pure Bayesian Equilibrium, (T, ℓ). It is not possible to sustain (B, r) as a SCE since Player 2 will always want to deviate, knowing he can gain for certain.

FIGURE 4: A simple game with payoff uncertainty.

1 \ 2	ℓ	r
T	$1, -1$	$-1, -2$
B	$-1, \theta$	$0, 0$

Suppose now instead that θ is private knowledge possessed by player 1, and player 2 is only told (reliably) that $\theta \in [-0.5, 1]$. The game does not satisfy Own-Strategy Independence anymore. With a uniform conjecture over all available hypotheses, expected utility is 0.25, but if $\varphi_2(u) = -\frac{1}{\alpha}e^{-\alpha u}$ with $\alpha > 3.1$, it can be verified that the private value of the deviation to Player 2 is negative. Hence (B, r) is a SSCE of this game, given this ambiguity attitude. Interestingly, note that the status quo bias here ends up favoring Player 2, since this equilibrium is subjectively better than (T, ℓ), and profile (B, ℓ) cannot be sustained. Sometimes, certainty traps may turn out to be beneficial to some parties involved. ◁

Many games of interest with no payoff uncertainty can be put into a formal equivalence with games with payoff uncertainty, so the latter does not always add significant insights into the analysis. The literature of choice under uncertainty has often drawn connections with games played against a malevolent nature. Indeed, games against nature can easily be reduced to one-player, payoff uncertain games.

Example 4. Go back to the game of Example 1. The same game against Nature can be thought of as the one-player game where the utility profiles of the subgame MP are determined by the residual uncertainty parameter $\theta_0 \in \{\theta_0', \theta_0''\}$, and $U_{\mathbf{A}}(In.\mathbf{H}, \theta_0') = U_{\mathbf{A}}(In.\mathbf{T}, \theta_0'') = 2, U_{\mathbf{A}}(In.\mathbf{H}, \theta_0'') = U_{\mathbf{A}}(In.\mathbf{T}, \theta_0') = 1$. Adapting the reasoning on priors to this payoff uncertain setting does not alter the results, as long as feedback is such that realized utility is observed, but an off-path realization of θ_0 is not. ◁

It must be kept in mind, however, that while in complex games substituting a payoff uncertain representation with one that only features moves of chance and no payoff uncertainty may simplify the initial mathematical structure, it would also require more intricate descriptions of players' knowledge of the rules of the game and feedback. In Example 3, if θ were determined through a fictitious move of chance $q_0(\theta = 0.1) = 1$, in order to achieve a formal equivalence of the analysis with the initial interpretation as private knowledge of Player 2, it would be necessary to introduce a multistage structure, and it would not make sense to analyze the strategic form of the game. This is because Player 2 should always have knowledge of the actual value of the parameter before choosing her strategy.

4. AN APPLICATION

In this section, I present a simple application of the above framework to a contracting game, highlighting how certainty crises and non-Bayesian decision making may lead to inefficient outcomes. The game purposefully exhibits strong links with models of Adverse Selection and Moral Hazard (Akerlof, 1995; Mirrlees, 1999). Traditional versions of these theories typically rely on Bayes-Nash Equilibrium analysis. This application instead serves to appreciate the implications of Self-Confirming Equilibrium and ambiguity aversion for negotiation. Strategic interaction in the context of surplus sharing is also the fundamental feature of bargaining problems. A survey of models and equilibrium analysis (particularly, subgame-perfect equilibrium) for bargaining problems is provided by Osborne and Rubinstein (1991).

Section 4.1 describes the context of the game. Variables are introduced, together with specific numerical values, which are then used to clarify key insights. To this end, an example is conducted in Section 4.2, together with a more general comparative statics exercise, where Theorem 1 is applied to identify bounds on perception of ambiguity which may prevent certainty traps, as functions of the original variables.

4.1. Contract game description

The setting echoes the illustration of Section 3.2. A *thick* job market is composed of two large populations: Workers $\mathcal{W}$ and Firms $\mathfrak{F}$. At every repetition of the game, some $W \in \mathcal{W}$ and $F \in \mathfrak{F}$

are matched at random. W moves first, and must decide whether to apply for a job (J) at F at cost $c = 1$,[24] or not (NJ). If he applies, then he must choose whether to exert effort (E) at private cost $e = 3$, or shirk (S).

Simultaneously[25], the firm F decides whether to Collaborate (C) with the worker, which involves upfront cost $m = 4$, or Not (NC), which is costless. If F does Not collaborate, worker productivity is stale, and independent of effort. In this case, the worker receives a fixed payment schedule, making shirking optimal. Collaboration instead provides advanced technology, that crucially becomes productive *only if the worker has previously exerted effort.* Delivery of this technology occurs near the end of the contract, ensuring that W must commit to his effort level beforehand, without knowing F's choice. The payment structure in case of Collaboration is designed to incentivize effort: shirkers (S) receive no payment, while diligent workers (E) are compensated with a bonus above their effort cost.

Let R, P denote the Revenue and Payment (wage) functions, respectively. For notational convenience, also denote the share of revenues appropriated by the Firm by $\Pi_s := R_s - P_s$, where s denotes the (reduced) strategy profile. For the numerical example, assume the following specifications:

$$R(s) := \begin{cases} 12 & s = (J.E, C) \\ R_0 & s = (NJ, \cdot) \\ 3 & \text{otherwise} \end{cases} \qquad\qquad P(s) := \begin{cases} 6 & s = (J.E, C) \\ 3 & s_F = NC \\ 0 & \text{otherwise} \end{cases}$$

Let feedback be given by[26] $\overline{F}_i(s) = (R_s, P_s)$ for $i \in \{W, F\}$. One can quickly verify that this game satisfies Observed Payoffs in the sense of Definition 1, while it does not satisfy Definition 6. This is sensible as firms observe revenues, reward workers accordingly, and workers observe their salary and effort. As stressed above, $m > \Pi_{S,C}$, that is, if F recognizes a shirker then it does not reward him, but nonetheless incurs a loss. Aside from the necessary condition of Firm Cooperation to achieve higher profits, other assumptions on payoffs are consistent with typical models of Moral hazard. The firm is concerned with the level of effort on the part of the worker, and without monitoring technologies it cannot ex-ante ensure itself against shirking.

Let us introduce an additional strategy for the Firm: Cooperation with Insurance (I). If F wishes to cooperate, but is sufficiently worried that the worker might shirk, it can stipulate a parallel insurance contract[27] at premium $p = 3$, with potential payout $b = 6 > p$ in case of low revenue. By choosing I, worker compensation is the same as in the case of cooperation, while Firm profits change. If W exerts effort, the firm bears the insurance cost and incurs a loss equal to $\Pi_{E,C} - m - p = -1$. On the other hand, in case of shirking the insurance payout offsets losses and the Firm obtains $\Pi_{S,C} - p + b = 2$. In principle, insurance allows Firms to always achieve greater expected profits than by choosing NC. The tree of the game is shown in Figure 5, while Figure 6 depicts the strategic form, where numerical values for all relevant variables have been substituted in for clearer interpretation of the example.

The key strategic challenge is mutual commitment. Neither party can observe the other's choice in advance. The firm only wants to collaborate when confident of worker effort, while the worker would always exert effort if certain of collaboration. Post-contract, the firm can verify technology usage, creating a monitoring mechanism that aligns incentives. If the worker shirks, the advanced technology becomes worthless, and the firm incurs additional losses—even without

[24]The application cost c can be alternatively interpreted as reservation utility of the worker.

[25]Simultaneity is in fact non-essential and shall here be interpreted as a technical property of the rules of the game; it is key, however, that W's choice cannot be observed by F.

[26]I use $\overline{F}$ to denote feedback in this game in order to distinguish it from player index F used for Firms.

[27]This assumption allows flexible interpretation: in reality, firms may devise a number of financial and nonfinancial strategies to insure themselves against the possibility of hiring a bad employee. One simple strategy we can imagine is conducting parallel hiring processes at a cost to be able to replace the departing employee more quickly.

FIGURE 5: Contract game tree.

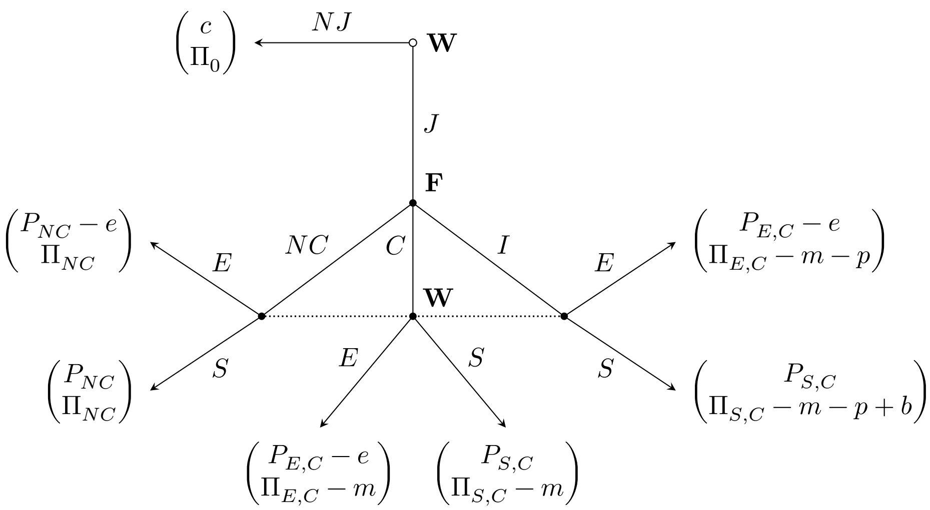

FIGURE 6: Strategic form of the Contract game.

W \ F	NC	C	I
NJ	$1, \Pi_0$	$1, \Pi_0$	$1, \Pi_0$
$J.E$	$0, 0$	$3, 2$	$3, -1$
$J.S$	$3, 0$	$0, -1$	$0, 2$

paying the worker. After the negotiation, the contract is carried out, W and F are separated, and the procedure is repeated.

The game is designed to capture *relational contracting*, a fundamental coordination challenge in labor markets.[28] Both firms and workers face incentives to make relation-specific investments only when they can credibly anticipate reciprocal commitment from their counterparty. The assumption that neither party can privately influence the market outcome adds a further layer to this bilateral hold-up problem, which characterizes many medium-term employment relationships.

4.2. **Comparative statics and market shutdown**

To fix ideas, consider the numerical example, whose strategic-form representation can be seen in Figure 6, and focus on SEU-maximizing players first. As in Section 3.2, a Bayesian worker W would never choose NJ, since for any $\sigma_F \in \Delta(\{NC, C, I\})$:

$$\max_{s_W \in \{J.E, J.S\}} V_W(s_W, \sigma_F) = \max\{3\sigma_F(NC), 3(1 - \sigma_F(NC))\} \geq 1.5 > 1 = V_W(NJ, \sigma_F)$$

Showing that a Bayesian F has the possibility of insuring itself against shirkers and therefore it would never pick NC, regardless of its belief, is analogous. As a result workers, who do not pick NJ, perfectly observe that firms are always choosing either C or I. As the game repeats, by observing their payoffs they quickly learn that NC is not being chosen and thus it is optimal for them to exert effort. In turn, and also by observing their payoffs, firms can indirectly observe $J.E$, and the only rational reply is to choose C. Again, the same reasoning applies to firms.

[28]The issue of relational contracts and relation-specific investments has been widely discussed. For more articulate theoretical insights, see Crawford (1990) and Baker *et al.* (2002), and the references therein.

Whatever their conjecture, they always choose C or I. Although they initially may choose to stipulate the insurance, in the long run they would observe that all workers are exerting effort, and begin collaborating. Simply put, the game necessarily converges toward the efficient coordination outcome for both parties, i.e. $(J.E, C)$. This is the unique Bayes-Nash equilibrium of the game, and the unique SCE for Bayesian players. Note that, although insurance is never picked, its presence is key. Without it, equilibrium sets would be quite different: let (G, F) denote the game with feedback described in Section 4.1, and let $\overline{G}$ denote the modification of G where strategy I is removed from the choice set of players $F \in \mathfrak{F}$. Then

$$\begin{aligned} BNE(G) = SCE(G) = \{(J.E, C)\} \\ &\subseteq BNE(\overline{G}) = \left\{(J.S, NC), (J.E, C), \left(\frac{1}{3}\delta_{J.E} + \frac{2}{3}\delta_{J.S}, \frac{1}{2}\delta_{NC} + \frac{1}{2}\delta_{C}\right)\right\} \\ &\subseteq SCE(\overline{G}, F) = BNE(\overline{G}) \cup \{(J.E, \alpha\delta_{NC} + (1-\alpha)\delta_C) : \alpha < 0.5\} \\ &\qquad\qquad \cup \left\{\left(\beta\delta_{J.E} + (1-\beta)\delta_{J.S}, \frac{1}{2}\delta_{NC} + \frac{1}{2}\delta_{C}\right) : \beta > \frac{1}{3}\right\} \end{aligned}$$

Now turn to ambiguity averse players. The structure of payoffs is such that uncertainty averse agents may pick NJ: e.g. for $\varphi_W = U_W^{1/\alpha}$ and $\alpha > \log_2 3$, NJ is justifiable as a best reply to $\mu_W = \frac{1}{2}\delta_{NC} + \frac{1}{2}\delta_C$. The same applies to firms: if they are sufficiently ambiguity averse and aware, they may choose NC, thereby cancelling the efficacy of the presence of insurance. The set of $SSCE$ is thus much larger, including mixed profiles which assign positive probabilities to subjectively optimal, yet objectively suboptimal strategies. For instance, sets $\{NJ\} \times \Delta(S_F)$ and $\Delta(S_W \backslash \{J.E\}) \times \{NC\}$ are both subsets of $SSCE(G, \omega)$. Simply introducing ambiguity aversion into to the model can easily make insurance useless and lead to market shutdowns, in the form of chronic worker shortages or pathological underinvestment. Insurance is unhelpful in this case because ambiguity aversion will either lead workers not to seek employment at all, or firms to underestimate the objective support offered by the decoy strategy to reach the coordination equilibrium.

It is also possible to make Theorem 1 operational, and obtain conditions for the occurrence of certainty traps, i.e., equilibria that are only possible when players are sufficiently ambiguity averse. We will conduct this exercise from the perspective of a worker; the reasoning is exactly the same for firms. Recall that the principle of reduction of compound lotteries applies to risky decisions. To avoid market shutdown it is necessary that

$$\begin{aligned} c < V_{\mu_W}(J) &= \max\{V_{\mu_W}(J.E), V_{\mu_W}(J.S)\} \\ &= \max\left\{\varphi_W^{-1}\left(\sum_{\hat{\pi}\in\operatorname{supp}\mu_W} \varphi_W(P_{E,C} + \hat{\pi}(NC)(P_{NC} - P_{E,C}) - e)\mu_W(\hat{\pi})\right),\right. \\ &\qquad \left.\varphi_W^{-1}\left(\sum_{\hat{\pi}\in\operatorname{supp}\mu_W} \varphi_W\left(P_{S,C} + \hat{\pi}(NC)(P_{NC} - P_{S,C})\right)\mu_W(\hat{\pi})\right)\right\} \end{aligned}$$

This condition simply requires that the entry cost c that would be saved by choosing NJ be smaller than the certainty equivalent of the subjectively optimal choice. If we consider the maximal degree of uncertainty perception, i.e. a second-order belief evenly split among the most extreme outcomes compatible with objective evidence, bounds on ambiguity aversion can again be found analogously to Section 3.2:

$$c < \varphi_W^{-1}\left(\frac{1}{2}\max\left\{\varphi_W(P_{NC} - e) + \varphi_W(P_{E,C} - e), \varphi_W(P_{NC}) + \varphi_W(P_{S,C})\right\}\right) \tag{15}$$

Explicit conditions on parameters can be derived through (15) by considering specific functional forms for φ_W. Conversely, we can consider a given degree of ambiguity aversion and calculate bounds on perception to avoid certainty traps. For instance, if W is extremely averse to ambiguity, we have the following sufficient condition:

$$c < \max\left\{\min_{\hat{\pi}\in\text{supp}\mu_W} P_{S,C} + \hat{\pi}(NC)(P_{NC} - P_{S,C}), \min_{\hat{\pi}\in\text{supp}\mu_W} P_{E,C} + \hat{\pi}(NC)(P_{NC} - P_{E,C}) - e\right\}$$

which reduces to:[29]

$$\min_{\hat{\pi}\in\text{supp}\mu_W} \hat{\pi}(NC) > \frac{c - P_{S,C}}{P_{NC} - P_{S,C}} \qquad \vee \qquad \max_{\hat{\pi}\in\text{supp}\mu_W} \hat{\pi}(NC) < \frac{P_{E,C} - c - e}{P_{E,C} - P_{NC}} \tag{16}$$

In the more general case of a given profile of smooth functionals φ, we can again consider the maximally uncertain belief within some given interval of utility values compatible with the information constraint. Then, we can find a sufficient condition on the maximal distance of models in the support to prevent certainty traps. Let

$$\overline{\pi} := \max_{\hat{\pi}\in\text{supp}\mu_W} 1 - \hat{\pi}(NC) \qquad\qquad \underline{\pi} := \max_{\hat{\pi}\in\text{supp}\mu_W} \hat{\pi}(NC)$$

denote the objective models which make subjectively most appealing choosing $J.E$ and $J.S$, respectively. Then $|\overline{\pi} - \underline{\pi}|$ must satisfy:

$$\begin{aligned} 2\varphi_W(c) < \max\{&\varphi_W\left(P_{E,C} - e + \overline{\pi}(P_{NC} - P_{E,C})\right) + \varphi_W\left(P_{E,C} - e + \underline{\pi}(P_{NC} - P_{E,C})\right) \\ &\varphi_W\left(P_{S,C} + \overline{\pi}(P_{NC} - P_{S,C})\right) + \varphi_W\left(P_{S,C} + \underline{\pi}(P_{NC} - P_{S,C})\right)\} \end{aligned}$$

Also in this case, an explicit solution for this sufficient condition can be derived only when a specific form for φ_W is considered.

5. CONCLUSIONS AND LIMITATIONS

An extensive general framework for Self-Confirming Equilibrium analysis of the strategic form of games with feedback, ambiguity attitudes, and payoff uncertainty has been provided. The results of BCMM are shown to generalize seamlessly to the case of payoff uncertainty. Most importantly, I have given a rigorous account of the game-theoretic interplay between uncertainty aversion and perception, clarifying that even in presence of mild aversion to ambiguity, agents who are more aware of the uncertainty present in their decision environment can get stuck in certainty traps.

The intuition for this fact is a direct consequence of the representation of KMM, who characterized ambiguity aversion as dread toward mean-preserving spreads of second-order beliefs in the space of expected utility values. A game-theoretic framework proves particularly illuminating in demonstrating how this feature can lead to objectively inferior decisions. Section 3 reveals the conditions required for this result are remarkably modest. It is sufficient to pose mild restrictions on the width of supports of beliefs to control flexibly the perception of ambiguity, while preserving the game's fundamental structure.

The application of Section 4 makes these theoretical results operational, demonstrating their relevance for microeconomic theory and particularly for models of adverse selection and relational contracting. Ambiguity aversion can fundamentally alter strategic interactions in employment relationships, potentially explaining observed patterns of inefficient contracting and underinvestment in human capital. This framework provides new insights into why seemingly beneficial long-term relationships might fail to materialize, even when they would objectively improve upon current arrangements. These results offer practical insights for organizational design and policy-making. They suggest that mechanisms for reducing strategic uncertainty, including reputation systems,

[29]Existence is taken for granted provided that $c + e \leq P_{E,C}$ is necessary for participation.

standardized practices, or industry norms, might be more valuable than previously recognized, as they could help overcome the coordination failures identified in the analysis.

The framework of Battigalli et al., 2019b would be most germane to the exploration of consequences on steady-state convergence of early conjectural revisions occurring in learning environments. Additionally, a more robust analysis would refer to extensive-form games where players can change continuation strategy at any non-terminal history (Battigalli et al., 2019a), in light of the implications of dynamic inconsistency pointed out in Section 2. In both cases, however, the complexity of mathematical objects involved grows noticeably. An exploration of comparative ambiguity perception in these settings is open to further research. In addition, a deeper understanding of the normative foundations of ambiguity perception would strengthen the economic significance of the analysis of games with restrictions on beliefs conducted in Section 3.

Nevertheless, core insights remains robust: heightened ambiguity aversion generates a status quo bias, whose likelihood increases with the perception of such ambiguity. Certainty crises, or deeply embedded uncertainty-averse cultures, can trap economic agents in inefficient equilibria. While multiple-prior models primarily address subjective rather than objective rationality, the existence of an objective probability distribution—if such exists—must be unique. Thus, although uncertainty aversion may be consistent with subjective rationality, and uncertainty averse decision-makers may never be convinced their choices are suboptimal, this trait can inadvertently become self-defeating.

REFERENCES

Akerlof, G. A. (1995). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Akerlof, G. A. and R. J. Shiller (2009). *Animal Spirits. How human psychology drives the economy, and why it matters for global capitalism*. Princeton University Press.

Almelhem, A., M. Iyigun, A. Kennedy, and J. Rubin (2023). Enlightenment ideals and belief in progress in the run-up to the industrial revolution: A textual analysis. IZA Discussion Paper No. 16674, IZA Institute of Labor Economics.

Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematics and Statistics 34*, 199–205.

Battigalli, P., E. Catonini, G. Lanzani, and M. Marinacci (2019a). Ambiguity attitudes and self-confirming equilibrium in sequential games. *Games and Economic Behavior 115*, 1–29.

Battigalli, P., E. Catonini, and N. D. Vito (2023). Game theory. an analysis of strategic thinking. Typescript, Bocconi University.

Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2015). Self-confirming equilibrium and model uncertainty. *American Economic Review 105*(2), 646–677.

Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2016). Analysis of information feedback and selfconfirming equilibrium. *Journal of Mathematical Economics 66*, 40–51.

Battigalli, P., A. Francetich, G. Lanzani, and M. Marinacci (2019b). Learning and self-confirming long-run biases. *Journal of Economic Theory 183*, 740–785.

Battigalli, P., M. Gilli, and M. C. Molinari (1992). Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey. *Research in Economics 46*(3), 335–378.

Battigalli, P. and D. Guaitoli (1988). Conjectural equilibria and rationalizability in a game with incomplete information. In P. Battigalli, A. Montesano and F. Panunzi (Eds.), Decisions, Games and Markets, Springer, Boston, MA, pp. 97–124.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica 52*, 1007–1028.

Bewley, T. (2002). Knightian decision theory: Part i. *Decisions in Economics and Finance 25*, 79–110.

Burn-Murdoch, J. (2024). Is the west talking itself into decline? *Financial Times.* January 5.

Choquet, G. (1953). Theory of capacities. *Annales de l'Institut Fourier 5*, 131–295.

Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics 75*, 643–69.

Fishburn, P. C. (1991). Nontransitive preferences in decision theory. *Journal of Risk and Uncertainty 4*(2), 113–34.

Fudenberg, D. and D. K. Levine (1993). Self-confirming equilibrium. *Econometrica 61*(3), 523–545.

Fudenberg, D. and D. K. Levine (1998). *The Theory of Learning in Games.* Cambridge, MA: MIT Press.

Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud (2008). Attitude toward imprecise information. *Journal of Economic Theory 140*, 27–65.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler (2010). Objective and subjective rationality in a multiple prior model. *Econometrica 78*(2), 755–770.

Gilboa, I. and M. Marinacci (2013). Ambiguity and the bayesian paradigm. In D. Acemoglu, M. Arellano and E. Dekel (Eds.), Advances in Economics and Econometrics: Tenth World Congress, Cambridge University Press, pp. 179–242.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Gustafsson, J. E. (2022). *Money-Pump Arguments.* Cambridge: Cambridge University Press.

Harsanyi, J. C. (1967). Games with incomplete information played by "bayesian" players, i–iii part i. the basic model. *Management Science 14*(3), 159–182.

Harsanyi, J. C. (1968a). Games with incomplete information played by "bayesian" players, i–iii part ii. bayesian equilibrium points. *Management Science 14*(5), 320–334.

Harsanyi, J. C. (1968b). Games with incomplete information played by "bayesian" players, i–iii part iii. the basic probability distribution of the game. *Management Science 14*(7), 486–502.

Iyengar, S. S. and M. R. Lepper (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology 79*(6), 995–1006.

Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money.* Macmillan.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Knight, F. H. (1921). *Risk, Uncertainty and Profit.* Houghton Mifflin.

Mirrlees, J. A. (1999). The theory of moral hazard and unobservable behaviour: Part i. *The Review of Economic Studies 66*(1), 3–21.

Mokyr, J. (2017). *A Culture of Growth: The Origins of the Modern Economy*. Princeton University Press.

Nash, J. (1950). Non-cooperative games. *Annals of Mathematics 54*(2), 286–295.

Osborne, M. J. and A. Rubinstein (1991). *Bargaining and Markets*. Academic Press.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica 52*, 1029–1050.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. a definition. *Journal of Economic Theory 2*(3), 225–243.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley and Sons.

Schmeidler, D. (1986). Integral representation without additivity. *Proceedings of the American Mathematical Society 97*, 255–261.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*(3), 571–587.

Seuanez-Salgado, M. (2006). Choosing to have less choice. *FEEM Working Paper* (37).

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Wald, A. (1949). Statistical decision functions. *Annals of Mathematical Statistics 20*(2), 165–205.

APPENDIX

Lemma 1

Proof. Let $i \in I, s_i \in S_i, \sigma^*_{-i} \in \Delta(S_{-i})$. Payoff observability implies the section of payoff function U_{s_i,θ_i} must be $\mathcal{F}_{s_i}$-measurable: that is, any payoff value corresponds to a measurable set in partition $\mathcal{F}_{s_i,\theta_i}$. This measurability implies that if two distributions $\hat{\pi}(\theta_0, \theta_{-i}, s_{-i})$ and $q_0(\theta_0)q_{-i}(\theta_{-i})\sigma^*_{-i}(s_{-i}|\theta_{-i})$ have the same $\mathcal{F}_{s_i}$ measure, they will induce the same conditional distribution (i.e., conditional on $\mathcal{F}_{s_i}$, 5$^{\text{th}}$ equality below). Formally, for any $\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}$,

$$
\begin{aligned}
U_i(s_i, \theta_i, \hat{\pi}) &= \sum_{(\theta_0,\theta_{-i},s_{-i})\in\Theta_0\times\Theta_{-i}\times S_{-i}} U_i(s_i, \theta, s_{-i})\hat{\pi}(\theta_0, \theta_{-i}, s_{-i}) \\
&= \int_{\Theta_0\times\Theta_{-i}\times S_{-i}} U_i(s_i, \theta, s_{-i})\ \hat{\pi}(\mathrm{d}\theta_0, \mathrm{d}\theta_{-i}, \mathrm{d}s_{-i}) \\
&= \int_{\Theta_0\times\Theta_{-i}\times S_{-i}} U_{s_i,\theta_i}\ \hat{\pi}(\mathrm{d}\theta_0, \mathrm{d}\theta_{-i}, \mathrm{d}s_{-i}) \\
&= \int_{\Theta_0\times\Theta_{-i}\times S_{-i}} U_{s_i,\theta_i}\ \hat{\pi}(\mathrm{d}\theta_0, \mathrm{d}\theta_{-i}, \mathrm{d}s_{-i})|\mathcal{F}_{s_i,\theta_i} \\
&= \int_{\Theta_0\times\Theta_{-i}\times S_{-i}} U_{s_i,\theta_i,q_{-i}}\ \sigma^*_{-i}(\mathrm{d}s_{-i}|\theta_{-i})q_0(\mathrm{d}\theta_0)q_{-i}(\mathrm{d}\theta_{-i})|\mathcal{F}_{s_i,\theta_i} \\
&= \int_{\Theta_{-i}}\int_{\Theta_0}\int_{S_{-i}} U_{s_i,\theta_i,q_{-i}}\sigma^*_{-i}(\mathrm{d}s_{-i}|\theta_{-i})q_0(\mathrm{d}\theta_0)q_{-i}(\mathrm{d}\theta_{-i}) \\
&= \sum_{(\theta_0,\theta_{-i},s_{-i})\in\Theta_0\times\Theta_{-i}\times S_{-i}} U_i(s_i, \theta_i, q_{-i}, s_{-i})\sigma^*_{-i}(s_{-i}|\theta_{-i})q_0(\theta_0)q_{-i}(\theta_{-i}) \\
&= U_i(s_i, \theta_i, q_{-i}, \sigma^*_{-i})
\end{aligned}
$$

□

Lemma 2

Proof. I show how to adapt the standard equivalence result on second-order stochastic dominance to a framework with subjective beliefs, rather than lotteries. I prove the "only if" part: as the random variables are shown to satisfy SOSD for all concave functions in the sense of Rothschild & Stiglitz (1970), the reverse implication holds due to their Theorem 2. Let $s \in S$ denote a state of the world and Σ the σ-algebra defined on $\Delta(S)$. Let $U_{s_i,\theta_i} : \Delta(\Theta_0 \times \Theta_{-i} \times S_{-i}) \to \mathbb{R}$ be (the (s_i, θ_i)-section of) a Bernoulli utility function defined on the space of objective lotteries. To simplify notation, we simply write U_i in place of U_{s_i,θ_i}. Let φ_i be a strictly increasing and weakly concave function, and ν^i, μ^i be two subjective beliefs such that μ^i is an MPS of ν^i on the space of expected utility values. This immediately implies $\int_\Sigma U_i(\hat{\pi}(s))\nu^i(\mathrm{d}\hat{\pi}) = \int_\Sigma U_i(\hat{\pi}(s))\mu^i(\mathrm{d}\hat{\pi})$. An objective probability law on unknown states $\hat{\pi}$ and the utility function jointly induce objective expected utility (a lottery) with value

$$u := U_i(\hat{\pi}) \in \mathbb{U} := \mathrm{Im} U_i = \left[\min_{\pi\in\Sigma} U_i(\pi), \max_{\pi\in\Sigma} U_i(\pi)\right] \subseteq \mathbb{R}$$

where the minimum and maximum are well-defined by compact-valuedness of Σ. Since when payoffs are observable U_{s_i} is always $\mathcal{F}_{s_i}$-measurable, we can write — with a slight abuse of notation, i.e. identifying a belief μ on objective lotteries with its push-forward $\mu \circ U^{-1}$ on the space of expected utility values

$$\int_\Sigma \varphi\left(U(\hat{\pi}(s))\right)\nu\left(\mathrm{d}\hat{\pi}\right) = \int_{\mathbb{U}} \varphi\left(u\right)\nu\left(\mathrm{d}u\right)$$

Since μ is a mean-preserving spread of ν, there exists some random variable z such that, for $v \sim \nu$ and $u \sim \mu$, $u \stackrel{D}{=} v + z$ and $\mathbb{E}(z|v) = 0$ for every $v \in \mathbb{U}$. By the law of iterated expectations and the MPS relation: $\mathbb{E}_{\mu^i}(\varphi_i(u)) = \mathbb{E}_{\nu^i}(\varphi_i(u)) = \mathbb{E}_{\nu^i}(\varphi_i(v+\varepsilon)) = \mathbb{E}_{\nu^i}\left[\mathbb{E}(\varphi_i(v+\varepsilon)|v)\right]$. Therefore we may write

$$\begin{aligned}\int \varphi_i(u)\mu^i(\mathrm{d}u) &= \int \mathbb{E}\left[\varphi_i(v+\varepsilon)|v\right]\nu^i(\mathrm{d}v)\\ &\leq \int \varphi_i\left(\mathbb{E}\left[v+\varepsilon|v\right]\right)\nu^i(\mathrm{d}v)\\ &= \int \varphi_i(\underbrace{\mathbb{E}\left[v|v\right]}_{=v} + \underbrace{\mathbb{E}\left[\varepsilon|v\right]}_{=0})\nu^i(\mathrm{d}v)\\ &= \int \varphi_i(v)\nu^i(\mathrm{d}v)\end{aligned}$$

The inequality is an application of Jensen's inequality valid for any weakly concave function, and we can separate the sum in the expected value by linearity of the operator. Hence, ν^i second-order stochastically dominates μ^i under φ_i. □

Lemma 3

Proof. I prove that μ is an MPS of Σ-measurable function (probability distribution) ν on the space of expected utility values if

1. $\int_\Sigma \hat{\pi}(s)\nu\,(\mathrm{d}\hat{\pi}) = \int_\Sigma \hat{\pi}(s)\mu\,(\mathrm{d}\hat{\pi})$
2. conv supp$\nu \subset$ conv suppμ

Step 1: Condition (1) is equivalent to $\int_\Sigma U(\hat{\pi}(s))\nu\,(\mathrm{d}\hat{\pi}) = \int_\Sigma U(\hat{\pi}(s))\mu\,(\mathrm{d}\hat{\pi})$. The two measures are necessarily absolutely continuous (with respect to the Lebesgue measure), so that we can invoke the Radon-Nykodym theorem and, together with the results of Lemma 2 restate condition 1 as

$$\int_{\mathbb{U}} u\ \nu\,(u)\,\mathrm{d}u = \int_{\mathbb{U}} u\ \mu\,(u)\,\mathrm{d}u$$

That is, $\mathbb{E}_\nu(u) = \mathbb{E}_\mu(u)$.

Step 2: Let $g = \mu - \nu$ denote the function which describes the probability spread between the two distributions , with $G(u) = \int_{-\infty}^{u} g(t)\mathrm{d}t$. We need to prove that $T(y) := \int_{-\infty}^{y} G(u)\mathrm{d}u \geq 0\ \forall y \in \mathbb{R}$ with $T(\infty) = 0$. By Lemma 1* in Deb & Seo (2011), Lemma 2 and Theorem 2 in Rothschild & Stiglitz (1970), this proves that μ is a MPS of ν.

Let $\underline{u} := \inf \mathrm{supp}\mu$ and $\overline{u} := \sup \mathrm{supp}\mu$, $\underline{v} := \inf \mathrm{supp}\nu$ and $\overline{v} := \sup \mathrm{supp}\nu$. Condition (2) guarantees $\underline{u} < \underline{v}$ or $\overline{u} > \overline{v}$. Consider the first case; the other is symmetric. Clearly, then, $\int_{-\infty}^{\underline{u}} \mu(t)\mathrm{d}t = \int_{-\infty}^{\underline{u}} \nu(t)\mathrm{d}t = 0$ and $\int_{-\infty}^{\overline{u}} \mu(t)\mathrm{d}t = \int_{-\infty}^{\overline{u}} \nu(t)\mathrm{d}t = 1$, so that $g(u) = G(u) = 0\ \forall u \notin [\underline{u}, \overline{u}]$. This automatically grants $T(\infty) = yG(y)]_{-\infty}^{\infty} - \int_{-\infty}^{\infty} ug(u)\mathrm{d}u = 0$ (since in this case 0 is rapidly reached within g on a bounded subset of the real line, the limit is not indeterminate). Additionally, $u \in [\underline{u}, \underline{v}) \Rightarrow g(u) > 0 \Rightarrow G(u) > 0$.

There is a unique point $c \geq \underline{v}$ such that $G(u) \geq 0\ \forall u \leq c$ and $G(u) \leq 0\ \forall u > c$. Assume by way of contradiction this were not the case, i.e. G changes sign more than once. Then there exists (at least one) point $c' > c$ such that

- $G(u) \geq 0$ if $u \leq c$;
- $G(u) \leq 0$ if $u \in (c, c')$;
- $G(c') > 0$.

But this cannot be: either (i) $G(u) \geq 0$ for every $u > c'$, in which case μ conserves more probability density until the end of the distribution, contradicting $T(\infty) = 0$, or (ii) the latter is satisfied, which requires $G(c'') < 0$ for some other $c'' \in (c', \overline{v}]$, in which case the mean-preservation assumption is necessarily contradicted, as probability mass or density has been moved towards the lower end of the distribution on net. In accordance with this, it follows that $T(y) \geq 0\ \forall y \leq c$. For $y > c$, note that

$$T(y) = \underbrace{T(c)}_{\geq 0} + \int_c^y G(u)\mathrm{d}u$$

Since $\forall u > c, G(u) \leq 0$, $T(y)$ must be nonincreasing for all $y > c$. Since also $T(y) = 0\ \forall y \geq \overline{u}$, it follows that $T(y) \geq 0\ \forall y$, which completes the proof. □

Proposition 1

Proof. If $SSCE(BG, \varphi, \nu) = \emptyset$, the statement vacuously holds. Assume $SSCE(BG, \varphi, \nu) \neq \emptyset$. Let $\sigma^* \in SSCE(BG, \varphi, \nu)$ and consider some θ_i and $s_i^* \in \mathrm{supp}\sigma^*_{\theta_i}$. The game is assumed to satisfy Def. 1; by Lemma 1,

$$V_{\theta_i}^{\varphi_i}(s_i^*, \nu^i_{s_i,\theta_i}) = V_{\theta_i}^{\varphi_i}(s_i^*, \mu^i_{s_i,\theta_i}) = U_i(s_i^*, \theta_i, \hat{\pi}) = U_i(s_i^*, \theta_i, q_{-i}, \sigma^*_{-i})\ \forall \hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s_i^*, q_{-i}, \sigma^*_{-i})$$

By Lemma 2, if φ_i is weakly concave and $\mu^i_{s_i,\theta_i}$ is a mean-preserving spread of $\nu^i_{s_i,\theta_i}$, for every $s_i \in S_i$,

$$\begin{aligned} V_{\theta_i}^{\varphi_i}(s_i^*, \mu^i_{s_i,\theta_i}) &= U_i(s_i^*, \theta_i, q_{-i}, \sigma^*_{-i}) \\ &= V_{\theta_i}^{\varphi_i}(s_i^*, \nu^i_{s_i,\theta_i}) \geq V_{\theta_i}^{\varphi_i}(s_i, \nu^i_{s_i,\theta_i}) \\ &= \int \varphi_i\left(U_i(s_i, \theta_i, \hat{\pi})\right) \nu^i_{s_i,\theta_i}(\mathrm{d}\hat{\pi}) \\ &\geq \int \varphi_i\left(U_i(s_i, \theta_i, \hat{\pi})\right) \mu^i_{s_i,\theta_i}(\mathrm{d}\hat{\pi}) = V_{\theta_i}^{\varphi_i}(s_i, \mu^i_{s_i,\theta_i}) \end{aligned}$$

Since the support of $\mu^i_{s_i,\theta_i}$ is contained in $\hat{\Sigma}_{-i,\theta_i}(s_i^*, q_{-i}, \sigma^*_{-i})$ by assumption, the statement holds. □

Theorem 1

Proof. First note that the second inclusion always holds by construction and, if $\varepsilon = 0$, then $SSCE(BG, \varphi, \varepsilon) = SCE(BG, F)$. If $\gamma = 0$, the whole statement vacuously holds. Consider then the generic case $\varepsilon, \gamma > 0$, and a profile of beliefs $\mu \in R_\varepsilon$. Either $\mu \in R^*_\varepsilon$ or $\mu \in R_\varepsilon \backslash R^*_\varepsilon$. In the first case, $\mu \in R^\gamma_\varepsilon$, hence the equilibrium set for this belief is preserved. Assume it is the nontrivial case (otherwise, the belief is irrelevant) where $SSCE(BG, \varphi, \mu) \neq \emptyset$. As mentioned in the text, ignorance of the matching process ensures convexity of $\hat{\Sigma}_{-i,\theta_i}$ for any strategy profile. Let $\mu \in R_\varepsilon \backslash R^*_\varepsilon$. Then by definition there exist $i \in I, \theta_i \in \Theta_i, \sigma^* \in SSCE(BG, \varphi, \mu)$ and $s_i \in \mathrm{supp}\sigma^*_i$ such that we can construct a new profile $\tilde{\mu}$ as follows.

Let $\underline{u} := \inf_{\hat{\pi} \in \mathrm{supp}\mu^i_{s_i,\theta_i}} U_{s_i,\theta_i}(\hat{\pi})$ and $\overline{u} := \sup_{\hat{\pi} \in \mathrm{supp}\mu^i_{s_i,\theta_i}} U_{s_i,\theta_i}(\hat{\pi})$. These utility values can be attained within the support of $\mu^i_{s_i,\theta_i}$ as the game is finite, hence the suprema are maxima. Denote the expected utility profile by $\hat{u} := \int_{\underline{u}}^{\overline{u}} u\ \mu^i_{s_i,\theta_i}(\mathrm{d}u)$. Then there exists $\beta \in (\underline{u}, \overline{u})$ such that $\int_{\underline{u}}^{\beta} u\ \mu^i_{s_i,\theta_i}(\mathrm{d}u) = \int_{\beta}^{\overline{u}} u\ \mu^i_{s_i,\theta_i}(\mathrm{d}u) = \hat{u}/2$. Let $a = \varepsilon - (\overline{u} - \underline{u}) \in [0, \varepsilon]$, and $t \in (a, a + \gamma/2]$. Then, define

$$\tilde{\mu}^i_{s_i,\theta_i}(u) := \begin{cases} \mu^i_{s_i,\theta_i}(u + t) & u \in [\underline{u} - t, \beta - t] \\ \mu^i_{s_i,\theta_i}(u - t) & u \in [\beta + t, \overline{u} + t] \\ 0 & \text{otherwise} \end{cases}$$

For every other strategy, type and player, assume $\tilde{\mu}$ and μ coincide. Then, $\tilde{\mu} \in (R_{\varepsilon+\gamma} \backslash R_\varepsilon)$. By convexity of $\hat{\Sigma}_{-i,\theta_i}$, and again by definition of $R^*_\varepsilon \not\ni \mu$, an appropriate choice of i, θ_i and s_i ensures that for $t \to a^+$, $\tilde{\mu}^i_{s_i,\theta_i} \in \Delta\left(\hat{\Sigma}_{-i,\theta_i}(s_i, \sigma^*_{-i})\right)$. It is immediate to verify that $\int_{\underline{u}-t}^{\beta-t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u) = \frac{\hat{u}-t}{2}$, $\int_{\beta+t}^{\overline{u}+t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u) = \frac{\hat{u}+t}{2}$, so that

$$
\begin{aligned}
\int_{\mathrm{supp}\tilde{\mu}^i_{s_i,\theta_i}} u \, \tilde{\mu}^i_{\theta_i}(\mathrm{d}u) &= \int_{\underline{u}-t}^{\overline{u}+t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u) \\
&= \int_{\underline{u}-t}^{\beta-t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u) + \underbrace{\int_{\beta-t}^{\beta+t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u)}_{=0} + \int_{\beta+t}^{\overline{u}+t} u \, \tilde{\mu}^i_{s_i,\theta_i}(\mathrm{d}u) \\
&= \frac{\hat{u}}{2} - \frac{t}{2} + \frac{\hat{u}}{2} + \frac{t}{2} = \hat{u}
\end{aligned}
$$

By Lemma 3 and Proposition 1, $SSCE(BG, \varphi, \mu) \subseteq SSCE(BG, \varphi, \tilde{\mu})$. Since μ was arbitrary, for any $\mu \notin R^*_\varepsilon$ we can always find some $\tilde{\mu} \in (R_{\varepsilon+\gamma} \backslash R_\varepsilon) \subseteq R^\gamma_\varepsilon$ as above. □

Lemma 4

Proof. By Proposition 1 in BCMM, we know that if a game satisfies observable payoffs then $MSCE \supseteq SSCE$, thus including those determined by any exogenous restriction on priors. Analogously to their setting, it is sufficient to show that every MSCE is a Bayesian Equilibrium to prove that also every SSCE is a Bayesian Equilibrium, regardless of the restrictions we put on beliefs as long as we allow at least singleton support priors − which is the case since $C_1 \subset R_\varepsilon \ \forall \varepsilon > 0$ – considering that $SSCE(BG, \varphi, C_1) = SCE(BG, f)$. Hence we simply extend their proof to the payoff uncertain case. Let σ^* be a MSCE of (BG, φ) in accordance with Definition 4. By Lemma 1, observability of payoffs implies that for any strategy s_i of any θ_i, $U_i(s_i, \theta_i, \hat{\pi}) = U_i(s_i, \theta_i, q_{-i}, \sigma^*_{-i})$ for every $\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})$. Consider then some θ_i and $s^*_i \in \mathrm{supp}\sigma^*_{\theta_i}$; OSI implies that $\mathcal{F}_{s^*_i,\theta_i} = \mathcal{F}_{s_i,\theta_i}$ for every s_i that is a best reply to some conjecture. We therefore have

$$
\begin{aligned}
&\hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i}) = \\
&\quad = \left\{\hat{\pi} \in \Delta(\Theta_0 \times \Theta_{-i} \times S_{-i}) : \hat{\pi}(\theta_0, \theta_{-i}, s_{-i})|\mathcal{F}_{s_i,\theta_i} = q_0(\theta_0) q_{-i}(\theta_{-i}) \sigma^*_{-i}(s_{-i}|\theta_{-i})|\mathcal{F}_{s_i,\theta_i}\right\} \\
&\quad = \left\{\hat{\pi} \in \Delta(\Theta_0 \times \Theta_{-i} \times S_{-i}) : \hat{\pi}(\theta_0, \theta_{-i}, s_{-i})|\mathcal{F}_{s^*_i,\theta_i} = q_0(\theta_0) q_{-i}(\theta_{-i}) \sigma^*_{-i}(s_{-i}|\theta_{-i})|\mathcal{F}_{s^*_i,\theta_i}\right\} \\
&\quad = \hat{\Sigma}_{-i,\theta_i}(s^*_i, q_{-i}, \sigma^*_{-i})
\end{aligned}
$$

Hence we can prove that σ^* is a Bayesian equilibrium:

$$
\begin{aligned}
U_i(s^*_i, \theta_i, q_{-i}, \sigma^*_{-i}) &= \min_{\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s^*_i, q_{-i}, \sigma^*_{-i})} U_i(s^*_i, \theta_i, \hat{\pi}) \geq \min_{\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})} U_i(s^*_i, \theta_i, \hat{\pi}) \\
&= \min_{\hat{\pi} \in \hat{\Sigma}_{-i,\theta_i}(s_i, q_{-i}, \sigma^*_{-i})} U_i(s_i, \theta_i, \hat{\pi}) = U_i(s_i, \theta_i, q_{-i}, \sigma^*_{-i})
\end{aligned}
$$

□