LEMMA Working Paper

n° 2025-07

Ambiguity Aversion Equilibrium in Games

Youcef Askoura

Université Paris-Panthéon-Assas, LEMMA

Antoine Billot

Université Paris-Panthéon-Assas, LEMMA

Moussa Larbani

School of Mathematics and Statistics, Carleton University

LEMMA EA 4442 - 4 rue Blaise Desgoffe 75006 Paris

Ambiguity Aversion Equilibrium in Games

Youcef Askoura[*], Antoine Billot[†] and Moussa Larbani[‡]

September 19, 2025

Abstract

We introduce a new solution concept for games with ambiguity-averse players, termed the *Ambiguity Aversion Equilibrium* (**AAE**). In contrast to existing models that treat ambiguity as endogenous, we model ambiguity as an exogenous feature entirely derived from the game's structure. Players' ambiguity aversion is incorporated via a non-Bayesian state-dependent expected utility framework, capturing worst-case reasoning over uncertain states. The **AAE** preserves strategic interaction among players by ensuring that ambiguity affects their evaluations but not the definition of equilibrium itself. This approach provides new insights into how structural ambiguity shapes behavior without relying on belief adjustments to construct equilibria. We compare the **AAE** to well-known concepts in the literature, using standard finite games to illustrate its distinctive features. Furthermore, we establish existence results for the **AAE** in both pure and mixed strategies.

Keywords : ambiguity aversion, equilibrium, maxmin expected utility.
JEL classification : C70, C72, C79.

1 Introduction

It is widely acknowledged that the classical assumptions of game theory—which posit that players are fully rational and perfectly informed about the parameters of the game—constitute a useful but often unrealistic simplification (see, eg. Luce and Raiffa (1957), Dominiak and Eichberger (2016)). In particular, conventional models often fail to incorporate ambiguity and strategic uncertainty, despite their pervasive influence on real-world economic and social interactions. Addressing these dimensions within game-theoretic frameworks remains a central theoretical challenge.

[*]Lemma, Université Paris Panthéon-Assas. Email: youcef.askoura@u-paris2.fr
[†]Lemma, Université Paris Panthéon-Assas. Email: antoine.billot@assas-universite.fr
[‡]School of Mathematics and Statistics. Carleton University, Ottawa, Canada, Email: larbani61@hotmail.com

The integration of ambiguity aversion into strategic models began with seminal works in the 1990s (Klibanoff (1996), Lo (1996), Dow and Werlang (1994)). A common insight behind these contributions lies in the reinterpretation of mixed-strategy Nash equilibria: in two-player games, one player's mixed strategy reflects the other's subjective beliefs over possible pure strategies (Dow and Werlang (1994)). Building on this, Dow and Werlang (1994) generalized Nash equilibrium under uncertainty by replacing probabilities with Choquet capacities, thereby moving from additive to non-additive beliefs. Subsequent contributions, notably by Klibanoff (1996), Lo (1996), and Marinacci (2000), modeled ambiguity aversion through sets of priors or belief functions, enriching the theoretical foundations of games with uncertainty. Further refinements have incorporated attitudes toward ambiguity, including models of optimism and pessimism based on neo-additive capacities (Eichberger, Kelsey and Schipper (2009); Eichberger and Kelsey (2014)).

Despite these advances, most existing models treat ambiguity as endogenous to the equilibrium concept: players' beliefs are manipulated or constrained within the solution framework itself, often diminishing the role of strategic interaction. Players effectively interact through their beliefs rather than directly through strategies, leading to equilibrium definitions in which belief consistency substitutes for classical strategic reasoning.

This paper proposes an alternative approach. We introduce an equilibrium concept for games with ambiguity-averse players, the *Ambiguity Aversion Equilibrium* (**AAE**), in which ambiguity is modeled as exogenous—derived entirely from the game's structure and independent of the equilibrium construction. Players interact via their pure or mixed strategies, as in classical game theory, while their ambiguity aversion is embedded in their payoff evaluations through a maxmin expected utility framework. Importantly, players form beliefs based on the normal-form structure of the game and a *status quo* outcome, but these beliefs do not define the equilibrium set. As a result, strategic interaction is preserved, and equilibrium stability rests solely on players' rationality, consistent with traditional game-theoretic principles.

In this framework, players reason along the lines of: "If the current *status quo* is x, what should I do?" rather than adjusting beliefs to rationalize outcomes. This preserves both strategic logic and ambiguity aversion in a coherent manner. Our analysis focuses on complete-information games, but the approach provides a natural foundation for extending ambiguity models to more general settings (see, Kajii and Ui (2005), He and Yannelis (2013) for related work on incomplete information games).

The paper proceeds as follows. Section 2 introduces the framework, Section 3 defines the **AAE** concept, highlights its distinct features by comparing it to established equilibrium notions using standard finite games, and analyzes the existence of pure-strategy equilibria under ambiguity aversion. Section 4 extends the results to mixed strategies. Section 5 concludes with perspectives for future research.

2 Framework

Let $N = \{1, \ldots, n\}$ denote the set of *players*, and let X_i, for each $i \in N$, be the set of *strategies available* to player i. We adopt standard notation from game theory: for any player $i \in N$, we denote by $-i = N \setminus \{i\}$ the set of all players except i, and x_{-i} denotes a strategy profile of all players other than i. The set of strategy profiles of the game is denoted by $X = \prod_{i \in N} X_i$. We assume that each X_i is a compact metric topological space equipped with its Borel σ-algebra $\mathcal{B}(X_i)$, and denote by $\mathcal{P}(X_{-i})$ the set of Borel probability measures on X_{-i}. Throughout this paper, any subspace of a topological space inherits the induced topology and the corresponding σ-algebra.

2.1 Non-Bayesian state-dependent preferences

We assume that all players are fully informed about the structure of the game, and that cooperation is not permitted—the game is non-cooperative. Each player is confident in his own rationality and, based on the known structure of the game, forms beliefs regarding the strategies of his opponents. In other words, the $(n-1)$-tuple formed by the decisions of the other players, insofar as these constitute a source of uncertainty, serves as the "state" for each player and directly affects the evaluation of all his possible strategies. In addition, we assume that all players exhibit ex-ante ambiguity aversion in the sense of Gilboa and Schmeidler (1989), with a state-dependent utility specification adapted from Hill (2019), where the state-dependence of player's i utility reflects the dependence of i's decisions on $-i$'s decisions (but where preferences, unlike in Hill's model, are not imprecise).

Each player i is endowed with a preference relation $\succsim_{i,x}$ that depends on the *status quo strategy profile* $x = (x_i, x_{-i}) \in X$. Here, x_i, the *individual status quo strategy* of player i, refers to maintaining his current strategic choice x_i, given the full profile x. Relative to this *status quo*, player i evaluates and compares *acts*, ie. measurable functions $h : X_{-i} \to X_i$, via the preference relation $\succsim_{i,x}$. We denote by $\mathcal{H}$ the set of all such acts. The set X_{-i} of opponents' strategies defines the *state space* for player i, while X_i represents his *consequence space*. An act $h_c \in \mathcal{H}$ is called a *constant act* if there exists some $x_i \in X_i$ such that $h_c(y_{-i}) = x_i$, for all $y_{-i} \in X_{-i}$. Let $h \in \mathcal{H}$, $y_{-i}, z_{-i} \in X_{-i}$ and $\tilde{y}_i \in X_i$. Denote by $h_{(y_{-i},\tilde{y}_i)}$, the act such that:

$$h_{(y_{-i},\tilde{y}_i)}(z_{-i}) := \begin{cases} \tilde{y}_i & \text{if } z_{-i} = y_{-i}, \\ h(z_{-i}) & \text{otherwise.} \end{cases}$$

Now, accordingly, for every $i \in N$ and every $x \in X$, we assume that the preference relation $\succsim_{i,x}$ satisfies standard and adapted axioms of Hill (2019):

Weak order. $\succsim_{i,x}$ is complete and transitive over $\mathcal{H}$.
Non-degeneracy. $\exists h_1, h_2 \in \mathcal{H}$ st. $h_1 \succ_{i,x} h_2$.
Continuity. $\forall h' \in \mathcal{H}$, the subsets $\{h \in \mathcal{H} \mid h \succsim_{i,x} h'\}$ and $\{h \in \mathcal{H} \mid h' \succsim_{i,x} h\}$ are closed under the uniform convergence topology on $\mathcal{H}$.

Uncertainty aversion. $\forall h_1, h_2 \in \mathcal{H}$, $\forall \alpha \in [0,1]$, if $h_1 \sim_{i,x} h_2$, then $\alpha h_1 + (1-\alpha) h_2 \succsim_{i,x} h_2$.

Most papers dealing with state-dependent utility adopt a framework close to that of expected utility (Karni, Schmeidler and Vind (1983)), whereas the axiomatic approaches to ambiguity explicitly rely on constant acts (Gilboa and Schmeidler (1989)), which lose their meaning as soon as utilities are state-dependent. The following two axioms thus reflect the specificity of a model that draws from both strands of the literature. First, we adopt the c-independence axiom of Gilboa and Schmeidler (1989) in place of ec-independence from Hill (2019), since imprecise preferences are not considered here.

C-independence. $\forall h, \tilde{h}, h_c \in \mathcal{H}$, $\forall \alpha \in [0,1]$, $h \succsim_{i,x} \tilde{h}$ iff $\alpha h + (1-\alpha) h_c \succsim_{i,x} \alpha \tilde{h} + (1-\alpha) h_c$.

Second, the following axiom allows for a coherent definition of the notion of preference conditional on a *status quo*.

State consistency. $\forall y \in X$, $\forall y_i, \tilde{y}_i \in X_i$, $\forall h \in \mathcal{H}$, if, $h_{(y_{-i}, \tilde{y}_i)} \succsim_{i,x} h_{(y)}$, then, $\forall \tilde{h} \in \mathcal{H}$, $\tilde{h}_{(y_{-i}, \tilde{y}_i)} \succsim_{i,x} \tilde{h}_{(y)}$.

The following theorem can then be established.

Theorem 1. *Let $i \in N$ be a player. Let $\succsim_{i,x}$ be a preference relation over acts $\mathcal{H}$, defined relative to a status quo profile $x \in X$. The following are equivalent:*
(i) *The preference $\succsim_{i,x}$ satisfies weak order, non-degeneracy, continuity, uncertainty aversion, c-independence and state consistency.*
(ii) *There exist a continuous state-dependent payoff function $u_i : X_i \times X_{-i} \to \mathbb{R}$, and a null-consistent, closed, convex, nonempty set $C_i(x) \subseteq \mathcal{P}(X_{-i})$ of probability measures with compact support such that, for all acts $h_1, h_2 \in \mathcal{H}$,*

$$h_1 \succsim_{i,x} h_2 \iff \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(h_1(y_{-i}), y_{-i})\, dp(y_{-i}) \geq \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(h_2(y_{-i}), y_{-i})\, dp(y_{-i}).$$

Proof. Assume $\succsim_{i,x}$ satisfies the axioms of weak order, non-degeneracy, continuity, uncertainty aversion, c-independence, and state consistency. Fix $x \in X$ and $x_{-i} \in X_{-i}$. The state consistency axiom allows us to define a local preference $\succsim_{i,x_{-i}}$ over X_i such that, for any $x_i, y_i \in X_i$, $x_i \succsim_{i,x_{-i}} y_i$ if and only if there exist $h, g \in \mathcal{H}$ such that:

$$h^x(y_{-i}) := \begin{cases} x_i & \text{if } y_{-i} = x_{-i}, \\ h(y_{-i}) & \text{otherwise} \end{cases}, \quad g^y(y_{-i}) := \begin{cases} y_i & \text{if } y_{-i} = x_{-i}, \\ g(y_{-i}) & \text{otherwise} \end{cases},$$

$$\text{and } h^x \succsim_{i,x} g^y.$$

In short, $x_i \succsim_{i,x_{-i}} y_i$ iff there exist acts h, g that coincide everywhere except at x_{-i}, and such that $h \succsim_{i,x} g$. In this way, the existence of a local preference $\succsim_{i,x_{-i}}$ guarantees a coherent interpretation of the state space X_{-i} from the perspective of player i, and ensures satisfaction of Hill's state consistency axiom (A7). Under the axioms given (in particular, weak order, uncertainty aversion, c-independence and continuity), we can then invoke the representation theorem of Gilboa and Schmeidler (1989): there exists a utility function $u_i : X_i \times X_{-i} \to \mathbb{R}$ and a closed, convex, nonempty set $C_i(x) \subseteq \mathcal{P}(X_{-i})$ such that, for all $h_1, h_2 \in \mathcal{H}$:

$$h_1 \succsim_{i,x} h_2 \iff \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(h_1(y_{-i}), y_{-i}) \, dp(y_{-i}) \geq \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(h_2(y_{-i}), y_{-i}) \, dp(y_{-i}).$$

From standard results, u_i is unique up to positive affine transformations. The belief set $C_i(x)$ is uniquely determined in the sense of minimality and closed convex support of the infimum representation. Null-consistency follows from the monotonicity implied by the preference over constant acts and ensures that all priors in $C_i(x)$ agree on the set of null events.

Conversely, suppose that such a representation exists. Then all the axioms in (i) are satisfied: weak order follows from the ordering over real numbers; non-degeneracy from the strict inequality between integrals; continuity from continuity of u_i and weak*-compactness of $C_i(x)$; uncertainty aversion from the convexity of $C_i(x)$; c-independence from the linearity of the expectation and state consistency from the fact that preferences over acts agree on consequences that differ only at one state. $\square$

If player i maintains his *status quo* strategy x_i and exhibits ambiguity aversion regarding potential deviations by the other players, he evaluates the constant act $h_c \equiv x_i$ as:

$$\inf_{p \in C_i(x)} \int_{X_{-i}} u_i(x_i, y_{-i}) \, dp(y_{-i}).$$

Our framework can be viewed as a special case of the model developed and axiomatized by Hill (2019), where an act h is evaluated according to:

$$\min_{p \in C} \sum_{s \in S} p(s) \min_{u \in v(s)} u(h(s)),$$

Here, we adapt this model, which was originally formulated in the finite case with multi-utility described by the set $v(s)$ (to capture imprecise preferences), to the standard setting of game theory (infinite case with single utility or payoff function), where each individual preference $\succsim_{i,x}$ is represented by a continuous *expected utility* $U_i : \mathcal{H} \to \mathbb{R}$ such that:

$$U_i(h) = \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(h(y_{-i}), y_{-i}) \, dp(y_{-i}).$$

2.2 The normal form of the game

The functions u_i are of the von Neumann–Morgenstern (vNM) type, they serve to evaluate the actual payoffs of the players. By contrast, the utilities U_i capture the expected payoffs, here understood as cautiously anticipated utilities, thus reflecting the players' ambiguity aversion. We assume that all the payoff functions $u_i : X \to \mathbb{R}$ are bounded. By Theorem 1, they are also continous, however upper semicontinuity is actually sufficient for the below results. It is important to note that the utility U_i continues to depend on x_{-i}, since the argument x appears in the belief correspondence C_i over which the infimum is computed. Consequently, player i simultaneously accounts for both his beliefs and the strategies of his opponents. In other words, strategic interaction remains embedded within the expected utilities. This feature marks a fundamental distinction between our model and the related equilibrium concepts discussed earlier. We denote the resulting normal-form game by:

$$G = (N, X_i, u_i, U_i, C_i,\ i \in N).$$

By construction of the game G, ambiguity is entirely incorporated within the expected utilities U_i and the belief correspondences C_i. From this point onward, ambiguity is treated as an exogenous characteristic of the model.

3 The pure strategy case

In the presence of ambiguity-averse players, it is essential to adapt the classical equilibrium concept to account for how players form cautious assessments of payoffs based on their beliefs. We now introduce an equilibrium notion that incorporates exogenous ambiguity aversion in the case where players adopt pure strategies. This definition extends the standard best-response criterion by integrating the players' perceived payoffs, which reflect their worst-case evaluations consistent with their ambiguity aversion.

Definition 1. *An* exogeneous ambiguity aversion pure strategy equilibrium ($\mathbf{AAE}_{ps}$) *for G is a strategy profile $x \in X$ such that: for all $i \in N$, for all $x'_i \in X_i$, $U_i(x'_i, x_{-i}) \leq U_i(x)$.*

A particularly relevant case arises when players are also confident in the rationality of their opponents. In this context, each player anticipates only individually rational deviations from any given strategy profile. This situation is further illustrated through the examples provided below. We denote the *best response correspondence of player i* by:

$$\beta_i : X \to X_i, \text{defined at each } x \in X, \text{ by } \beta_i(x) = \arg\max_{y_i \in X_i} u_i(y_i, x_{-i}),$$

which represents the set of optimal strategies available to player i in response to the *status quo* strategy profile x. This correspondence is well-defined and guaranteed to be nonempty under the structural assumptions imposed on the game G.

When player i evaluates a given strategy profile x, he anticipates that his opponents may choose deviations from x within the product set $\beta_{-i}(x) = \prod_{j \neq i} \beta_j(x)$, in accordance with a probability measure belonging to his belief set $C_i(x)$. The assumption that opponents are rational, together with the independence of their deviations, is then formalized by the following condition:

[**A1**] $\forall i \in N, C_i(x) \subset \otimes_{j \neq i} \mathcal{P}(\beta_j(x))$, where we use the following notation: given two measurable spaces E and F, the product set of probability measures is defined by: $\mathcal{P}(E)\otimes\mathcal{P}(F) = \{p\otimes q : p \in \mathcal{P}(E),\ q \in \mathcal{P}(F)\}$. Thus, to each i, we associate a state-contingent belief correspondence $C_i : X \to \otimes_{j\neq i}\mathcal{P}(X_j)$, such that $C_i(x) \subset \otimes_{j\neq i}\mathcal{P}(\beta_j(x))$, for every $x \in X$.

In standard analysis, a strategy is *rationalizable* if it is a best response to some belief about the opponents' strategies, where those beliefs themselves only put weight on other rationalizable strategies (Bernheim (84)). Rationalizability is defined iteratively by eliminating strategies that are never best responses. Every Nash equilibrium is rationalizable, but rationalizability often shrinks the set of plausible equilibria when an equilibrium relies on implausible beliefs. Our concept of a pure-strategy equilibrium, $\mathbf{AAE}_{ps}$, considers players insofar as they evaluate strategies through worst-case payoffs (ambiguity aversion), assume opponents are rational, ie. only deviations within their best-response correspondences $\beta_j(x)$ are possible, and form beliefs $C_i(x)$ restricted to rational opponents' deviations. Hence, players choose cautious best responses to rationalizable deviations of others.

3.1 Examples

Our structure aligns with rationalizability analysis in the sense where the set $\beta_j(x)$ corresponds to the rationalizable strategies of player j at x. [**A1**] ensures players only believe others play within their rationalizable sets. However, the worst-case evaluation is stricter than standard rationalizability: instead of a best response to some belief about rationalizable strategies, players choose a best response to the least favorable belief in $C_i(x)$. Thus, $\mathbf{AAE}_{ps}$ can be seen as a refinement of rationalizability under ambiguity aversion.

We now present several examples to illustrate this concept of equilibrium and to compare it with both the standard Nash equilibrium and well-established equilibrium concepts under ambiguity found in the literature. Throughout these examples, we adopt the extreme cautious case consistent with assumption **A1**, ie. for every $x \in X$, $C_i(x) = \otimes_{j\neq i}\mathcal{P}(\beta_j(x))$. As a result, the modified utility functions U_i simplify to the following expression:

$$\forall x \in X, U_i(x) = \inf_{y_{-i}\in\beta_{-i}(x)} u_i(x_i, y_{-i}).$$

Example 1. In the example below, the left matrix displays the actual payoffs u_1 and u_2 of the original two-player game, while the right matrix presents the perceived utilities U_1 and U_2, computed under the assumption that $C_i(x) = \mathcal{P}(\beta_{-i}(x))$, for each player $i \in N = \{1, 2\}$. Player 1 chooses among rows $l1$, $l2$, and $l3$ while player 2 chooses among columns $c1$, $c2$, and $c3$.

original payoffs (u_1, u_2)			game payoffs (U_1, U_2)		
$(-1,-1)$	$(-1,-1)$	$(-1,-1)$	$(-1,-1)$	$(-1,3)$	$(-1,4)$
$(-1,-1)$	$(1,1)$	$(2,2)$	$(2,-1)$	$(2,3)$	$(2,4)$
$(-1,-1)$	$(3,3)$	$(4,4)$	$(4,-1)$	$(4,3)$	$(4,4)$

We see that the original game has two Nash equilibria: $(l1, c1)$ and $(l3, c3)$. However, the equilibrium $(l1, c1)$ appears highly implausible. Even if player 1 lacks access to player 2's payoffs, strategy $l3$ strictly dominates both $l1$ and $l2$, making it the rational choice for player 1. By symmetry, the same reasoning applies to player 2, for whom $c3$ strictly dominates $c1$ and $c2$. Thus, both players are led to choose their dominant strategies $l3$ and $c3$, reinforcing the relevance of $(l3, c3)$ as the only credible Nash equilibrium. In other words, $(l1, c1)$ is a Nash equilibrium but not rationalizable in a strong sense. Rationalizability and $\mathbf{AAE}_{ps}$ both select $(l3, c3)$ as the credible outcome.

We observe that, similar to the original game, the models of Klibanoff (1996), Lo (1996), and Dow and Werlang (1994) can also select $(l1, c1)$ as an equilibrium.[1] Let δ_z denote the *degenerate probability concentrated at z*. In this setting, by choosing $\beta_1 = \{\delta_{C_1}\}$, $\beta_2 = \{\delta_{L_1}\}$, $\sigma_1 = \delta_{L_1}$, and $\sigma_2 = \delta_{C_1}$, we find that the quadruple $(\sigma_1, \sigma_2, \beta_1, \beta_2)$ constitutes an equilibrium under uncertainty aversion in the sense of Klibanoff (1996), for the original game. Moreover, it is straightforward to verify that (β_1, β_2) defines a Nash equilibrium under uncertainty according to Dow and Werlang (1994), as well as a *belief equilibrium* in the sense of Lo (1996). In contrast, within our framework, $(l3, c3)$ uniquely emerges as the only equilibrium outcome, as shown in the matrix on the right.

A more instructive case arises when we consider the following modification of the original game:

original payoffs (u_1, u_2)			game payoffs (U_1, U_2)		
$(-1,-1)$	$(-1,-1)$	$(-1,-1)$	$(-1,-1)$	$(-1,4)$	$(-1,4)$
$(-1,-1)$	$(4,6)$	$(6,4)$	$(4,-1)$	$(4,4)$	$(4,4)$
$(-1,-1)$	$(6,4)$	$(4,6)$	$(4,-1)$	$(4,4)$	$(4,4)$

where the sub-matrix

$(4,6)$	$(6,4)$
$(6,4)$	$(4,6)$

represents a modification of the *matching pennies game*. The unique Nash equilibrium of this game is $(l1, c1)$, which also qualifies as an equilibrium under ambiguity à la Klibanoff (1996); Lo (1996); Dow and Werlang (1994). However, under our model, the set of equilibria is given by $\{(l2, c2), (l2, c3), (l3, c2), (l3, c3)\}$. In other words, player 1 is indifferent between playing $l2$ or $l3$, and player 2 between $c2$ or $c3$. In this modified matching pennies game, the unique Nash equilibrium $(l1, c1)$ fails under ambiguity-averse rationalizability. The model instead yields the equilibrium set $\{(l2, c2), (l2, c3), (l3, c2), (l3, c3)\}$. This outcome appears more realistic, as it reflects the intuitive resolution of the matching pennies game: players are indifferent among their options, and all resulting strategy profiles constitute equilibria in the sense of Definition 1.

The following example presents the *minimum-effort game*, as discussed in Dominiak and Eichberger (2016), which itself simplifies an earlier example from Huyck, Battalio, and Beil (1990).

[1] It is worth noting that Dow and Werlang (1994) originally formulated their equilibrium concept in terms of capacities rather than traditional probability measures.

Example 2. There are two players, $N = \{1, 2\}$, each of whom chooses an effort level of 1, 2, or 3—represented by actions $l1, l2, l3$ for player 1 and $c1, c2, c3$ for player 2. Each player's payoff is given by twice the minimum of the two efforts minus their own contribution, as shown in the original payoff matrix:

original payoffs (u_1, u_2)			game payoffs (U_1, U_2)		
$(1,1)$	$(1,0)$	$(1,-1)$	$(1,1)$	$(1,2)$	$(1,3)$
$(0,1)$	$(2,2)$	$(2,1)$	$(2,1)$	$(2,2)$	$(2,3)$
$(-1,1)$	$(1,2)$	$(3,3)$	$(3,1)$	$(3,2)$	$(3,3)$

It is easy to verify that the original game admits the following set of Nash equilibria:

$$\{(l1, c1), (l2, c2), (l3, c3)\}.$$

However, our model refines this set by selecting only the dominant equilibrium, namely $(l3, c3)$, as the unique solution.

The Nash equilibrium arises through the selection of all fixed points of mutual best responses. Rationalizability selects strategies surviving iterated deletion of dominated strategies, consistent with common knowledge of rationality. Finally, $\mathbf{AAE}_{ps}$ proposes a refinement that restricts beliefs to rationalizable deviations and evaluates them cautiously via worst-case payoff. Therefore, $\mathbf{AAE}_{ps}$ can be reframed as an *ambiguity-sensitive rationalizability refinement of Nash equilibrium.*

3.2 Properties of $\mathbf{AAE}_{ps}$

We now present some basic and straightforward properties of the equilibrium introduced in Definition 1. To ensure that all involved objects are properly defined, we assume throughout the following properties that:

[**A2**] All the utilities u_i are continuous, the strategy spaces X_i are compact, and the beliefs $x \mapsto C_i(x) \subset \mathcal{P}(X_{-i})$ have weak*-closed nonempty values.

The properties of the $\mathbf{AAE}_{ps}$ depend fundamentally on the structure of the players' belief correspondences. As we will illustrate, the $\mathbf{AAE}_{ps}$ framework encompasses and generalizes several well-known equilibrium concepts, contingent on the specific form of these correspondences. Importantly, our formulation is intentionally broad and not limited to particular specifications of beliefs. Therefore, any limitations or potential drawbacks associated with the $\mathbf{AAE}_{ps}$ equilibrium should be interpreted as stemming from specific modeling choices regarding belief correspondences, rather than from intrinsic shortcomings of the general framework itself.

We begin by examining the following case (where $\delta_{x_{-i}}$ denotes the degenerate probability concentrated at x_{-i}):

[**C1**] $\forall i \in N, \forall x \in X, C_i(x) = \{\delta_{x_{-i}}\}$.

Property 1. *Under* **C1***, the* $\mathbf{AAE}_{ps}$ *coincides with the Nash equilibrium.*

The proof is omitted as it is straightforward.

Thanks to this property, it is legitimate to view the $\mathbf{AAE}_{ps}$ as a natural generalization of the Nash equilibrium to settings where players exhibit ambiguity aversion. In particular, the classical Nash equilibrium is recovered when the players' belief correspondences satisfy **C1**.

Consider, now, the next condition:

[**C2**] $\forall i \in N, \forall x \in X, \delta_{x_{-i}} \in C_i(x)$.

The condition **C2** is equivalent to:

$$U_i(x) = \min \left\{ \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}),\ u_i(x) \right\}, \forall x \in X.$$

In other words, each player i perceives a payoff that is at most equal to the actual payoff, meaning $U_i(x) \leq u_i(x)$. This condition can be interpreted as reflecting an exaggerated form of ambiguity aversion.

Denote by $\mathbf{AAE}(G)$ the set of all $\mathbf{AAE}_{ps}$ of G and by $\mathbf{I}(G)$ the set of individually rational strategy profiles in G, that is, strategy profiles where no player can unilaterally guarantee themselves a strictly higher payoff, regardless of the actions of others:

$$\mathbf{I}(G) = \left\{ x \in X : \forall i \in N, \forall x'_i \in X_i, \exists y_{-i} \in X_{-i},\ u_i(x'_i, y_{-i}) \leq u_i(x) \right\}.$$

Property 2. *Under* **C2***,* $\mathbf{AAE}(G) \subset \mathbf{I}(G)$*.*

This property means that any $\mathbf{AAE}_{ps}$ is individually rational.

Proof. If $x \in \mathbf{AAE}(G)$, then by definition $U_i(x'_i, x_{-i}) \leq U_i(x)$, for all $i \in N$ and all $x'_i \in X_i$. Under **A2**, it follows directly that, for every $x'_i \in X_i$, there exists $y_{-i} \in X_{-i}$ such that $u_i(x'_i, y_{-i}) \leq U_i(x'_i, x_{-i}) \leq U_i(x) \leq u_i(x)$. This implies that $x \in \mathbf{I}(G)$, meaning x is individually rational. $\square$

Let us denote by $\mathbf{sN}(G)$ the set of strict Nash equilibria of G, that is the set of strategy profiles $x \in X$ such that: $\forall i \in N, \forall x'_i \in X_i \setminus \{x_i\}, u_i(x'_i, x_{-i}) < u_i(x)$.

For the purpose of establishing the next result, we need the following condition, stronger than **C2**:

[**C3**] $\forall i \in N, \forall x \in X, C_i(x) = C_i^1(x) \cup \{\delta_{x_{-i}}\}$ and $C_i^1(x) \subset \mathcal{P}(\beta_{-i}(x))$.

Property 3. *Under* **C3***,* $\mathbf{sN}(G) \subset \mathbf{AAE}(G)$*.*

Proof. First, observe that **C3** implies that, for all $i \in N$, all $x \in X$, we have $u_i(x) \geq U_i(x)$. Now, let $x \in \mathbf{sN}(G)$ and assume by contradiction that $x \notin \mathbf{AAE}(G)$. Then, there exists $x'_i \in X_i$ such that $U_i(x'_i, x_{-i}) > U_i(x)$. From **C3**, we know that $u_i(x'_i, x_{-i}) \geq U_i(x'_i, x_{-i}) > U_i(x)$. Since $x \in \mathbf{sN}(G)$, by definition $\beta_i(x) = \{x_i\}$, for every $i \in N$. Applying **C3** again, we obtain $C_i(x) = \{\delta_{x_{-i}}\}$, so: $U_i(x) = u_i(x)$. Thus, the inequality becomes: $u_i(x'_i, x_{-i}) \geq U_i(x'_i, x_{-i}) > U_i(x) = u_i(x)$, which is a contradiction. Therefore, we conclude that $\mathbf{sN}(G) \subset \mathbf{AAE}(G)$. $\square$

3.3 Existence of $\mathbf{AAE}_{ps}$

Given the complexity of the modified utility functions U_i, it is reasonable to anticipate that the existence conditions for equilibrium may be restrictive. We first establish equilibrium existence within a specific class of games—namely, finite games satisfying **A1**—before presenting a more general existence result.

Proposition 1. *Under* **A1***, suppose $N = \{1, 2\}$, each strategy set X_i is finite an endowed with the discrete topology, and $C_i(x) = \mathcal{P}(\beta_{-i}(x))$, for $i \in N$. Then, the game G admits at least one* $\mathbf{AAE}_{ps}$.

Proof. Let us first observe that under the assumptions of the proposition, C_i depends only on x_i. Specifically, for every $x_i \in X_i$ and all $x^1_{-i}, x^2_{-i} \in X_{-i}$, we have: $C_i(x_i, x^1_{-i}) = C_i(x_i, x^2_{-i})$. This follows from the fact that $\beta_{-i}(x_i, x_{-i}) = \arg\max_{y_{-i} \in X_{-i}} u_{-i}(x_i, y_{-i})$ does not depend on x_{-i}. Consequently, $C_i(x_i, x_{-i})$ depends only on x_i, and we can simply write $C_i(x_i)$ instead of $C_i(x_i, x_{-i})$. The modified utilities introduced earlier thus reduce to:

$$U_i(x_i, x_{-i}) = \inf_{p \in C_i(x_i)} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}) = U_i(x_i).$$

In other words, U_i depends solely on player i's own strategy x_i. Since each X_i is finite, the functions U_i attain their maximum at some $\bar{x}_i \in X_i$. The profile $(\bar{x}_i)_{i \in N}$ is clearly an $\mathbf{AAE}_{ps}$. $\square$

Next, we present a general existence result. In what follows, the spaces of probability measures $\mathcal{P}(X_{-i})$ are endowed with the weak*-topology, and all corresponding product spaces are equipped with the associated product-topology.

Theorem 2. *Assume that the following conditions are satisfied. For all $i \in N$:*
(i) *The strategy set X_i is a convex, compact subset of a locally convex topological vector space.*
(ii) *The payoff functions $u_i(\cdot)$ are continuous, and for every $x_{-i} \in X_{-i}$, the mapping $x_i \mapsto u_i(x_i, x_{-i})$ is concave.*
(iii) *The belief correspondences C_i are continuous.*[2]
(iv) *For all $x, t \in X$, the following inequality holds:*

$$U_i(x) = \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}) \leq \inf_{p \in C_i(t_i, x_{-i})} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}).$$

Then, the game G admits at least one $\mathbf{AAE}_{ps}$.

[2] A correspondence T from a topological space Z to another topological space W is said to be *continuous* if it is both lower semicontinuous and upper semicontinuous; see Castaing and Valadier (1977) for further details.

Condition Th.2(iv) expresses the idea that, given a *status quo* strategy profile $x \in X$, player i always perceives a lower or equal payoff when evaluating x using the belief set $C_i(x)$ corresponding to the *status quo*, compared to the perceived payoff obtained by evaluating the same strategy x_i under the belief set $C_i(t_i, x_{-i})$ associated with the alternative strategy profile (t_i, x_{-i}) resulting from a unilateral deviation by player i from the *status quo*.

Proof. We begin by introducing the following real-valued functions V and W, both defined on X^2:

$$V(t,x) = \sum_{i=1}^{n} U_i(t_i, x_{-i}) - \sum_{i=1}^{n} U_i(x), \tag{1}$$

and

$$W(t,x) = \sum_{i=1}^{n} \inf_{q \in C_i(x)} \int_{X_{-i}} u_i(t_i, y_{-i})\, dq(y_{-i}) - \sum_{i=1}^{n} U_i(x).$$

From condition Th.2(i), it follows immediately that:

$$\forall\, t, x \in X, V(t,x) \leq W(t,x). \tag{2}$$

Conditions Th.2(i-iii) imply that the requirements for applying Fan's minimax inequality are satisfied for the function $W(t,x)$. Specifically, the mapping $x \mapsto W(t,x)$ is continuous[3] by virtue of the Berge maximum principle, and the mapping $t \mapsto W(t,x)$ is concave. Therefore, there exists a strategy profile $\bar{x} \in X$ such that $W(t,\bar{x}) \leq 0$, for all $t \in X$. Using inequality (2), it follows that: for all $t \in X$, $V(t,\bar{x}) \leq 0$.
Now, fix an arbitrary player $i \in N$, and for any $t_i \in X_i$, set $t_j = \bar{x}_j$, for all $j \neq i$. Substituting into the definition of $V(t,\bar{x})$, we obtain: $V(t,\bar{x}) = U_i(t_i, \bar{x}_{-i}) - U_i(\bar{x}) \leq 0, \forall t_i \in X_i$. This implies that $\bar{x}$ is an $\mathbf{AAE}_{ps}$. □

Even though Theorem 2 is stated without explicitly assuming **A1**, one might wonder whether condition Th.2(iii) is easily satisfied under the **A1** setting. In particular, for each $i \in N$, the mapping $x \mapsto \prod_{j \neq i} \mathcal{P}(\beta_j(x))$ is not generally continuous. Consequently, assuming the existence of a continuous set-valued selection for C_i may be a demanding requirement. It is also worth noting that condition Th.2(iv) is always satisfied when $C_i(\cdot)$ does not depend on x_i, ie. $C_i(x) = C_i(y_i, x_{-i})$ for all $x \in X$ and $y_i \in X_i$. The following example illustrates that Th.1(iv) can still hold even without imposing this restriction.

Example 3. Let $N = \{1, 2\}$, $X_i = [-1, 1]$, for $i \in N$, and define the utility functions as follows: $u_1(x_1, x_2) = x_1 - x_2$ and $u_2(x_1, x_2) = \min\{x_1, x_2\}$. Assume that, for every $x \in X$, the belief correspondences are given by $C_i(x) = \mathcal{P}(\beta_{-i}(x))$. It is straightforward to verify that $\beta_2(x_1, x_2) = [x_1, 1]$, which depends on x_1. Consequently, $C_1(x_1, x_2) = \mathcal{P}([x_1, 1])$ also depends on x_1. The modified utility for player 1 becomes:

$$U_1(x_1, x_2) = \inf_{p \in \mathcal{P}([x_1,1])} \int_{X_2} u_1(x_1, y_2)\, dp(y_2) = x_1 - 1.$$

[3]A detailed proof of the continuity of the modified utilities U_i, for all $i \in N$, is given in the proof of Theorem 2 below.

For any $t_1 \in X_1$, we similarly obtain:

$$\inf_{p \in C_1(t_1, x_2)} \int_{X_2} u_1(x_1, y_2)\, dp(y_2) = \inf_{p \in \mathcal{P}([t_1, 1])} \int_{X_2} u_1(x_1, y_2)\, dp(y_2) = x_1 - 1.$$

Thus, Th.2(iv) is satisfied for U_1. For U_2, observe that: $C_2(x_1, x_2) = \mathcal{P}(\beta_1(x_1, x_2)) = \mathcal{P}(\{1\}) = \{\delta_1\}$, meaning $C_2(x_1, x_2)$ does not depend on x. Consequently, Th.2(iv) is satisfied for U_2 as well.

4 The mixed strategy case

In this section, each player $i \in N$ adopts a mixed strategy, represented by a probability measure $\mu_i \in \mathcal{P}(X_i)$. To accommodate mixed strategies, we associate with each player i a *canonical mixed-extension utility function* V_i, defined for every mixed strategy profile $\mu = (\mu_1, \ldots, \mu_n) \in \prod_{i \in N} \mathcal{P}(X_i)$ by:

$$V_i(\mu_1, \ldots, \mu_n) = \int_X \left[\inf_{p \in C_i(x)} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}) \right] d \otimes_{j \in N} \mu_j(x) = \int_X U_i(x)\, d \otimes_{j \in N} \mu_j(x),$$

where the expression is assumed to be well-defined. The technical details ensuring the well-posedness of V_i will be addressed later in this section.

Definition 2. *An* exogeneous ambiguity aversion mixed strategy equilibrium ($\mathbf{AAE}_{ms}$) *for G is a strategy profile $\mu \in \prod_{i \in N} \mathcal{P}(X_i)$ such that: for all $i \in N$, for all $\mu_i' \in \mathcal{P}(X_i)$, $V_i(\mu_i', \mu_{-i}) \leq V_i(\mu)$.*

Let Y be a compact topological space. We denote by $ca(Y)$ the space of countably additive set functions defined on the Borel σ-algebra of Y, and by $\mathcal{C}(Y)$ the space of continuous real-valued functions on Y, endowed with the sup-norm-topology. More generally, unless stated otherwise, all topological spaces considered are equipped with their respective Borel σ-algebras generated by the prevailing topologies, and all product spaces are endowed with the corresponding product σ-algebra.

Theorem 3. *Assume that the following conditions are satisfied. For all $i \in N$:*
(i) *The strategy set X_i is a compact metric space.*
(ii) *The payoff function u_i is continuous.*
(iii) *The belief correspondences C_i are continuous with closed, nonempty values.*
(iv) *The range spaces of probabilities are equipped with the weak* topologies $\sigma(ca(X_{-i}), \mathcal{C}(X_{-i}))$, for each $i \in N$.*
Then, all the mixed extension payoff functions V_i are well-defined and continuous, and the game G admits a $\mathbf{AAE}_{ms}$.

Proof. Let us first endow the probability spaces containing players' beliefs $\mathcal{P}(X_{-i})$, for each $i \in N$, with their corresponding induced weak* topologies $\sigma(ca(X_{-i}), \mathcal{C}(X_{-i}))$. It is worth noting that these topologies are metrizable since the spaces $\mathcal{C}(X_{-i})$ are separable.

Consequently, the spaces $\mathcal{P}(X_{-i})$, for all $i \in N$, are compact metric spaces. Under these assumptions, consider the functional $F : (x_i, \mu) \mapsto \int_{X_{-i}} u_i(x_i, y_{-i})\, d\mu(y_{-i})$ defined on $X_i \times \mathcal{P}(X_{-i})$. We claim that F is continuous. To see this, let $(x_i^n)_{n\in\mathbb{N}}$ be a sequence in X_i converging to x_i, and let $(\mu_n)_{n\in\mathbb{N}}$ be a sequence in $\mathcal{P}(X_{-i})$ converging to μ in the weak*-topology. Then, for every $n \in \mathbb{N}$,

$$
\begin{aligned}
|F(x_i^n, \mu_n) - F(x_i, \mu)| \leq\ & \underbrace{\left| \int_{X_{-i}} u_i(x_i^n, y_{-i})\, d\mu_n(y_{-i}) - \int_{X_{-i}} u_i(x_i, y_{-i})\, d\mu_n(y_{-i}) \right|}_{(I)} \\
& + \underbrace{\left| \int_{X_{-i}} u_i(x_i, y_{-i})\, d\mu_n(y_{-i}) - \int_{X_{-i}} u_i(x_i, y_{-i})\, d\mu(y_{-i}) \right|}_{(II)}.
\end{aligned}
$$

It is clear that term (II) converges to zero by the weak* convergence of μ_n to μ. To show that term (I) converges to zero, observe that the family $\mathcal{H} = \{\mu_n : n \in \mathbb{N}\}$ defines a family of continuous linear functionals on $\mathcal{C}(X_{-i})$ endowed with the sup-norm. Since each μ_n is a probability measure,

$$
\forall \varphi \in \mathcal{C}(X_{-i}), \left| \int_{X_{-i}} \varphi\, d\mu_n \right| \leq \|\varphi\|_{\mathcal{C}(X_{-i})}.
$$

This implies that $\mathcal{H}$ is uniformly bounded. By the Banach-Steinhaus theorem, $\mathcal{H}$ is equicontinuous. Furthermore, $u_i(x_i^n, \cdot)$ converges to $u_i(x_i, \cdot)$ uniformly over X_{-i}, implying that:

$$
\lim_{n\to\infty} \sup_{k\in\mathbb{N}} \left| \int_{X_{-i}} u_i(x_i^n, y_{-i})\, d\mu_k(y_{-i}) - \int_{X_{-i}} u_i(x_i, y_{-i})\, d\mu_k(y_{-i}) \right| = 0.
$$

Therefore, term (I) also converges to zero, establishing the continuity of F. Next, since C_i is continuous with compact values in a metric space, Berge's Maximum Principle ensures that:

$$
U_i(x) = \inf_{p \in C_i(x)} \int_{X_{-i}} u_i(x_i, y_{-i})\, dp(y_{-i}) = \inf_{p\in C_i(x)} F(x_i, p),
$$

which implies that U_i is continuous. The integral of U_i over X defines a linear functional on $ca(X)$, which is continuous with respect to the weak*-topology $\sigma(ca(X), \mathcal{C}(X))$. Moreover, the embedding $(\mu_1, \ldots, \mu_n) \mapsto \mu_1 \otimes \ldots \otimes \mu_n$ from $\prod_{i\in N} ca(X_i)$ into $ca(X)$ is continuous for the weak* topologies.[4] Thus, the extended utilities V_i, for all $i \in N$, are continuous on $\prod_{i\in N} ca(X_i)$ endowed with the product weak* topologies. Furthermore, for each fixed μ_{-i}, the mapping $\mu_i \mapsto V_i(\mu_i, \mu_{-i})$ is linear. Finally, since the strategy spaces $\mathcal{P}(X_i)$, for all $i \in N$, are convex, weak* compact subsets of locally convex spaces, the existence of a $\mathbf{AAE}_{ms}$ follows from classical Nash equilibrium existence results. $\square$

[4]Indeed, functions in $\mathcal{C}(X)$ can be uniformly approximated by sequences of finite linear combinations of functions belonging to $\prod_{i\in N} \mathcal{C}(X_i)$. As a consequence, the space $\prod_{i\in N} \mathcal{C}(X_i)$ is dense in $\mathcal{C}(X)$ (see, Balder (1988), Lemma 2.6).

5 Conclusion

In this paper, we introduced a novel equilibrium concept for games with ambiguity-averse players. Unlike related approaches in the literature, our equilibrium treats ambiguity as an exogenous feature derived from the structure of the game itself, while preserving the core of strategic interaction among players. We have demonstrated that this equilibrium exists under relatively weak conditions in the mixed-strategy setting. In contrast, existence in pure strategies may require stronger assumptions due to the complex structure of the modified utilities and belief correspondences.

Although not our primary objective, we also illustrated through several well-known finite games that our framework offers a natural mechanism for refining and selecting desirable Nash equilibria in the original game. This suggests that our equilibrium concept can serve both as a tool for analyzing ambiguity and as a refinement device within strategic environments.

The present work constitutes a first step toward a broader research agenda. We believe that further investigation of the proposed equilibrium is both promising and necessary. In particular, uncovering additional properties may reveal its potential for applications to real-world strategic situations, which are inherently shaped by uncertainty and ambiguity. On the theoretical side, several directions remain open, such as weakening the assumptions underlying Theorems 1 and 2, or establishing sufficient conditions for the existence and uniqueness of the **AAE**. These questions offer fertile ground for future research.

References

Bade, S., Ambiguous Act Equilibria, *Games and Economic Behavior* 71 (2011), 246-60.

Balder, E. J., Generalized Equilibrium Results for Games with Incomplete Information. *Mathematics of Operations Research* 13 (1988), 265-76.

Bernheim, D., Rationalizable Strategic Behavior, *Econometrica* 52 (1984), 1007-28.

Castaing, C., and Valadier, M., Convex Analysis and Measurable Multifunctions, *Lecture Notes in Mathematics* 580, Springer-Verlag Berlin Heidelberg 1977.

Dominiak, A., and Eichberger, J., Equilibrium under Ambiguity (EUA) for Belief Functions, [https://www.wiwi.tu-dortmund.de/wiwi/de/forschung/kolloquien/vwk/VWL-Kolloquium-2016-Eichberger.pdf].

Dow, J., and Werlang, S. R. C., Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction, *Journal of Economic Theory* 64 (1994), 305-24.

Eichberger, J., and D. Kelsey, D., Optimism and Pessimism in Games, *International Economic Review* 55 (2014), 483-505.

Eichberger, J., Kelsey, D., and Schipper, B.C., Ambiguity and Social Interaction, *Oxford Economic Papers* 61 (2009), 355-79.

Ellsberg, D., Risk, Ambiguity and the Savage Axioms, *Quarterly Journal of Economics* 75 (1961), 643-69.

Gilboa, I., and Schmeidler, D., Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics* 18 (1989), 141-53.

He, W., and Yannelis, N.C., *Equilibrium theory under ambiguity*, The School of Economics Discussion Paper Series 1307(2013), Economics, The University of Manchester, [http://hummedia.manchester.ac.uk/schools/soss/economics/discussionpapers/EDP-1307.pdf].

Hill, B. A non-Bayesian Theory of state-dependent Utility, *Econometrica* 87 (2019), 1341-66.

Huyck, J.V., Battalio, R., and Beil, R., Tacit Coordination Games, strategic uncertainty and coordination failure, *American Economic Review* 80 (1990), 234-48.

Kajii, A., and Ui, T., Incomplete information games with multiple priors, *The Japanese Economic Review* 56 (2005), 332-51.

Karni, E., Schmeidler D., and Vind, K., On State Dependent Preferences and Subjective Probabilities, *Econometrica* 51 (1983), 1021–31.

Klibanoff, P., *Uncertainty, Decision, and Normal Form Games*, Technical report, Northwestern University, 1996.

Lo, K. C., Equilibrium in Beliefs under Uncertainty, *Journal of Economic Theory* 71 (1996), 443-84.

Luce, R. D., and Raiffa, H., *Games and Decisions. Introduction and Critical Survey*, Wiley, New York 1957.

Marinacci, M., Ambiguous games, *Games and Economic Behavior* 31 (2000), 191-219.