On Hurwicz–Nash Equilibria of Non–Bayesian Games under Incomplete Information*

Patrick Beißner† and M. Ali Khan‡

December 5, 2018

Abstract

We consider finite-player simultaneous-play games of private information in which a player has no prior belief concerning the information under which the other players take their decisions, and which he therefore cannot discern. This dissonance leads us to develop the notion of Hurwicz-Nash equilibria of non-Bayesian games, and to present a theorem on the existence of such an equilibrium in a finite-action setting. Our pure-strategy equilibrium is based on non-expected utility under ambiguity as developed in Gul and Pesendorfer (2015). We do not assume a linear structure on the individual action sets, but do assume private information to be "diffused" and "dispersed." The proof involves a multi-valued extension of an individual's prior to the join of the finest σ-algebra $\mathcal{F}$ of the information of the other players, and hinges on an absolute-continuity assumption on an individual's belief with respect to the extended beliefs on $\mathcal{F}$. (149 words)

Key words and phrases: Non-Bayesian games, Hurwicz-Nash equilibria, Knightian uncertainty, private information, private beliefs, ambiguous beliefs.

JEL Classification Numbers: C62, D50, D82, G13.

*The authors are grateful to Yongchao Zhang for ongoing collaboration, and to Achile Basile, Ying Chen, Amanda Friedenberg, Simon Grant, Faruk Gül, Peter Hammond, Osama Khan, Jong Jae Lee, Alejandro Manelli, Ed Prescott, Frank Riedel, Eddie Schlee, Vernon Smith, Ron Stauber, Max Stinchcombe and Metin Uyanik for stimulating correspondence and conversation. Preliminary versions of this paper were presented at departmental seminars at *The University of Queensland* (November 7, 2018), *Arizona State University* (October 30, 2017) and *Fundaçao Getulio Vargas* (April 10, 2018); and at the *17th SAET Conference* in Faro (June 30, 2017), the *6th Workshop in Equilibrium Analysis* in Naples (January 27, 2018), the *IMS-NUS Workshop in Game Theory*, and the *27th EWGET Workshop* in Paris (June 28, 2018). The authors warmly thank the Editor and her two anonymous referees for their careful reading: the second example, a fuller and more careful elaboration of the proof, and the connection to the literature on "imprecise probabilities" all presented herein, resulted from it.

†Research School of Economics, The Australian National University, Canberra Australia. **E-mail** patrick.beissner@anu.edu.au

‡Department of Economics, The Johns Hopkins University, Baltimore, MD 21218. **E-mail** akhan@jhu.edu

Contents

There is a tendency in our planning to confuse the unfamiliar with the improbable.[1]

Thomas Schelling (1962)

1 Introduction

Incomplete information and uncertainty in economic theory date at least to the classic texts of Knight (1921) and Von Hayek (1937), but their formal study in the context of non-cooperative game theory surely originates in Harsanyi (1967/68) and Aumann (1974), and, in the context of Walrasian general equilibrium theory, in Radner (1968). This subsequent, more formal, work led in settings with a finite number of agents to individual (private) σ-algebras as an added characteristic of each agent, and to the requirement that individual strategies in the context of games, and individual demands in the context of economies, be measurable with respect to individual sub-σ-algebras of the common and publicly-known sample space. However, in the modelling of game-theoretic situations in which one player has by necessity to infer and "integrate out" another player's influence on her payoffs before taking her own decision, the pervasive point of departure has been to rely on a common prior on the pooled information of all the players in the game. Put another way, however impressive recent advances in the theory of the existence of pure-strategy Bayesian-Nash equilibria in Bayesian games may have been,[2] they have all relied on the common prior assumption. As such, they are blemished by the contradiction that there are common public beliefs that a player can invoke on private information that by necessity and definition he cannot discern.

This unsatisfactory state of affairs led to rather vigorous questioning by Morris (1995). He asked:

> Why is (it that) common priors are implicit or explicit in the vast majority of the differential information literature in economics and game theory? Why has the economic community been unwilling, in practice, to accept and actually use the idea of truly private probabilities in much the same way that it did accept the idea of private utility functions? After all, in (Savage's expected utility theory), both the utilities and probabilities are derived separately for each decision maker. Why were the utilities accepted as private, and the probabilities not?

In the context of economies, Khan, Sun, Tourky, and Zhang (2008) developed a topology on the space of individual information and individual beliefs in an attempt to incorporate private probabilities, but the question of the existence of a Walrasian equilibrium foundered on the

[1]See his "Forward" to R. Wohlstetter (1962) *Pearl Harbor: Warning and Decision* (Stanford: Stanford University. Also cited on pages 41 and 111 in Mirowski and Nik-Khah (2017) but with the insertion ".. in our theories of planning to confuse ... "

[2]For these advances, see the announcement of Khan and Sun (1996), and the subsequent literature detailed in He, Sun, and Sun (2017) and in He, Sun, Sun, and Zeng (2015).

fact that agents, in the use of equilibrium prices in the updating of their private information, led to a mapping that could not be guaranteed to be upper hemi-continuous.[3] In the context of games, the thrust, and the achievements, of the analytical work have been in the service of mechanism design and auction theory, and therefore the common prior assumption was hardly an irritant.[4] In this paper, we leave economies and design problems aside and take a first step in answering Morris' call in the context of finite-player simultaneous-play games of private information.

Even with a limitation to this game-theoretic register, there is the conceptual question of how a player is to move from subjective (private probabilities) beliefs on his subjective (private σ–algebras) information regarding the set of states that he can discern, to objective (public) beliefs on the information available to the others in the game and which he cannot, by definition, discern. The answer that we pursue in this paper is simply that he has to proceed by inference, and that this inference necessarily involves ambiguity and imprecision. Each player has little option but to extend his private probability on his private information to a possible *set* of probabilities on *all* of the available information in the game; and rather than an expectation taken with respect to a *single* Bayesian prior, he has to modify his objective function in accordance with this extended set of probabilities. In an early research proposal,[5] Hurwicz had already written:

> The emphasis is on ... the *technology* of the processes whereby decisions are reached and the choices are made. [W]hen the information processing are taken explicitly into account it is found that the concept of "rational action" is modified. This is true when applied to the action of a single individual, but it becomes particularly interesting when considered in situations involving many persons ... The uncertainty need not be generated by external factors like weather prospects: it may be man-made.

This is to say that the solution concept is itself changed: in the context of this paper, from a Bayesian-Nash equilibrium to a Hurwicz-Nash equilibrium, and to the consequent focus on the *extension* of individual beliefs as opposed to the *restriction* or an *updating* of an exogenously-assumed universal public belief on the totality of the privately-available, and presumably secret, information. In the rest of this introduction, we further elaborate our contribution in the light of the antecedent classical and modern literature.

---

[3] This is so even with a common prior formalizing identical beliefs; see the papers of Cotter cited in Khan, Sun, Tourky, and Zhang (2008) – individual priors only exacerbated the analytical difficulties.

[4] This normative, and planning, motivation of the work is explicit in the writings of Hurwicz. Hurwicz (1979) distinguishes between frameworks arising out of *structural change* and *control theory*, and writes, "Both are *normative* in spirit. They do not accept the status quo, but rather look for modes of intervention that would bring the system as close to optimality as possible;" also see Figure 5 in Hurwicz (2007), and Reiter (2009).

[5] Mirowski and Nik-Khah (2017) date the writing to 1951; see their page 86 and its reference to Marschak's UCLA papers.

Whereas the tone and the letter of our exploration is, at least, partly in keeping with the classical texts of Knight and Hayek, some subtleties ought to be noticed and emphasized. To be sure, the Knightian difference between risk and uncertainty has been pervasive in the decision-theory literature, an obligatory citation, so to speak, but the underlying theme has been that the multiplicity of beliefs that emerge testify to the formalization of ambiguity and ignorance, and give a quantitative measurement to Knightian uncertainty. However, one can legitimately argue that Knight, and for that matter Keynes (1937), were making an argument for a phenomena that was resistant to the representation of any sort of a "calculative or purposive" disposition and conduct. And as such, the ambiguity literature that has emerged in the last twenty or so years ought to stand confidently on his own rather than needlessly drawing specious authority from Knight and Keynes. Nevertheless, given that we consider one-shot, simultaneous play games in this paper, our formalization may be consistent with the following qualification of Knight's:

> When an individual instance only is at issue, there is no difference for conduct between a measurable risk and an unmeasurable uncertainty. The individual ... throws his estimate of the value of an opinion into the probability form of 'a successes in b trials' ... and 'feels toward it as toward any other probability situation.[6]

But again, the matter is hardly so simple. The situation we study here is unlike "that in which prospect of a European war is uncertain or the price of copper and the rate of interest twenty years hence";[7] the uncertainty modelled here is the uncertainty of the other's ambiguous evaluation, perhaps even the certainty, of the other's information. If one follows the vernacular of Hayek's classic 1937 and 1945 essays in which he sighted the "division of knowledge" as being of equal importance for economic theory as the Smithian concept of the "division of labor", the question reduces to the uncertainty of how "local knowledge" regarding an individual trader's information can be stitched into, and acted upon, "global knowledge" representing the ensemble of society's knowledge. The theorem reported below is focused on this reduction, one that engages and copes with what Hurwicz terms "man-made uncertainty."[8]

To be sure, our question leads us naturally to the rich literature of the last thirty years on 'ambiguity and the Bayesian paradigm' so authoritatively surveyed in Gilboa and Marinacci (2016), and also to Bayesian games with diffused and dispersed information.[9] The former

[6] Cited in Al-Najjar and Weinstein (2015); also see Leroy and Singell (1987) for an imaginative interpretation pertaining to incomplete markets. We defer to future work a detailed textual exegesis of Knight's classic.

[7] These are the well-known and often-quoted words of Keynes (1937).

[8] Again, we defer to future work a detailed textual exegesis of Hayek's writings, but one can surely begin by engaging narratives with different points of origin and different takes on trajectories presented in Kamenca (2017), and in Chapters 6 and 14 in Mirowski and Nik-Khah (2017).

[9] This literature, initiated by Radner and Rosenthal (1982), is expressed in fuller maturity in Milgrom and

has emerged under the rubric of Knightian-Bayesian (ambiguous) games, and rests on the aftermath of Machina's (1982) move of expected utility theory away from the independence axiom, Bewley's (1986) initiation of a research program of integrating Knightian concerns in to the broad gamut of economic theory, and the jettisoning of the additivity postulate in the formalization of subjective probability by Schmeidler (1982, 1989) and Gilboa (1987).[10] In what is perhaps the originary benchmark on "ambiguous games," Marinacci (2000) quotes from Knight and presents his motivation work as follows:

> The action which follows upon an opinion depends as much upon the amount of confidence in that opinion as it does upon the favorableness of the opinion itself. Given the importance of ambiguity at the individual level, it seems natural to investigate the role it plays, if any, in strategic interactions.

In an important footnote,[11] he refers to previous work that attempted to "integrate into game theory recent advances in nonexpected utility theory." Azrieli and Teper (2011), and the work of Kajii and Ui (2005) and Bade (2011) that they refer to, can be seen as important updates on Marinacci's work. One may also add in this connection Ellsberg games of Riedel and Sass (2014) and the Savage games of Grant, Meneghel, and Tourky (2016). Stinchcombe (2008) is another early attempt.

This is undoubtedly useful and important work, but the result reported in this paper is not primarily motivated by it. Rather than yet another incorporation of non-expected utility in the canonical models of non-cooperative game theory, it is to address the dissonance between differential information formalized as personalized σ-algebras and non-differentiated beliefs in terms of a common prior. The distinctive characteristic of our approach is to address this dissonance by casting the recent work of Gul and Pesendorfer (2015) in terms of a viable equilibrium notion – that of a *Hurwicz-Nash* equilibrium – and the concept of group- or societal-rationality that is thereby articulated therein.[12] . As is by now well-understood, a Hurwicz-expected-utility maximizer works under a tripartite parametrization: a personal σ-algebra, a personal prior on the personal σ-algebra, and a personalized parameter of uncertainty aversion. He resorts to subjective expected utility theory along with the Hurwicz-criterion.[13] We return to the question as to why the Hurwicz-criterion rather than another

Weber (1985), Khan and Sun (1999) and Khan and Zhang (2014); also see the references in Footnote 2, the paragraphs below and Section 4.1.

[10]See Gilboa (2009) for a comprehensive treatment and additional framing and references.

[11]See the work of Crawford and others as cited in Footnote 1, page 192 in Marinacci (2000).

[12]See Levi (1987) for an insightful discussion of the subtleties of the meaning of group-rationality.

[13]We refer the interested reader to the axiomatization and discussion in Gul and Pesendorfer (2015); also see the antecedent treatment in Gul and Pesendorfer (2014). In this connection of representation of Hurwicz preferences, one may also cite the earlier work of Ghirardato, Maccheroni, and Marinacci (2004) and the follow-up by Eichberger, Grant, and Kelsey (2008).

in the sequel; what is relevant at this point is that our approach is decidedly non-Bayesian, one that forgoes successive updating by a single, one-time extension.

Returning to *ambiguous* games, the principal impediment to the type of result that we report lies in the fact that integration with respect to a set of priors, rather than a single prior, renders moot the linearity property of the payoffs. Thus, in their recent work on the existence of equilibrium in games with uncertainty aversion and incomplete information, Azrieli and Teper (2011) write:

> We characterize equilibrium existence in terms of the preferences of the participating players. It turns out that, given continuity and monotonicity of the preferences, equilibrium exists in every game if and only if all players are averse to uncertainty (i.e., all the functionals are quasi-concave).

The authors furnish an example of a game in which there is no pure-strategy equilibrium when there is a player with a preference functional that is not quasi-concave, and who therefore "exhibits a weak form of uncertainty loving over some range." This is surely a useful result, but one that nevertheless opens a lacuna in the literature by tying equilibrium existence issues to risk-taking behavior. In short, it necessitates a search for a context that makes possible an existence result without any such assumption on the preference functionals of the players. The result reported here underscores that a natural and viable context is readily available in the appeal to diffused and dispersed information, a formalization stemming from Aumann (1974), and pioneered by Radner and Rosenthal (1982). It is a setting that is tractable enough to face up to the lack of quasi-concavity in the preference functionals, and to show the existence of such an equilibrium in this setting. At the risk of repetition, we again emphasize that in our focus on pure-strategy equilibria, there is no role for the requirement that a player's payoffs are quasi-concave in his or her actions. This raises its own set of problems when it comes to extensions to many possible priors rather than a restriction of a single prior to many individual domains. The fact all these difficulties can be so effortlessly handled by the techniques that are already available in the antecedent literature that they need not be relegated to an appendix, is, we feel, an added advantage of our approach.

The plan of the paper then is straightforward: (i) the model and the result, (ii) two examples that illustrate how the Hurwicz-Nash equilibrium notion differs from the usual Bayesian-Nash formulation, and suggested parametrizations under which it yields identical outcomes,[14] (iii) a further framing of the result in light of alternative non-Bayesian perspectives drawn from both epistemic game theory and the literature on artificial intelligence concerning "vague and imprecise" probabilities, pointing to possible extensions that may possibly bring our basic idea to a fully mature theoretical expression, (iv) a proof of the result and the impediments that

[14]However, it bears emphasis that, as expected, the Hurwicz parameter α_i plays a crucial role in these examples, and that none of the cases is generic with respect to this α-register parametrizing the degree of optimism and pessimism.

it overcomes, and finally, (v) proofs of ancillary results that are necessary for the execution of the argument.

2 The Model and Result

For perspective, we recall some basic definitions for games with incomplete information and a common prior. In terms of the textbook categorization of Maschler, Solan, and Zamir (2013), as presented in their Chapter 9, we formulate the game in the setting due to Aumann rather than the equivalent one of Harsanyi. The rationale for this relies on the terminological simplification that results in our reliance of Hurwicz expected payoffs in the setting of a non–Bayesian game; also see Gul and Pesendorfer (2015). This is to say that instead of describing (incomplete) information as a finite partition, we represent it by means of a possibly infinite σ–algebra.

2.1 Bayesian Games in the Aumann Setting

This section recalls the basic concepts to define a Bayesian game and Bayesian–Nash equilibria under incomplete information.

Definition 1 (Basic Primitives) *A model of incomplete information is a specification* $(\mathbb{I}, \Omega, (\mathcal{F}_i)_{i\in\mathbb{I}}, P)$ *where*

- $\mathbb{I} = \{1, \ldots, I\}$ *is the set of players.*
- $\Omega = \times_{i\in\mathbb{I}} T_i$ *consists of the states of the world.* T_i *denotes the set of types of player* i.
- $\mathcal{F}_i$ *is a* σ*-algebra on* Ω*, the private information of player* $i \in \mathbb{I}$.
- P *is a probability measure on* $\mathcal{F}$*, the coarsest* σ*-algebra containing* $\cup_{i\in\mathbb{I}}\mathcal{F}_i$.

In this model *only* the private σ-algebra $\mathcal{F}_i$ of each player i reflects her incomplete informational structure. The players have a common belief on the finest σ-algebra which describes the uncertainty of any player about the types of the other players. A crucial implication behind the specification in Definition 1 is the assertion that any player has the ability to assign a probability also to those events A that are outside of the domain of her private information; that is, for every $A \in \mathcal{F}$ the number $P(A)$ is common knowledge.

According to Bayes rule, each player i is now in the position to update her belief, employ the posterior $P(\cdot|\mathcal{F}_i)$ and compute expected payoffs by virtue of a conditional probability measure. Formally, we have $P(A|\mathcal{F}_i) = E^P[1_A|\mathcal{F}_i]$ for any $A \in \mathcal{F}$.

Definition 2 (Type of Game) *A* Bayesian game with incomplete information *is a specification* $(\mathbb{I}, \Omega, (\mathcal{F}_i)_{i \in \mathbb{I}}, P)$ *where:*

- *Each agent i can take actions from a finite set A_i.*
- *Any pure strategy $s_i : T_i \to A_i$ of agent i is $\mathcal{F}_i$–measurable.*
- *The $\mathcal{F}_i$–conditional P–expected payoffs are computed by a state–dependent utility index $u_i : \Omega \times A \to \mathbb{R}$, where $A = \times_{i \in \mathbb{I}} A_i$:*

$$\mathbb{E}^P\big[u_i(\cdot, s_i, s_{-i})|\mathcal{F}_i\big] = \int_\Omega u_i\big(\omega, s_i(\omega), s_{-i}(\omega)\big)\mathrm{d}P(\omega|\mathcal{F}_i).$$

To receive a well–defined conditional expectation, we have to assume that $u_i(\cdot, a)$ is $\mathcal{F}$–measurable and P–integrable for any $i \in \mathbb{I}$ and $a \in A$.

The following equilibrium notion is standard.

Definition 3 (Equilibrium Concept) *A Bayesian–Nash equilibrium for a game with incomplete information is a set of strategies $s_i^{\mathsf{B}*} : \Omega \to A_i$, $\mathcal{F}_i$–measurable for each $i \in \mathbb{I}$, that satisfies*

$$\mathbb{E}^P\big[u_i\big(s_i^{\mathsf{B}*}, s_{-i}^{\mathsf{B}*}\big)|\mathcal{F}_i\big] \geq \mathbb{E}^P\big[u_i\big(s_i, s_{-i}^{\mathsf{B}*}\big)|\mathcal{F}_i\big] \quad P\text{–almost surely}$$

for all pure and $\mathcal{F}_i$–measurable strategies s_i of player i.

2.2 Non–Bayesian Games with Incomplete Information

To define a non–Bayesian game à la Aumann, we change the specification of Definition 1 in one crucial aspect.

Definition 4 (Basic Primitives) *A model of probabilistically incomplete information is a specification* $(\mathbb{I}, \Omega, (P_i, \mathcal{F}_i)_{i \in \mathbb{I}})$ *where:*

- *$\mathbb{I} = \{1, \ldots, I\}$ is the set of players.*
- *T_i is the set of types of player i. $\Omega = \times_{i \in \mathbb{I}} T_i$ consists of the states of the world.*
- *For any $i \in \mathbb{I}$, $\mathcal{F}_i$ is a σ-algebra of Ω.*
- *P_i is a probability measure on $(\Omega, \mathcal{F}_i)$ for each $i \in \mathbb{I}$.*

The main difference from Definition 1 stems from the change of each player's belief P_i. In the specification of Definition 4, a player can only assign a probability to those events that she is aware of. As discussed in the introduction, the ignorance about certain events $A \notin \mathcal{F}_i$

then also implies ignorance about the very chance with which that said event may occur. To put it in a different way, a non–Bayesian incomplete information specification allows for the absence of subjective probabilities that are defined on the grand σ-algebra.

The following simple example specifies a probabilistically incomplete specification for player 1.

Example 1 *For simplicity, let there be only three states $\Omega = \{a, b, c\}$. An incomplete information structure of player 1 is then $\mathcal{F}_1 = \{\{a\}, \{b, c\}\}$ with $P_1(\{a\}) = \frac{1}{3}$ and $P_1(\{b, c\}) = \frac{2}{3}$.*

To maintain the common prior assumption on common knowledge $\cap_{i\in\mathbb{I}}\mathcal{F}_i$, we fix a probability measure $\mathbb{P}$ on the finest information structure $\mathcal{F} = \sigma(\cup_{i\in\mathbb{I}}\mathcal{F}_i)$ and assume that each P_i of Definition 4 is induced by the restriction of $\mathbb{P}$ to $\mathcal{F}_i$.[15] Formally we assume

$$P_i = \mathbb{P}_{\restriction \mathcal{F}_i} \quad \text{for all } i \in \mathbb{I}.$$

We import Definition 2 to the present setting with some changes. For the definition of an equilibrium, the $\mathcal{F}_i$ measurability of the opponents' strategies s_i, does not allow player i to compute the expectation as in the Bayesian game, see the third bullet point of Definition 2. The reason is simply that her prior P_i only assigns probabilities to events in $\mathcal{F}_i$. To overcome the inability to evaluate expected payoffs depending on the opponents' strategies, we define for each player i the set

$$\mathcal{P}(\mathcal{F}_i) = \{P \in \Delta(\Omega, \mathcal{F}) : P = P_i \text{ on } \mathcal{F}_i\} \tag{1}$$

of possible priors being consistent with the given prior P_i on $\mathcal{F}_i$. Here, $\Delta(\Omega, \mathcal{F})$ denotes the set of probability measures on $(\Omega, \mathcal{F})$. We invoke Example 1 again.

Example 2 *Let there be only three states $\Omega = \{a, b, c\}$. Information $\mathcal{F}_1$ and belief P_1 are from Example 1. The set of possible extensions on $\mathcal{F} = \{\{a\}, \{b\}, \{c\}\}$ is then*

$$\mathcal{P}(\mathcal{F}_1) = \left\{ \left(\frac{1}{3}, p, \frac{2}{3} - p \right) \in \Delta(\Omega) : p \in \left[0, \frac{2}{3}\right] \right\}.$$

The expected payoffs of any player become ambiguous, since any element $\tilde{P}_i$ in $\mathcal{P}(\mathcal{F}_i)$ is a possible choice in the determination of the expected payoff $\mathbb{E}^{\tilde{P}_i}[u_i(s_i, s_{-i})]$ of player i. In contrast to Definition 2, the notion of a game is not yet fully specified. The choice of $\tilde{P}_i \in \mathcal{P}(\mathcal{F}_i)$ remains arbitrary. Consequently, we have to extend expected utility in a way to accommodate preferences for this new type of uncertainty. We do so by assuming that each player applies a Hurwicz expected payoff (henceforth HEP) W_i, based on Gul and Pesendorfer

[15] This common prior assumption on the common knowledge can be relaxed.

(2015). Towards this end, we introduce for each player the additional parameter $\alpha_i \in [0,1]$ and define an HEP by

$$W_i(s_i, s_{-i}) = \alpha_i \min_{\tilde{P} \in \mathcal{P}(\mathcal{F}_i)} \mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})] + (1-\alpha_i) \max_{\tilde{P} \in \mathcal{P}(\mathcal{F}_i)} \mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})]. \tag{2}$$

HEP of each player i has three parameters: the "prior" $P_i : \mathcal{F}_i \to [0,1]$ captures the perception of uncertainty, the usual utility index u_i encodes preferences for risk, and α_i quantifies preferences for ambiguity. Lemma 1 below establishes that HEP is well–defined under our hypotheses.

In the following formalization of a non-cooperative game-form, each player i employs a prior in $\mathcal{P}(\mathcal{F}_i)$. The main difference of this notion from that of the more conventional Bayesian game stems from the different use of private information. Instead of computing the posterior, the players in a non–Bayesian environment have to employ the individual sets of consistent extensions.

Definition 5 (Type of Game) *A game Γ with probabilistically incomplete information is a specification $(\mathbb{I}, \Omega, (\mathcal{F}_i, P_i)_{i\in\mathbb{I}})$ where:*

- *Each player i can take actions from a finite set A_i.*
- *Any pure strategy $s_i : T_i \to A_i$ of player i is a $\mathcal{F}_i$–measurable mapping.*
- *The $\tilde{P}_i$–expected payoff are computed by a state dependent utility index $u_i : \Omega \times A \to \mathbb{R}$, where $A = \times_{i\in\mathbb{I}} A_i$:*

  $$\mathbb{E}^{\tilde{P}_i}[u_i(\cdot, s_i, s_{-i})] = \int_\Omega u_i\big(\omega, s_i(\omega), s_{-i}(\omega)\big) \mathrm{d}\tilde{P}_i(\omega)$$

  where $\tilde{P}_i : \mathcal{F} \to [0,1]$ is an extension of P_i and $(s_i, s_{-i}) = (s_1, \ldots, s_I)$ denotes a given profile of pure strategies.
- *Each player has a preference for ambiguity, indexed by $\alpha_i \in [0,1]$.*

In the following equilibrium concept, every player employs the Hurwicz–expected payoff functional specified above as (2).

Definition 6 (Equilibrium Notion) *A Hurwicz–Nash equilibrium for a game with probabilistically incomplete information is a list of $\mathcal{F}_i$–measurable strategies $s_i^{\mathtt{H}*} : \Omega \to A_i$ that satisfies*

$$W_i\big(s_i^{\mathtt{H}*}, s_{-i}^{\mathtt{H}*}\big) \geq W_i\big(s_i, s_{-i}^{\mathtt{H}*}\big)$$

for all pure $\mathcal{F}_i$–measurable strategies s_i of player $i \in \mathbb{I}$.

As in a Bayesian–Nash equilibrium, we continue to assume that the strategies of each player i are measurable with respect to her private information $\mathcal{F}_i$.

2.3 Remarks

The first two remarks discuss possible variations of the definition of consistent extension, stated in (1). This change affects the notion of HEP, and consequently the equilibrium notion in Definition 6. The last two remarks discuss the limits of alternatives to HEP and the nature of non–Bayesian games.

Remark 1: The set of $\mathcal{F}_i$–consistent extensions $\mathcal{P}(\mathcal{F}_i)$ in (1) considers all possible extensions in $\Delta(\Omega, \mathcal{F})$. In fact, this may be considered in some situations as too extreme. For proving existence of a Hurwicz–Nash equilibrium, we restrict the set of possible extension of probability measures $\tilde{P}$ that are mutually absolutely continuous (sharing the same null sets) with respect to $\mathbb{P}$, denoted by $\tilde{P} \sim \mathbb{P}$. In other words: the domain of what is possible is determined by $\mathbb{P}$ and assumed to be common knowledge. $\square$

Remark 2: Departing from Remark 1 above, we can now go one step further and interpret $\mathbb{P}$ as a common reference belief. Under such an interpretation, we can assume that each player is focused on only those extensions that are "sufficiently close" to $\mathbb{P}$, where the notion of "closeness" or similarity is the classic one given by the distance of relative entropy $\mathtt{Ent}(\tilde{P}, \mathbb{P}) = \int_\Omega \ln \frac{d\tilde{P}}{d\mathbb{P}} d\mathbb{P}$ and consider only those $\tilde{P}$ that satisfy

$$\mathtt{Ent}(\tilde{P}, \mathbb{P}) \leq \eta, \qquad \text{for some number } \eta \geq 0.$$

This results in the following modified set

$$\mathcal{P}^{\mathtt{Ent}}(\mathcal{F}_i) = \left\{ \tilde{P} \sim \mathbb{P} : \tilde{P} = \mathbb{P} \text{ on } \mathcal{F}_i \text{ and } \mathtt{Ent}(\tilde{P}, \mathbb{P}) \leq \eta \right\}.$$

The extreme case $\eta = 0$, then corresponds to $\{\mathbb{P}\} = \mathcal{P}^{\mathtt{Ent}}(\mathcal{F}_i)$ and results in a Nash–Equilibrium with complete information. In the existence proof of Hurwicz–Nash equilibria that we present as Theorem 1 below, it is the case that only the convexity and closedness of $\mathcal{P}^{\mathtt{Ent}}(\mathcal{F}_i)$ matters, which is the case. Nevertheless, HEP of player i then depends on an additional parameter η_i and results in a different equilibrium. $\square$

Remark 3: The specific functional form of HEP above relies on Gul and Pesendorfer (2015). One can of course consider possible extensions of the payoff structure that are in keeping with the existence of a Hurwicz–Nash equilibrium. For example, consider a double variational payoff structure:

$$\hat{W}_i(s_i, s_{-i}) = \alpha_i \min_{\tilde{P} \in \mathcal{P}(\mathcal{F}_i)} \left(\mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})] + c^{min}(\tilde{P} \| \mathbb{P}) \right) + (1 - \alpha_i) \max_{\tilde{P} \in \mathcal{P}(\mathcal{F}_i)} \left(\mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})] - c^{max}(\tilde{P} \| \mathbb{P}) \right).$$

for some convex penalty terms $c^{max}(\cdot \| \mathbb{P})$ and $c^{min}(\cdot \| \mathbb{P})$. If $\alpha_i = 1$ we are back in the classic case of variational preferences axiomatized by Maccheroni, Marinacci, and Rustichini (2006).

In view of Remark 2, for the entropy based set of extensions, we may rely on the robust control form of Hansen and Sargent (2001) via $c^{min}(\cdot\|\mathbb{P}) = \mathtt{Ent}(\cdot, \mathbb{P})$.

In order to follow the proof strategy of the present paper a modification of the HEP, say $\hat{W}_i$, needs to satisfy

$$\hat{W}_i(s_i, s_{-i}) = \mathbb{E}^{\overline{\underline{P}}^s}\left[u_i(s_i, s_{-i})\right] = \mathbb{E}^{\mathbb{P}}\left[\rho^s \cdot u_i(s_i, s_{-i})\right]$$

for any strategy profile $s = (s_i, s_{-i})$ and some Radon–Nykodym density ρ^s that has sufficient continuity on s, see Lemma 1 below. Hence, the above generalization of HEP, the case of $c^{min}(\cdot\|\mathbb{P}), c^{max}(\cdot\|\mathbb{P})$ being both linear penalties, again allows to prove existence of a corresponding equilibrium concept.[16] □

We conclude with a final observation that relates to the kind of information that can be managed within our model, and may possibly go towards distinguishing our non-Bayesian perspective from the Bayesian one.

Remark 4: The rules of the non-Bayesian game exist in the model as elements in the ensemble of aggregate information, the join $\mathcal{F}$ of the individual information $\mathcal{F}_i$, and as far as any rule outside the model is concerned, it has zero probability, so to speak. As such, these rules are surely constituted by events outside of $\mathcal{F}$, and known as such by every player. However, if some particular player does not know in any concreteness some particular rule, but nevertheless "knows" in some sense to be a potential rule, it could be an event in $\mathcal{F} \setminus \mathcal{F}_i$. It formalizes an event of "deep incomplete information" whereby it is known to have a probability, but to which no particular number, including the number zero can be assigned. □

2.4 The Main Result

For the main result presented in this paper, we need the following portmanteau assumption:

Assumption 1 *(i) Players agree on the null sets determined by $\mathbb{P} : \mathcal{F} \to [0, 1]$, where $\mathcal{F} = \sigma(\mathcal{F}_1, \ldots, \mathcal{F}_I)$, $\mathbb{P} = \prod_{i\in\mathbb{I}} \lambda_i$ and each λ_i is an atomless probability measure on $(T_i, \mathcal{T}_i)$.*

(ii) For each player i, the utility index $u_i(\cdot, a)$ is $\mathbb{P}$–square integrable and continuous for any $a = (a_1, \ldots, a_I) \in A$.

(iii) For each player i, u_i depends only on the i-th component of $\omega = (t_1, \ldots, t_I)$.

[16] This would be a variation of Theorem 1, and for whose existence we would be following Steps 1a and 1b in the proof of Theorem 1 below.

(iv) For each player i, there is a $\sigma_i \in L^2 = L^2(\Omega, \mathcal{F}, \mathbb{P})$ such that only those extensions $\tilde{P}$ are considered that satisfy $\frac{\mathrm{d}\tilde{P}}{\mathrm{d}\mathbb{P}} \leq \sigma_i$.

In relation to earlier work on Bayesian games with private information, the novel element is the additional hypothesis required on the set of all extensions. Towards this end, Assumption 1(iv) restricts the set of possible extensions allowed to a player to cope with his or her uncertainty. As already mentioned in Remark 2, with a given reference measure, extensions with a strong deviation from $\mathbb{P}$ are forbidden to the players. In the same vein, σ_i then determines the acceptable deviation from $\mathbb{P}$. As such, the assumption solely serves to guarantee weak compactness of $\mathcal{P}(\mathcal{F}_i)$ as specified in (3) below. An alternative to Assumption 1(iv) is to assume directly an upper bound for each player $i \in \mathbb{I}$, say $var_i \in \mathbb{R}_+$. This guarantees that for any density we have $\frac{\mathrm{d}\tilde{P}}{\mathrm{d}\mathbb{P}} \leq var_i$ for some extension $\tilde{P}$, and thereby to rely on Alaoglu's Theorem to ensure the weak compactness of each $\mathcal{P}(\mathcal{F}_i)$.

The following example illustrates Assumption 1(iv) in a standard setup.

Example 3 *Let $T_i = \mathbb{R}$ and $\lambda_i = \mathrm{d}x$ denote the Lebesgue integral on each coordinate of $\Omega = \mathbb{R}^{|I|}$. In anticipating Proposition 1.2 from Section 5.2 below, any extension $\tilde{P} = (\mu_1, \ldots, \mu_I)$ is of the product form. Assumption 1(iv) then translates into a condition that $\rho_{ij} = \frac{\mathrm{d}\mu_j}{\mathrm{d}\lambda_i} \leq \sigma$ for all $i, j \in \mathbb{I}$. For simplicity the bound σ is common for all players and constant. As such, Assumption 1(iv) boils down to the boundedness of all probability density functions ρ_{ij}.*

In keeping with our departure from Definition 4 and a fixed probability measure P_i on $(\Omega, \mathcal{F}_i)$ for some player i, the set of all such $\mathcal{F}_i$–consistent extensions to $(\Omega, \mathcal{F})$ is defined by

$$\mathcal{P}(\mathcal{F}_i) = \left\{ \tilde{P} \in \Delta(\Omega, \mathcal{F}) : \tilde{P} = P_i \text{ on } \mathcal{F}_i \text{ and } \tilde{P} \ll \mathbb{P} \text{ with } \frac{\mathrm{d}\tilde{P}}{\mathrm{d}\mathbb{P}} \leq \sigma_i \right\}. \tag{3}$$

By Assumption 1(i), the agreement of the null sets imply that any extension $\tilde{P}$ of P_i must be absolutely continuous with respect to $\mathbb{P}$ and is usually denoted by $\tilde{P} \ll \mathbb{P}$. As mentioned in Remark 1, the (reference) probability measure $\mathbb{P}$ determines all events that are possible, and thereby implicitly assumed to be common knowledge.

We now turn to a more detailed discussion of the implications of Assumption 1: it guarantees that HEP are well-defined.

Lemma 1 *Under Assumption 1(i), (ii) and (iv), the HEP $W_i(s_i, s_{-i})$ is finite for every feasible strategy $s = (s_i, s_{-i})$ and player i. In particular there is a $\underline{\overline{P}}^{\alpha_i, s} \in \mathcal{P}(\mathcal{F}_i)$ with*

$$W_i(s_i, s_{-i}) = \mathbb{E}^{\underline{\overline{P}}^{\alpha_i, s}} \left[u_i(\cdot, s_i, s_{-i}) \right],$$

depending on α_i and s.

As we shall see in the proof of Lemma 1, essential to the proof of Theorem 1, the assumed square integrability of the extension's Radon–Nikodym densities, in combination with Assumption 1(ii), then guarantees the finiteness of each Hurwicz-expected payoff W_i, defined in (2). But we now present the principal theorem of the paper.

Theorem 1 *Under Assumption 1, Hurwicz–Nash equilibria exist.*

We conclude this subsection with several remarks on the method and scope of the proof, leaving a more detailed account to a subsection.

Remark 5: A standard proof via a fixed–point of best–response correspondences with mixed strategies fails, as we have no concavity assumption on each utility index u_i. Moreover, Hurwicz expected utility, defined in (2), contains an optimistic part where a maximum of expectations is considered. This is a further source of difficulty that results in best–response correspondences with non–convex values. □

One has a clear sense of the relation between Bayesian and non–Bayesian games already from a direct comparison between the formal Definitions 1–3 and Definitions 4–6 and the case studies presented as Section 3. Remark 6 departs from Remark 1 and Definition 6 and simply encapsulates this difference.

Remark 6: Under a Bayesian-Nash equilibrium of a one-shot Bayesian game, the belief of any Bayesian player i, in combination with his or her private information, can be summarized by a well–understood probability function $p_i : T_i \to \Delta(T_{-i})$. The point is that the resulting consistency of the common prior assumption (recall the discussion in Section 1 above) requires the existence of a $\mathbb{P} \in \Delta(T)$ such that each p_i is induced via Bayes rule from $\mathbb{P}$. □

We conclude this subsection with a preliminary remark on the impossibility of Bayesian updating in a non-Bayesian framework, but we shall return to this issue in Section 4 below.

Remark 7: As continually emphasized in the exposition, players in the non–Bayesian game cannot know the probability of those events that are not part of (contained in) their private information, and as such, private information can no longer be used to update beliefs. In particular, within a non-Bayesian world, the role of a common prior assumption is rendered totally irrelevant. The only assumption that comes close in the (present) set-up is a requirement that all players agree on what is possible; which is to say, are aware of, and agree with, the null-sets of $\mathbb{P}$; see Assumption 1(i) and Remarks 1 and 2. At this point it is clear that the non-Bayesian approach operates "one stage" before the modeling stage of fixing a (common) prior, and can then be considered as an alternative, and perhaps rather complementary, way of modeling information asymmetries. □

3 Two Case Studies

This section presents two examples that directly compare Bayesian-Nash equilibria and Hurwicz–Nash equilibria. Both examples are structured in the same way.

In both case studies we start with the basic primitives shared in the Bayesian and non–Bayesian game let each example begin. Bayesian–Nash equilibria (Definition 3) and Hurwicz–Nash equilibria (Definition 6) are computed first in a simple two–type example (Section 3.1) and then in a more involved Battle of Sexes case study. In the latter, the continuous type space is the unit interval equipped with the (atomless) Lebesgue measure. At the end of each case we consider some further comparative statics for the Hurwicz-Nash equilibrium. In all games of this section there are only $I = 2$ players.

3.1 A Two–Type Example

In the first example a simple type space allows to model information as partitions. Only one player has incomplete information.

3.1.1 Basic Primitives in the Non-Bayesian Game

Let there be two players $i = 1, 2$ and both action sets $A_1 = \{I, D\}$ and $A_2 = \{E, F\}$ contain only two elements. Suppose that player 1 can only be of type θ, which is known by player 2. Player 2 can be of two different types, say $T_2 = \{\theta_a, \theta_b\}$. We have $\Omega = \{\theta\} \times \{\theta_a, \theta_b\}$.

Let payoffs be given by Table 1. The private information of player 2 is complete $\mathcal{F}_2 =$

for $t_2 = \theta_a$

	I	D
E	3,3	3,1
F	3,6	2,3

for $t_2 = \theta_b$

	I	D
E	3,3	3,6
F	3,1	4,3

Table 1: Payoff structure for the respective types

$\{\emptyset, \{\theta\}\} \times \{\emptyset, \{\theta_a\}, \{\theta_b\}, \{\theta_a, \theta_b\}\}$, while player 1 has no information about player 2, that is $\mathcal{F}_1 = \{\emptyset, \{\theta\}\} \times \{\emptyset, \{\theta_a, \theta_b\}\}$. Player 2 is of type θ_b.

3.1.2 Bayesian–Nash Equilibria

With the above primitive it remains to fix a prior to describe the Bayesian game with incomplete information, see Definition 2: fix the prior P on $T_1 \times T_2$, say $P(t_2 = \theta_a) = \frac{2}{3}$, that is, the probability that player 2 is of type θ_a. In view of Definition 2, conditioning in the present example yields deterministic beliefs. Specifically the conditional beliefs of player 1 about player 2, $P(\cdot|t_2 = \theta_c) = \delta_{\theta_c}$, $c = a, b$, are Dirac measures. Hence, payoffs are summarized in Table 1 and relates to classical Nash equilibria, on the respective event $\{t_2 = \theta_a\}$ and $\{t_2 = \theta_b\}$. From Table 1 we see that pure strategy equilibria are given by $s_2^{\mathsf{B}*} = E$ and $(s_1^{\mathsf{B}*}(\theta_a), s_1^{\mathsf{B}*}(\theta_b)) = (I, D)$.

For completeness, we also consider here unconditional expected payoffs. With the payoff data from Table 1, we then get any expected payoff for player 1:

$$\mathbb{E}^P[u_1(s_1, s_2)] = \frac{2}{3} u_1\Big(s_1(\theta_a), s_2(\theta)\Big) + \frac{1}{3} u_1\Big(s_1(\theta_b), s_2(\theta)\Big).$$

From this we get a (unconditional) Bayesian–Nash equilibrium given by $s_2^* = E$ and $(s_1^*(\theta_a), s_1^*(\theta_b)) = (I, D)$. As listed in Table 2, this can be seen by evaluating for both players all possible expected payoffs. The first row specifies the expected payoffs of Player 1, if player 2 chooses $E \in A_2$, given that Player 1 plays one of the four strategies being induced by her information set.[17]

	I,I	*I,D*	*D,I*	*D,D*
E	3	3	3	3
F	3	10/3	7/3	8/3

Table 2: Expected payoffs of player 1 in the unconditioned game

3.1.3 Hurwicz–Nash Equilibria

When moving to a non–Bayesian game with probabilistically incomplete information (Definition 4), the belief of player 1 is only defined on $\mathcal{F}_1$. For the present example this means that any probability on $T_2 = \{\theta_a, \theta_b\}$ is possible, hence $\mathcal{P}(\mathcal{F}_1) = \Delta(T_2)$. In view of Definition 5, it remains to specify the Hurwicz parameter α_1 for player 1. To keep the specification close to the unconditional Bayesian game from Subsection 3.1.2, set for the moment $\alpha_1 = \frac{2}{3}$.

Clearly, since $\mathcal{P}(\mathcal{F}_2)$ is single valued, the choice of α_2 is irrelevant. The ambiguous beliefs of player 1 about player 2 results in the HEP from (2), that induces the game with probabilistically incomplete information (Definition 5) defined by

$$\begin{aligned} W_1(s_1, s_2) = \quad & \frac{2}{3} \min_{p \in [0,1]} \Big\{ p \cdot u_1\Big(s_1(\theta), s_2(\theta_a)\Big) + (1-p) \cdot u_1\Big(s_1(\theta), s_2(\theta_b)\Big) \Big\} \\ + \quad & \frac{1}{3} \max_{p \in [0,1]} \Big\{ p \cdot u_1\Big(s_1(\theta), s_2(\theta_a)\Big) + (1-p) \cdot u_1\Big(s_1(\theta), s_2(\theta_b)\Big) \Big\}. \end{aligned}$$

Computing all possible Hurwicz–expected payoffs of player 1 with data from Table 1 yields Table 3:

[17] Specifically, "I, D" means that player 1 plays I if player 2 is of type $t_2 = \theta_a$ and plays D if $t_2 = \theta_b$.

	I,I	*I,D*	*D,I*	*D,D*
E	3	3	3	3
F	3	10/3	7/3	8/3

Table 3: Hurwicz–expected payoff of player 1 in the non–Bayesian game

Consequently, we get the following Hurwicz–Nash equilibrium $s_2^{\mathtt{H}*} = E$ and $(s_1^{\mathtt{H}*}(\theta_a), s_1^{\mathtt{H}*}(\theta_b)) = (I, D)$. The outcome within this equilibrium notion coincides with the Bayesian–Nash equilibrium $(s_1^{\mathtt{B}*}, s_2^{\mathtt{B}*})$ and the unconditional Bayesian–Nash equilibrium (s_1^*, s_2^*) as specified above. However, as we see in the following, any such coincidences are *not robust*, if we parametrize the game with the index for ambiguity aversion of player 1, that is $\alpha_1 \in [0,1]$. As we see in the sequel, a change of optimism may lead to a different Hurwicz–Nash equilibrium strategy.

3.1.4 Comparative Statics of Hurwicz–Nash Equilibria

Departing from the payoff structure of Table 3, we may find conditions on α_1 which yields a coincidence of the equilibrium strategies for the two equilibrium concepts, that is, $s^{\mathtt{B}*} = s^{\mathtt{H}*}$. However such a coincidence of equilibrium outcomes relies on the choice of the a priori given parameter for ambiguity aversion of player 1, α_1, determining the Hurwicz-expected payoffs of player 1. This payoff dependency is summarized in Table 4, by generalizing Table 3.

	I,I	*I,D*	*D,I*	*D,D*
E	3	3	3	3
F	3	4-α_1	3-α_1	4-2α_1

Table 4: Expected payoffs of player 1 in the non–Bayesian game, with a dependency of the ambiguity index of player 1

From Table 4 we can now infer that if $\alpha_1 \neq 1$, then both Bayesian and non–Bayesian equilibrium strategies coincide with $s_2^{\mathtt{B}*} = E = s_2^{\mathtt{H}*}$ and $s_1^{\mathtt{H}*} = (I, D) = s_1^{\mathtt{B}*}$. In the case of extreme pessimis, $\alpha_1 = 1$, the Hurwicz-Nash equilibrium strategy of player 1 changes and is given by $s_1^{\mathtt{H}*} = (I, \{I, D\})$.

3.2 Continuous–Type Space Example - Battle of Sexes

The example in Subsection 3.1 differs from the setting of our main result, Theorem 1, in one crucial aspect: the measure on the type space was (implicitly) atomic. In the example of this subsection, we consider an atomless (continuous) type space. However, also in the present Battle of Sexes example only two actions for both players are possible: $\{\texttt{Opera}, \texttt{Fight}\}$.

3.2.1 Basic Primitives in the Non-Bayesian Game

Consider again two players $i = 1, 2$ with same action sets $A_1 = A_2 = \{\texttt{Opera}, \texttt{Fight}\} =: \{\texttt{O}, \texttt{F}\}$. Player i can be of some type $t_i \in T_i := [0, 1] = [0, 1]_i$ with $\Omega = [0, 1] \times [0, 1]$ and $[0, 1]$–Lebesgue product measure $\mathbb{P} = \lambda_{[0,1]} \otimes \lambda_{[0,1]}$. Here, $[0, 1]_i$ denotes the respective coordinate. The private information of player i is $\mathcal{F}_i = \mathcal{B}([0, 1]_i)$. For simplicity we ignore any enlargements via λ-null sets.

Player 2 \ Player 1	`Opera`	`Fight`
`Opera`	$\frac{1}{2} + t_1$,1	0,0
`Fight`	0,0	1, $\frac{1}{2} + t_2$

Table 5: Payoff structure for the battle of sexes with $t_1, t_2 \in [0, 1]$.

Table 5 describes the payoffs for the Bayesian and non-Bayesian game. In the following, we discuss a specific case of the basic primitives to get a further perspective about the differences between Bayesian and non-Bayesian games and the respective equilibrium outcomes.

3.2.2 Bayesian–Nash Equilibria

The common prior is given by the standard product Lebesgue measure $\mathbb{P} = \lambda_{[0,1]_1} \otimes \lambda_{[0,1]_2}$. After updating we now only identify the push–forward measure $\mathbb{P}\circ(s_i^{\texttt{B}*})^{-1}$ on A_i via a $\texttt{p}_i \in [0, 1]$ with respect to equilibrium strategy $s_i^{\texttt{B}*} : [0, 1]_i \to A_i$ in the present Bayesian game form. Each $(\texttt{p}_i, 1 - \texttt{p}_i) \in \Delta_{|A_i|}$ then represents a probability distribution on the action set $\{\texttt{O}, \texttt{F}\}$.

To obtain the pair of equilibrium strategies (s_1, s_2) without full measure on one action a_i, note that player 1 (player 2) will *not* play action $\texttt{O}$ (action $\texttt{F}$) as a low type $t_1 \in [0, t_1^*]$ ($t_2 \in [0, t_2^*]$) up to some notrivial threshold $t_1^* \in (0, 1)$ ($t_2^* \in (0, 1)$). Moreover, by the linearity of the payoff structure (as a function of each type), at most one threshold for player i can exist. Consequently, we have $\texttt{p}_1 = \mathbb{P}(s_1^{\texttt{B}*} = \texttt{O}) = \lambda([t_1^*, 1]) = 1 - t_1^*$ and $\texttt{p}_2 = \mathbb{P}(s_2^{\texttt{B}*} = \texttt{O}) = \lambda([0, t_2^*]) = t_2^*$.

Given type $t_1 \in [0, 1]_1$ of player 1 and a strategy $s_2 : [0, 1]_2 \to \{\texttt{O}, \texttt{F}\}$ of player 2, the payoffs of player 1 are then:

$$\begin{aligned}
U_1(\mathtt{F}, s_2|t_1) &= \mathbb{E}^{\mathbb{P}}[u_1(\mathtt{F}, s_2)|t_1] = E^{\mathbb{P}}(\cdot|t_1) \circ (s_2)^{-1}[u_1(\mathtt{F}, \cdot)] = \mathbb{P}(s_2^{\mathtt{B}*} = \mathtt{O})0 + \big(1 - \mathbb{P}(s_2^{\mathtt{B}*} = \mathtt{O})\big)1 \\
&= \mathtt{p}_2 \cdot 0 + (1 - \mathtt{p}_2) \cdot 1 = 1 - \mathtt{p}_2, \\
U_1(\mathtt{O}, s_2|t_1) &= \mathtt{p}_2\Big(\frac{1}{2} + t_1\Big) + (1 - \mathtt{p}_2) \cdot 0 = \mathtt{p}_2\Big(\frac{1}{2} + t_1\Big).
\end{aligned}$$

For player 1 with type t_1, the best response to play $\mathtt{O}$ is characterized by $\mathtt{p}_2 \geq \frac{2}{2t_1+3}$.

Similarly, for a given type $t_2 \in [0,1]_2$ and strategy $s_1 : [0,1]_1 \to \{\mathtt{O}, \mathtt{F}\}$ of player 1, we get for player 2:

$$\begin{aligned}
U_2(s_1, \mathtt{F}|t_2) &= \mathtt{p}_1 \cdot 0 + (1 - \mathtt{p}_1)\left(\frac{1}{2} + t_2\right) = (1 - \mathtt{p}_1)\left(\frac{1}{2} + t_2\right), \\
U_2(s_1, \mathtt{O}|t_2) &= \mathtt{p}_1 \cdot 1 + (1 - \mathtt{p}_1)0 = \mathtt{p}_1.
\end{aligned}$$

For player 2 with type t_2, the best response to play $\mathtt{O}$ is characterized by $\mathtt{p}_1 \geq \frac{2t_2+1}{2t_2+3}$.

With $\mathtt{p}_1 = 1 - t_1^*$ and $\mathtt{p}_2 = t_2^*$, we can now solve the pair of equilibrium conditions. This gives the equilibrium threshold values $t_1^* = \frac{1}{2}$ and $t_2^* = \frac{1}{2}$ and a pair of equilibrium strategies

$$s_1(t_1) := \begin{cases} \mathtt{F} & t_1 \in [0, \frac{1}{2}) \\ \{\mathtt{O}, \mathtt{F}\} & t_1 = \frac{1}{2} \\ \mathtt{O} & t_1 \in (\frac{1}{2}, 1] \end{cases} \qquad s_2(t_2) := \begin{cases} \mathtt{O} & t_2 \in [0, \frac{1}{2}) \\ \{\mathtt{O}, \mathtt{F}\} & t_2 = \frac{1}{2} \\ \mathtt{F} & t_2 \in (\frac{1}{2}, 1] \end{cases}$$

Incorporating now also equilibrium strategies s_i^* with full measure on one action a_i, that is $\mathbb{P}(s_i^* = a_i) \in \{0, 1\}$ for each a_i and i, gives the three pairs of Bayesian-Nash equilibrium strategies $(s_1^{\mathtt{H}*}, s_2^{\mathtt{H}*}) = \{(\mathtt{F}, \mathtt{F}), (\mathtt{O}, \mathtt{O}), (s_1, s_2)\}$.

3.2.3 Hurwicz–Nash Equilibria

In the non–Bayesian game, the basic primitive information structure $(\mathcal{F}_1, \mathcal{F}_2)$ gives the following beliefs $P_i = \lambda_{[0,1]_i}$ which are only defined on the sub σ-algebra $\mathcal{F}_i$, $i = 1, 2$. The set of extensions (on the opponents type space) is then given by

$$\mathcal{P}(\mathcal{F}_i) = \left\{\lambda_{[0,1]_i} \otimes \mu : \mu \in \Delta([0,1]_{-i}) \text{ and } \frac{\mathrm{d}\mu}{\mathrm{d}\lambda_{[0,1]_{-i}}} \in L^2([0,1])\right\} \approx \lambda_{[0,1]_i} \otimes \Delta([0,1]_{-i}).$$

For simplicity we allow here also extensions that are not absolute continuous with respect to λ. This is in line with our motivation to understand in a concise way the difference between Bayesian–Nash and Hurwicz–Nash equilibria. However, it is worth stating that by considering densities of the form $\rho^p(x) = c^p x^{-p}, p \geq 0$, we have $\rho^p \in L^2([0,1], \mathcal{B}([0,1]), \lambda)$ for any p and can therefor approximate any Dirac measure sufficiently close by a λ-absolutely continuous measure μ^p given by $\frac{\mathrm{d}\mu^p}{\mathrm{d}\lambda} = \rho^p$. The resulting equilibrium strategy with Dirac measures δ_t at $t = 0$ is then sufficiently close to such L^2-approximations.

As we consider a battle of sexes, let us fix $\alpha_1 = \alpha_2 = 0$ and consider optimistic (that is, ambiguity loving) players. Moreover, the complete ignorance of player 1 about player 2 results in a rather large set of possible beliefs. We fix these behavioral primitives. The HEP $W_i(s) = \max_P \mathbb{E}^P[u_i(s)]$ of player i is a linear problem. The optimal solution is on the boundary of the convex set $\mathcal{P}(\mathcal{F}_i)$ and contains at least one element in the set of extremal point $ext(\Delta([0,1]_{-i})) = \{\delta_t : t \in [0,1]\}$.

To ease computation, we focus only on those strategies where for each player i, there is a positive mass of types playing each action in A_i, i.e. $\lambda_i(s_i(T_i) = a_i) > 0$ for each $a_i \in A_i$, $i = 1, 2$. Under this restriction, we next compute payoffs W_i for given t_i.

With this we derive for player 1:

$$\begin{aligned} W_1(t_1, \mathtt{F}, s_2) &= \max_{P \in \mathcal{P}(\mathcal{F}_1)} \mathbb{E}^P[u_1(t_1, \mathtt{F}, s_2)] = \max_{\delta_t : t \in [0,1]_2} \delta_t 1 + 0 \cdot 0 = 1, \\ W_1(t_1, \mathtt{O}, s_2) &= \max_{P \in \mathcal{P}(\mathcal{F}_1)} \mathbb{E}^P[u_1(t_1, \mathtt{O}, s_2)] = \max_{\delta_t : t \in [0,1]_2} \delta_t \left(\frac{1}{2} + t_1\right) + 0 \cdot 0 = \frac{1}{2} + t_1. \end{aligned}$$

For player 2 we get in a very similar way:

$$\begin{aligned} W_2(t_2, s_1, \mathtt{O}) &= \max_{\delta_t : t \in [0,1]_1} \delta_t 1 + 0 \cdot 0 = 1, \\ W_2(t_2, s_1, \mathtt{F}) &= \frac{1}{2} + t_2. \end{aligned}$$

This together yields the Hurwicz–Nash equilibrium strategies

$$s_1^{\mathtt{H}^*}(t_1) = \begin{cases} \mathtt{F} & t_1 \in [0, \frac{1}{2}) \\ \{\mathtt{O}, \mathtt{F}\} & t_1 = \frac{1}{2} \\ \mathtt{O} & t_1 \in (\frac{1}{2}, 1] \end{cases} \qquad s_2^{\mathtt{H}^*}(t_2) = \begin{cases} \mathtt{O} & t_2 \in [0, \frac{1}{2}) \\ \{\mathtt{O}, \mathtt{F}\} & t_2 = \frac{1}{2} \\ \mathtt{F} & t_2 \in (\frac{1}{2}, 1] \end{cases}$$

As in Bayesian-Nash equilibrium, at $t_i = \frac{1}{2}$, the optimal strategies are multi valued $\{\mathtt{F}, \mathtt{O}\}$. Furthermore, strategies with full measure on one action, yields the following equilibrium strategies: $s_1^{\mathtt{H}^*} \equiv \mathtt{F} \equiv s_2^{\mathtt{H}^*}$ and $s_1^{\mathtt{H}^*} \equiv \mathtt{O} \equiv s_2^{\mathtt{H}^*}$.

From this, we directly see that the Bayesian-Nash and the Hurwicz-Nash equilibrium coincide. However, as we see in the following subsection, this coincidence is far away from being robust in the players primitives. Once we change the parameter α_1 different Hurwicz-Nash equilibrium strategies emerge.

3.2.4 Comparative Statics of Hurwicz–Nash Equilibria

So far we considered $\alpha_1 = \alpha_2 = 0$ for the HEU (both players are optimists) in the present non-Bayesian game. As in the comparative statics of the first example, a variation of the primitive may yield different equilibria. In particular, for certain parameters the Bayesian–Nash and Hurwicz–Nash equilibrium strategies and payoffs can match up. As shown in the following two variations of the example, such a coincidence is far away from being a generic property.

A Pessimistic Players Set $\alpha_1 = \alpha_2 = 1$. We now have the HEP $W_i(t,s) = \min_P \mathbb{E}^P[u_i(t,s)]$ with rather ambiguity averting players. In that case, we get a Hurwicz–Nash equilibrium with zero payoffs. Specifically, we can write payoffs via the effective or minimizing prior $\overline{\underline{P}}^{0,s} = \underline{P}^s$, see Lemma 1 in Section 2.3. This yields the equilibrium payoff to be $\mathbb{E}^{\underline{P}^s}[u_i(t_i, s_i^{\mathtt{H}*}, s_{-i}^{\mathtt{H}*})] = 0$. In particular, this is very much in line with the trade-reducing character of ambiguity aversion.

B Optimistic Player 2 and Varying Ambiguity Parameter for Player 1 Finally, we modify the present example to grasp a better understanding how the preference for ambiguity in the HEP affects the equilibrium strategy.[18] Let us assume that for some reason player 1 enjoys a payoff of $\frac{1}{3}$ when playing `Opera` and player 2 chooses `Fight`.[19] For the equilibrium with non-constant strategies, player 1 now plays `F` if and only if $t_1 < t^*_{\alpha_1} := \frac{3-\alpha_1}{6(1-\alpha_1)}$.

For player 2, the equilibrium strategies remain unaffected. The equilibrium strategy of player 1, denoted by $s_1^{\alpha_1,\mathtt{H}^*}$, now depends on α_1 and is given by:

$$s_1^{\alpha_1,\mathtt{H}^*}(t_1) = \begin{cases} \mathtt{F} & t_1 \in [0, t^*_{\alpha_1}) \\ \{\mathtt{O},\mathtt{F}\} & t_1 = t^*_{\alpha_1} \\ \mathtt{O} & t_1 \in (t^*_{\alpha_1}, 1]. \end{cases}$$

As a comparative-static exercise, if $\alpha_1 \geq \frac{3}{5}$ player will always play 1 `Fight` in equilibrium. We illustrate the modification in payoffs and the resulting equilibrium strategy of player as a function of α_1 in Figure 1.

4 Alternative Perspectives on the Result

However the basic equilibrium concept investigated in this work is motivated, be it by the dissonance between private information and public beliefs, or by an attempt at one concrete and viable notion of group-rationality that can face upto Hurwicz' man-made uncertainty, the point is that we end up in a situation in which a player does not have a single determinate value attached to a particular profile of types, or sets of types, of her opponents in the game. This

[18] Without this change of payoffs, a distortion in $(\alpha_1, \alpha_2) \in [0,1]^2 \setminus \{1,1\}$ would leave the strategy unaffected.
[19] For arbitrary $\alpha_1 \in [0,1]$, this results in: $W_1(t_1, \mathtt{F}, s_2) = (1-\alpha) + \frac{\alpha}{3}$ and $W_1(t_1, \mathtt{O}, s_2) = (1-\alpha)(\frac{1}{2} + t_1)$.

Player 2 \ Player 1	Opera	Fight
Opera	$\frac{1}{2}+t_1$,1	0,0
Fight	$\frac{1}{3}$,0	1, $\frac{1}{2}+t_2$

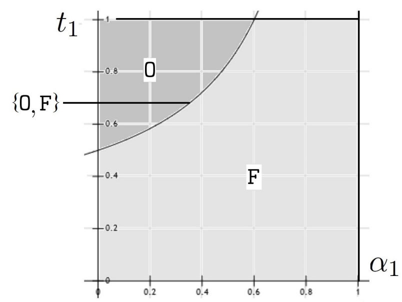

Figure 1: A modification of Table 5 for the Payoff of agent 1. The graphic illustrates player 1's equilibrium strategy as a function of preference for ambiguity $\alpha_1 \in [0,1]$ and type $t_1 \in [0,1]$.

probabilistic indeterminacy connects to an extensive philosophical and statistical literature that addresses itself to the difference between information and ignorance as it grapples with uncertainty. To be sure, any half-way adequate discussion of this work will take us well beyond the scope of this single-theorem paper, but a succinct pointing-out of the ways in which the theorem connects to a larger context of work may be worthwhile in gauging its importance. The question is whether Hurwicz-Nash equilibria of a non-Bayesian game have any relevance to the literature on imprecise and fuzzy probabilities? The following listing of alternative perspectives on our result answers this question and motivates this section.[20]

The illustrative case studies, presented in Section 3 and bringing out the distinctive autonomy of the Hurwicz-Nash equilibrium concept that we investigate, hinge on the fact that any $P \in \mathcal{P}(\mathcal{F}_i)$ gives an interval of values for any $A \in \mathcal{F}$; see (1) above. Seidenfeld and Wasserman (1993) ascribe an interval of probabilities to a particular event as arising from a variety of possible considerations: (i) formalization of group-rationality as discussed already above, (ii) the need for a rigorous mathematical framework for studying sensitivity and robustness in classical and Bayesian inference, (iii) a consequence of the weakening of the (Kolmogorov) axioms of classical probability theory, (iv) situations that necessarily have to make do with incomplete or partial elicitation,[21] (v) physical phenomena that inherently involve upper and lower probabilities.[22] The trajectory of the evolution of a sophisticated and nuanced litera-

[20]The authors are grateful to an anonymous referee for inspiring this section through a "throwaway remark" to the work of Peter Walley.

[21]In addition to Footnote 12, see the discussion of "direct inference" in Levi (1977) and Kyberg Jr. (1977). In referring to their exchanges regarding the issue, the latter writes, "we agree that the point at issue – at base an epistemological point – is of profound and critical importance for the philosophy of science and for rational decision theory." In this connection, also see Vicig and Sidenfeld (2012), Pellesoni and Vicig (2016) and their references.

[22]For a comprehensive and balanced discussion of these issues that does not short-change his own approach,

ture that they chart can be summarized in one nutshell sentence: from precise probabilities to imprecise probabilities to previsions to possibilities to the all-encompassing framework of fuzzy set theory and logic.[23] However, one important distinction from this literature that is implicit in our theorem must be carefully noted. The formulation of the problem that we study allows, indeed asks for, a setting in which the probabilities (beliefs) on *part* of the universal state space are precisely known, and it is only its complement in the state space that a player is ignorant of, or gives imprecise and ambiguous probabilities to. This distinction has implications.

The first implication is that the observation connects to the axiomatic underpinning of Hurwicz preferences provided by Gul and Pesendorfer (2014, 2015). Their theory also revolves on a division of the domain of the given preference relation into a part which satisfies Savage's axioms, in particular his sure-thing-principle, and a part that does not. The achievement of the theory is then to extend the representation on part of the domain to the entire domain, and it is this extension that brings in ambiguity, and leads to the Hurwicz decision-theoretic functional representation. It is thus no surprise that the authors are led directly, and naturally, to the Dempster-Shafer theory of evidence, a theory that also falls squarely within one of the frameworks listed in the paragraph above. In this connection, we note that Zhang, Luo, Ma, and Leung (2014) develop a theory of ambiguous Bayesian games under the rubric of the Dempster-Shafer theory executed under a finiteness (atomicity) assumptions on the information space and therefore necessarily coupled with the quasi-concave assumption on the payoffs. This paper then is focused on providing an extension of the original Nash theorem that simply incorporates the decision-theoretic ambiguity literature stemming from Marinacci (2000), and as such not directly relevant to the non-Bayesian result presented here.[24]

The second implication of the distinction arising from a bifurcation of the state space into a determinate and an indeterminate part is that it leads one to regard the equilibrium concept presented here as a temporary one in the sense of a rich Hicksian tradition in economic theory. As emphasized in the introduction, and now analytically transparent in the formal presentation, each player extends his or her subjectively-determinative, and hence objective, beliefs to those of his or her opponent. This inferred extension, by necessity, involves a set of beliefs, and the optimal actions are then taken based on this inferred set, with an attitude relating to optimism and pessimism measured by the Hurwicz signature parameter α. The

see Walley (1991) and the follow-up regarding the notion of *independence* by Couso, Moral, and Walley (1999) and Couso and Moral (2010). We refer to this direction again in the last two sentences of Section 5.1 below.

[23] We are now fortunate to have a comprehensive discussions in Bělohlávek, Dauben, and Klir (2017) (BDK) and in Walley (1991); see Section 3.7 concerning "possibility theory" and Section 4.5,2 on "possibilistic logic" in BDK on Shackle's work, and references to the Dempster-Shafer theory in Yager, Fedrizzi, and Kacprzyk (1994).

[24] This issue then takes us back to the work of Azrieli and Teper (2011) already discussed in the introduction.

equilibrium notion then emerges out of individual actions taken this way, and its fixed point is then sustained by each player's equilibrium inferences of the other players' information and beliefs. The question then is whether each agent can use these equilibrium inferences to update and refine his own information and beliefs, in the specific sense that is given to such an updating in the Walrasian theory of rational expectations.[25] Note that even though we are talking of updating of beliefs on the equilibrium path unlike that considered it Remark 6, this direction does *not* go against the non-Bayesian thrust of the work. The equilibrium does not involve individual players making inferences based on additional informational data being churned up at different stages of the game, but rather their updating with respect to a temporary equilibrium and affecting its transformation into a permanent one. This again, is an opening for the future, but one that may well prove productive for a fruitful synthesis of a Bayesian and non-Bayesian approaches.[26]

We now conclude this section by a look at our theorem from two viewpoints within non-cooperative game theory itself, at least as it is conventionally elaborated in the economic theory literature. This can be reduced to two questions in the formulation of equilibrium: (i) why the specific functional form of Hurwicz preferences as opposed to any other? and (ii) why the omission of constructed hierarchies of (ambiguous or indeterminate) beliefs as is by now a standard fare of epistemic game theory? Both questions can be succinctly answered, and both involve looking down the road to an outgrowth of the ideas presented here in their simplest setting. As regards (i), the maximin framework of Gilboa-Schmeidler is a special case of our approach, and we could surely have cast the result in terms of the smooth ambiguity model or one based on variational preferences, or indeed any other candidate of the many that are now available, but unlike the Gul-Pesendorfer (for us, tailor-made) framework, the others are *ad hoc* in the sense that they are not induced from a setting of incomplete information. As regards (ii), an equilibrium concept taking epistemic considerations into account would rely on the construction of hierarchies of beliefs based on sets of probabilities rather than on a single determinate one, and it would again fall under the rubric of a synthesis between Bayesian and non-Bayesian ideas. We sight the early paper of Heifetz and Mongin (2001) in this context, and note that such constructions are now available in the literature; see Friedenberg and Meier (2015) and Ahn (2007). This is surely yet another worthwhile investigation for the future.[27]

[25]See Khan's (2008) entry on *perfect foresight* and his references to the Jordan-Radner introduction to a JET symposium, and quotations from Chapter 12 from Keynes' *magnum opus* on employment, interest and money.

[26]The concept of *dilations* would fit in here with this equilibrium updating; see Seidenfeld and Wasserman (1993) and their follow-up by Pedersen and Wheeler (2014). The authors thank Osama Khan for emphasizing the relevance of "dilations" for our context.

[27]The authors thank Amanda Friedenberg, Peter Hammond and Alejandro Manelli for drawing their attention to this line of investigation.

5 A Proof of the Result

We preface the presentation of the proof by framing the general approach within which it is set, and then by bringing into salience the basic auxiliary results that it rests on. This constitutes Sections 5.1 and 5.2 below. Section 5.3 is a detailed elaboration of the argument that may be skipped on a first reading.

5.1 A Sketch of the Proof

There are two distinct approaches in the literature on the existence of pure-strategy equilibria that one can appeal to for the proof of Theorem 1. The first proceeds by showing the existence of behavioral-strategy equilibria, followed by an appeal to a the generalization of the Lyapunov theorem by Dvoretsky-Wald-Wolfowitz (DWW) to "purify" such equilibria. This was the original approach of Radner and Rosenthal (1982), and was subsequently followed by Milgrom and Weber (1985). Since these papers had already limited themselves to Bayesian private-information games with finite action sets in the case of pure-strategy equilibria, the existence argument was standard and the novelty hinged on the purification argument. A clear synthetic exposition is furnished in Khan, Rath, and Sun (2006), and it emphasized that in terms of mathematical technique, all that was needed was available as early as 1951. The other approach is a direct proof of existence, and in so far as the antecedent literature is concerned, this took some time in coming. In particular, it had to await results on the induced distribution of a set-valued random variable, results that relied on the Ballobás-Varopoulos extension of the marriage lemma (see Section 3 in Khan and Sun (1995) for the theory), and one that finds fuller expression in the context of finite action sets in Fu, Sun, Yannelis, and Zhang (2007).

5.1.1 Beyond a Convexification argument

In the Hurwicz-Nash setting elaborated above, there is no available theorem on the existence of a behavioral-strategy equilibrium; and so, by necessity, the proof of Theorem 1 presented here proceeds with the direct approach. This involves an application of Kakutani fixed-point theorem for finite-dimensional Euclidean spaces,[28] but several steps have to be filled in for the execution of the argument. We have to devise a compact convex set, and a upper-hemi–continuous mapping convex-valued mapping from that set to itself, such that the fixed point of the said mapping furnishes a Hurwicz-Nash equilibrium of the game. As is by

[28]See, for example, Section 17.9 in Aliprantis and Border (2006). It should perhaps be noted that as a consequence of the standing hypothesis of finite action sets, no infinite-dimensional fixed point theorem is required.

now usual, both in Walrasian equilibrium theory, and in non-atomic games, this mapping is some sort of "composition" of the best-response correspondence of each individual player, a correspondence from the set of joint distributions of the strategy profiles of the individual players to her own. Given such a profile, each player chooses a function from his or her type to what he or she considers her optimal actions, which is to say, a correspondence from her type space to her action set, and the induced distribution of this correspondence furnishes the individual correspondence that is then composed with that of the other players. The point is that this correspondence from the product of the other players' distributions to a player's own needs to be upper hemi-continuous, compact- and convex-valued. That it is indeed so is presented as Proposition 2 below. The atomlessness postulate guarantees convexification and that of independence ensures that a player can successfully "integrate out" from her expected utility functional the responses of the other players. But again, it should be noted that with there being no linear structures assumed on the individual action sets, the question of quasi-concavity of the associated pay-off functionals cannot even arise, and that therefore this convexity is not totally a straightforward matter. One has to enter by the back-door, so to speak.

5.1.2 The present Approach

Thus the only point that remains is to ensure that the actions that each player takes as a consequence of the Hurwicz-expected payoffs are well-defined and upper hemi-continuous. The latter follows from Berge's maximum theorem, and the former is presented as Lemma 1 below. As already emphasized both in the introduction and in Section 4, we do not have a single prior available to us, and have to work with an extended set of private beliefs on the universally-known sample space. It is here that there is an additional impediment to overcome. This concerns the assumption of *dispersedness,* which is to say, independence of information in the context of these extended set of priors, and it is this that is formalized as Assumption 1 above for the context in which we work. More specifically, since one is working with the extended priors, one has to ensure that they constitute a compact set, and it is this that is guaranteed by the existence of a Radon-Nikodym derivative of the extensions with respect to the product measure, as encapsulated in Assumption 1(iv). It is this that constitutes the novel supplementation of the standard argument has to be supplemented.[29]

We conclude this subsection with an observation concerning the extension of the result, and its proof, to action sets that are not finite. The antecedent literature in this case has proceeded by the division in two cases: (i) for countably-infinite action sets, the DWW theorem was

[29]The authors are grateful to an anonymous referee for emphasizing this point, and to Max Stinchcombe for showing us an example that suggests that the assumption of the existence of L^2-densities cannot be dispensed with.

extended by Khan and Sun (1995) and Khan and Rath (2009), and resulted in extensions that could be phrased in terms of arbitrary type spaces, (ii) for uncountable action sets, additional assumptions were required on the type spaces that were also found to be necessary; see Khan and Sun (1999), Khan and Zhang (2014) and their references. With these results at hand, the direct proof could be executed simply by the substitution of the Fan-Glicksberg fixed-point theorem for the Kakutani theorem; see Chapter 17 in Aliprantis and Border (2006). Indeed, one could move further in the direction of relaxing the independence constraint and non-interdependent types by appealing to the techniques detailed in the references in Footnote 2; also perhaps through the use of the extended Lebesgue as recently detailed in Khan and Zhang (2017), or the conceptual taxonomies of Couso, Moral, and Walley (1999). There is little doubt that the theorem we present and prove below can be extended to these more general contexts, and we leave this for future work.

5.2 The Auxiliary Results

The auxiliary results discussed in the subsection above are formally stated here: their proofs are relegated to Section 6.

Proposition 1 *Suppose Assumption 1 holds and fix a sub σ-algebra $\mathcal{G} \subset \mathcal{F}$.*

1. *The set $\mathcal{P}(\mathcal{G}) \subset \Delta(\Omega)$, defined in* (3), *is convex and $\sigma(L^2, L^2)$–compact.*[30]
2. *Any extension $P \in \mathcal{P}(\mathcal{G})$ is defined on $\otimes_{i\in\mathbb{I}}\mathcal{T}_i$ and is constituted by a product structure, i.e. $P = (\mu_1 \ldots, \mu_I)$, with $\mu_i \in \Delta(T_i)$ being atomless for each i.*

As constructed in the proof of Theorem 1 (see Step 1(c) below), the collection of all such $\overline{\underline{P}}^{\alpha_i,s}$, a convex combination minimizing and maximizing beliefs, ensured by Lemma 1, is denoted by $\overline{\underline{\mathcal{P}}}^{\alpha_i}(s)$. An element therein is called *effective prior* of player i at strategy s.

We now present the principal mathematical object and its properties below; see Theorems 7 to 9 in Khan and Sun (1995). This is applied in the crucial Step 3 of the proof of Theorem 1, and the reader should note that Assumption 1(i) enters in this application in a crucial way: without an atomless probability space such results are in general false.

Proposition 2 (Distribution of Correspondences) *Let A be a finite set, Y a metric space, $(T, \mathcal{T}, \lambda)$ an atomless probability space, and $\Xi : T \times Y \Rightarrow A$ be a correspondence. For each $y \in Y$, let $\Xi(\cdot, y) : T \Rightarrow A$ be $\mathcal{T}$–measurable. Define the correspondence from Y to $\Delta(A)$ by*

$$\mathscr{D}_{\Xi(\cdot,y)} = \left\{\lambda \circ \phi^{-1} : \phi \in \mathbb{M}^{sel}\Big(\Xi(\cdot, y)\Big)\right\} \quad \text{for all } y \in Y,$$

[30] As we shall see in the proof below, this notion of compactness relies heavily on the existence of the relevant Radon–Nikodym derivatives. Also see Footnote 29 above and the text it footnotes.

where $\mathtt{MB}^{sel}$ *denotes the collection of all measurable selections of a correspondence. Then,*

1. $\mathscr{D}_{\Xi(\cdot,y)}$ *is convex and compact valued;*

2. *if, in addition, the correspondence* $\Xi(t,\cdot)$ *is upper hemi-continuous on* Y *for each* $t \in T$*, then* $\mathscr{D}_{\Xi(\cdot,y)}$ *is upper hemi–continuous on* Y*.*

5.3 The Proof of Theorem 1

The proof is divided in three steps: the first step reformulates the expected payoffs of each player, the second prepares the fixed-point argument through the employment of the distribution of a correspondence, and the third executes it.

Let us introduce for each player an underlying (normed) space $L(T_i, \mathcal{F}_i; A_i)$ of strategies. It denotes the space of all $\mathcal{F}_i$–measurable mappings from T_i to A_i, equipped with the *sup*–norm, which in turn is denoted by $\|\cdot\|_{i,\infty}$. Set $L = \prod_i L(T_i, \mathcal{F}_i; A_i)$ and $\|\cdot\|_\infty = \sup_i \|\cdot\|_{i,\infty}$ on L.

Step 1 Reformulating of expected payoffs: We begin this step by rephrasing the payoffs of each player i, starting with the case where there is no ambiguity about the beliefs. This reformulation is crucial and itself proceeds in four steps.

Step 1(a) Single–prior case: For any player $i \in \mathbb{I}$ and any extension $\tilde{P} \in \mathcal{P}(\mathcal{F}_i)$ we have a product structure $\tilde{P} = (\tilde{\mu}_1, \ldots, \tilde{\mu}_I)$, by Proposition 1.2. We reformulate, for any $(\mathcal{F}_1, \ldots, \mathcal{F}_I)$–measurable strategy profile $(s_1, \ldots, s_I)$, the expected payoff $U_i^{\tilde{P}}$ under this possible prior $\tilde{P}$:

$$\begin{aligned}
U_i^{\tilde{P}}(s) = \mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})] &= \int_\Omega u_i(\omega, s_i(t_i), s_{-i}(t_{-i}))\mathrm{d}\tilde{P}(\omega) \\
&= \int_{T_i \times \prod_{j\neq i} T_j} u_i(\omega, s_i(t_i), s_{-i}(t_{-i}))\mathrm{d}\tilde{P}(t_1, \ldots, t_I) \\
&= \int_{T_i \times \prod_{j\neq i} T_j} u_i(\omega, s_i(t_i), s_{-i}(t_{-i}))\mathrm{d}\Big(\tilde{\mu}_i(t_i) \times \prod_{j\neq i} \tilde{\mu}_j(t_j)\Big) \\
&= \int_{T_i \times A_{-i}} u_i(t_i, s_i(t_i), a_{-i})\mathrm{d}\Big(\tilde{\mu}_i(t_i) \times \prod_{j\neq i} \tilde{\mu}_j \circ s_j^{-1}(a_j)\Big) \\
&= \int_{A_{-i}} \Big(\int_{T_i} u_i(t_i, s_i(t_i), a_{-i})\mathrm{d}\tilde{\mu}_i(t_i)\Big)\, \mathrm{d}\Big(\prod_{j\neq i} \tilde{\mu}_j \circ s_j^{-1}(a_j)\Big)
\end{aligned}$$

where the last but one equality applies Assumption 1(iii) and the last equality employs the fact that the strategy of player i only depends on her type. From the last calculation, we get with Assumption 1(i) and Fubini's Theorem:

$$
\begin{aligned}
\mathbb{E}^{\tilde{P}}[u_i(s_i, s_{-i})] &= \int_{A_{-i}} \Big(\int_{T_i} u_i(t_i, s_i(t_i), a_{-i}) \mathrm{d}\tilde{\mu}_i \Big) \mathrm{d}\Big(\prod_{j \neq i} \tilde{\mu}_j \circ s_j^{-1} \Big) \\
&= \int_{A_{-i}} \Big(\int_{T_i} u_i(t_i, s_i(t_i), a_{-i}) \frac{\mathrm{d}\tilde{\mu}_i}{\mathrm{d}\lambda_i} \mathrm{d}\lambda_i \Big) \mathrm{d} \prod_{j \neq i} \tilde{\mu}_j \circ s_j^{-1} \\
&= \int_{T_i} \Big(\int_{A_{-i}} \underbrace{u_i(t_i, s_i(t_i), a_{-i}) \frac{\mathrm{d}\tilde{\mu}_i}{\mathrm{d}\lambda_i}}_{=: \tilde{u}_i(t_i, s_i(t_i), a_{-i})} \mathrm{d} \prod_{j \neq i} \tilde{\mu}_j \circ s_j^{-1} \Big) \mathrm{d}\lambda_i =: U_i^{\tilde{\mu} \circ s_{-i}}(s_i).
\end{aligned}
$$

Here $\tilde{u}_i$ is defined in the obvious way, see also the end of Step 1(c). Moreover, note that $\tilde{\mu} \circ s_{-i} = \prod_{j \neq i} \tilde{\mu}_j \circ s_j^{-1}$ in the definition of $U_i^{\tilde{\mu} \circ s_{-i}}$ only matters for $j \neq i$. In the following step, payoffs of player i (under $\tilde{P}$) will only depend on his own action.

Step 1(b) Variation of opponents' actions: As a preparation of Step 1(d), we modify the function $U_i^{\tilde{P}}$ from step 1 (a) by abstracting from $\tilde{\mu} \circ s_{-i}$. As $U_i^{\tilde{\mu} \circ s_{-i}}$ only depends on s_i, we can substitute $\tilde{\mu} \circ s_{-i}$ by an arbitrary other product measure $\gamma \in \prod_{j \in \mathbb{I}} \Delta(A_j)$, where the i–th component of γ remains inactive. Set

$$
U_i^{\gamma}(s_i) = \int_{T_i} \left(\int_{A_{-i}} \tilde{u}_i(t_i, s_i(t_i), a_{-i}) \mathrm{d} \prod_{j \neq i} \gamma_j(a_{-i}) \right) \mathrm{d}\lambda_i(t_i) = \int_{T_i} \mathtt{u}_i(t_i, s_i(t_i), \gamma) \mathrm{d}\lambda_i(t_i)
$$

with

$$
\mathtt{u}_i(t_i, a_i, \gamma) = \int_{A_{-i}} \tilde{u}_i(t_i, a_i, a_{-i}) \mathrm{d} \prod_{j \neq i} \gamma_j(a_{-i}),
$$

where λ_i again denotes the projection of $\mathbb{P}$ on $(\Omega, \mathcal{F})$ to $(T_i, \mathcal{T}_i)$, by Assumption 1(i).[31]

Step 1(c) The α-maxmin case: In Step 1(a), we computed $U_i^P(s) = \mathbb{E}^P[u_i(s_i, s_{-i})]$ for any $P \in \mathcal{P}(\mathcal{F}_i)$ and feasible profile $s = (s_i, s_{-i})$, by employing the underlying atomless structure. In order to connect to the HEP W_i from (2), we introduce the convex and weakly compact (via Proposition 1.1 and the linearity of the expectation) sets of minimizers and maximizers

$$
\underline{\mathcal{P}}_i(s) = \arg\min_{P \in \mathcal{P}(\mathcal{F}_i)} \mathbb{E}^P[u_i(s_i, s_{-i})] \qquad \text{and} \qquad \overline{\mathcal{P}}_i(s) = \arg\max_{P \in \mathcal{P}(\mathcal{F}_i)} \mathbb{E}^P[u_i(s_i, s_{-i})].
$$

By applying Berge's maximum theorem, see page 570 in Aliprantis and Border (2006), both correspondences $L \ni s \mapsto \underline{\mathcal{P}}_i(s)$ and $L \ni s \mapsto \overline{\mathcal{P}}_i(s)$ are $\| \cdot \|_\infty$ to $\sigma(L^2, L^2)$ upper hemi–continuous. The weighted (Minkowski) sum of two convex and compacts sets

$$
\underline{\overline{\mathcal{P}}}^{\alpha_i}(s) = \alpha_i \underline{\mathcal{P}}_i(s) + (1 - \alpha_i) \overline{\mathcal{P}}_i(s) \tag{4}
$$

[31] In Step 1(d) the reason for the measure change in U_i^{γ} from $\tilde{\mu}_i$ to λ_i will become apparent.

is nonempty, convex, compact and a subset of $\mathcal{P}(\mathcal{F}_i)$. Moreover, $s \mapsto \underline{\overline{\mathcal{P}}}^{\alpha_i}(s)$ is upper hemi–continuous, as a sum of upper hemi–continuous correspondences.

In view of Step 1(a) and by Lemma 1, we always have $W_i(s_i, s_{-1}) = U_i^{\underline{\overline{P}}^s}(s)$ for some $\underline{\overline{P}}^s \in \underline{\overline{\mathcal{P}}}^{\alpha_i}(s)$ and any feasible strategy profile $s = (s_i, s_{-i})$. We can now repeat the same derivation from Step 1(a) and employ $\tilde{u}_i$. By Proposition 1.2, we get for any prior $(\underline{\overline{\mu}}_1^s, \dots, \underline{\overline{\mu}}_I^s) = \underline{\overline{P}}^s \in \underline{\overline{\mathcal{P}}}^{\alpha_i}(s)$:

$$\begin{aligned} W_i(s_i, s_{-i}) &= W_i^{\underline{\overline{P}}^s}(s_i, s_{-1}) := U_i^{\underline{\overline{P}}^s}(s) \\ &= \int_{T_i} \left(\int_{A_{-i}} \tilde{u}_i(t_i, s_i(t_i), a_{-i}) \mathrm{d} \prod_{j \neq i} \underline{\overline{\mu}^s}_j \circ s_j^{-1} \right) \mathrm{d}\lambda_i. \end{aligned}$$

Step 1(d) Variation of the α-maxmin case: With $\gamma \in \prod_{i \in \mathbb{I}} \Delta(A_i)$, the analog of U_i^γ in step 1 (b), denoted by W_i^γ and departing from $W_i^{\underline{\overline{P}}}$, has to incorporate the (s_i, s_{-i})-dependency on $\underline{\overline{\mathcal{P}}}^{\alpha_i}(s)$ when selecting an effective belief. As in step 1(c) set

$$W_i^P(s_i) = \int_{T_i} \left(\int_{A_{-i}} \tilde{u}_i(t_i, s_i(t_i), a_{-i}) \mathrm{d} \prod_{j \neq i} \mu_j \circ s_j^{-1} \right) \mathrm{d}\lambda_i(t_i) = \int_{T_i} \mathtt{w}_i(t_i, s_i(t_i), P) \mathrm{d}\lambda_i(t_i), \quad (5)$$

with $(\mu_1, \dots \mu_I) = P \in \underline{\overline{\mathcal{P}}}^{\alpha_i}(s)$, where

$$\mathtt{w}_i(t_i, a_i, P) = \frac{\mathrm{d}\mu_i}{\mathrm{d}\lambda_i}(t_i) \cdot \int_{A_{-i}} \tilde{u}_i(t_i, a_i, a_{-i}) \mathrm{d} \prod_{j \neq i} \mu_j \circ s_j^{-1}. \quad (6)$$

The Radon–Nikodym density $\frac{\mathrm{d}\mu_i}{\mathrm{d}\lambda_i}$ exists by Proposition 1.2 and Assumption 1(iv). Again by the product structure of P, $P \mapsto \mathtt{w}_i(a_i, t_i, P)$ is weakly $\sigma(L^2, L^2)$–continuous for any $(t_1, \dots, t_I) = \omega \in \Omega$ and $i \in \mathbb{I}$. Moreover, by the definition of $\tilde{u}_i$, $t_i \mapsto \mathtt{w}_i(t_i, a_i, P)$ is $\mathcal{F}_i$–measurable for any P. For the given $\mathbb{P} = \prod_{i \in \mathbb{I}} \lambda_i$ from Assumption 1 consider the correspondence

$$f : \prod_{i \in \mathbb{I}} \Delta(A_i) \rightrightarrows \prod_{i \in \mathbb{I}} L(T_i, \mathcal{F}_i; A_i), \quad (7)$$

that is induced by $\gamma_i = \lambda_i \circ s_i^{-1} \in \Delta(A_i)$, for any $i \in \mathbb{I}$. To see that f is upper hemi–continuous consider the function

$$f^{-1} : \prod_{i \in \mathbb{I}} L(T_i, \mathcal{F}_i; A_i) \to \prod_{i \in \mathbb{I}} \Delta(A_i),$$

that assigns component–wise the product law of each strategy profile, we have $f_i^{-1}(s_i) = \gamma_i$. Clearly, if $s_i^n \to s_i$ in the sup–norm, then $\gamma_i^n \to \gamma_i$ in $\mathbb{R}^{|A_i|}$. Hence, f^{-1} is a closed mapping and by Theorem 17.7 in Aliprantis and Border (2006) we have the upper hemi–continuity of f.

Consequently, as a composition of the upper hemi–continuous correspondences $\overline{\underline{\mathcal{P}}}^{\alpha_i}$ from (4), $\mathtt{w}_i$ from (6) and f from (7), the composed correspondence

$$\gamma = (\gamma_1, \ldots, \gamma_I) \mapsto \mathtt{w}_i\Big(t_i, a_i, \overline{\underline{\mathcal{P}}}^{\alpha_i}\big(f(\gamma)\big)\Big), \tag{8}$$

is again upper hemi–continuous for any $t_i \in T_i$ and $a_i \in A_i$.

By plugging (8) into (5), we have a set-valued analog W_i^γ of U_i^γ from Step 1(b) that additionally incorporates the (s_i, s_{-i})-dependency on the set of all effective priors $\overline{\underline{\mathcal{P}}}^{\alpha_i}(s)$.

Step 2 Preparation of the Fixed–Point argument: With the mapping from (8), define for each player i, the correspondence $\Psi_i : T_i \times [[0, \sigma_i]] \Rightarrow A_i$ by

$$\Psi_i(t_i, P) = \arg\max_{a_i \in A_i} \mathtt{w}_i(t_i, a_i, P),$$

where $[[0, \sigma_i]] = \{x \in L^2 : 0 \leq x \leq \sigma_i \ \mathbb{P}\text{-}a.s.\}$ denotes the weakly compact order interval of relevant Rodon-Nykodym densities that identify P, see Assumption 1(iv). For any $t_i \in T_i$, $\Psi_i(t_i, \cdot)$ is upper hemi–continuous on $[[0, \sigma_i]]$, by Berge's maximum theorem. Moreover, for any P with $\frac{\mathrm{d}P}{\mathrm{d}\mathbb{P}} \in [[0, \sigma_i]]$, $\Psi_i(\cdot, P)$ is $\mathcal{F}_i$–measurable, see Theorem 18.19 in Aliprantis and Border (2006) p.605.

Finally, define the composed correspondence $\Xi_i : \Omega \times \prod_{i \in \mathbb{I}} \Delta(A_i) \Rightarrow A_i$ by

$$\Xi_i(t_i, \gamma) = \Psi_i\Big(t_i, \overline{\underline{\mathcal{P}}}^{\alpha_i}\big(f(\gamma)\big)\Big). \tag{9}$$

Since $\overline{\underline{\mathcal{P}}}^{\alpha_i}$ (by step 1(c)) and f (by step 1(d)) are both upper hemi–continuous, the composition Ξ_i is then again upper hemi–continuous, see Theorem 17.23 Aliprantis and Border (2006) p.566.

Let γ_i^* denote the distribution $\lambda_i \circ (s_i^*)^{-1}$ of player i on A_i. A pure strategy equilibrium s^* for the non–Bayesian game Γ is equivalent to the maximality, for each player $i \in \mathbb{I}$, of s_i^* with respect to W_i^P, defined in (5), on $L(T_i, \mathcal{F}_i; A_i)$, as long as P is concatenated by a α_i convex combination of a minimizer and a maximizer at $s^* = (s_i^*, s_{-i}^*)$. In view of step 1(c), that is $P \in \overline{\underline{\mathcal{P}}}^{\alpha_i}(s^*)$. But this condition holds if s_i^* is a measurable selection of $\Xi_i(\cdot, \gamma^*)$, defined in (9), as Ξ_i monitors the usage of the correct composed prior.

Step 3 The Fixed–Point argument: We recover the equilibrium strategy from the fixed–point distribution of the correspondence. In order to show the existence of a pure-strategy equilibrium $(s_1^*, \ldots, s_I^*)$ for the game Γ, define the correspondence $\mathscr{G}_i : \prod_{j \in \mathbb{I}} \Delta(A_j) \Rightarrow \Delta(A_i)$ by

$$\mathscr{G}_i(\gamma) = \mathscr{G}_{\Xi_i(\cdot,\gamma)} = \Big\{\lambda_i \circ \phi^{-1} : \phi \in \mathtt{MB}^{sel}\Big(\Xi_i(\cdot, \gamma)\Big)\Big\}, \quad i \in \mathbb{I}$$

where $\texttt{MB}^{sel}$ denotes the collection of all measurable selections of a correspondence.

By the underlying atomless type space structure in Assumption 1(i), the combined correspondence $\mathscr{G} : \prod_{i\in\mathbb{I}} \Delta(A_i) \Rightarrow \prod_{i\in\mathbb{I}} \Delta(A_i)$, defined by $\mathscr{G}(\gamma) = \big[\mathscr{G}_1(\gamma), \ldots, \mathscr{G}_I(\gamma)\big]$, is convex and compact valued, and upper-hemicontinuous: Convex and compact valuedness holds for each component $\mathscr{G}_i$ by Proposition 2.1 and upper-hemicontinuity follows from Proposition 2.2, since $\Xi_i(t_i, \cdot)$ is upper-hemicontinuous for each player i, as shown in step 2. As a condition in Proposition 2, note that the space $(Y =) \prod_{j\in\mathbb{I}} \Delta(A_j)$ is a metric space, where the metric is induced by the Euclidean norm.

We apply the Kakutani fixed–point theorem and get the existence of a $\gamma^* \in \mathscr{G}(\gamma^*)$. In other words, for each player i there is an $\mathcal{F}_i$–measurable selection s_i^* of the correspondence $\Xi_i(\cdot, \gamma^*)$ satisfying $\gamma_i^* = \lambda_i \circ s_i^{*-1}$.

We can now appeal to Step 2 to assert the existence of a pure strategy Hurwicz–Nash equilibrium with strategies $(s_i^*)_{i\in\mathbb{I}}$.

6 Proofs of Auxiliary Results

We give the postponed proofs of Proposition 1 and Lemma 1 from Subsection 2.3.

Proof of Proposition 1: We take each assertion of the Proposition in turn.

1. Convexity follows immediately. To show compactness, it suffice to prove that $\mathcal{P}(\mathcal{G})$ is weakly closed, where

$$\mathcal{P}(\mathcal{G}) = \Big\{P \in \Delta(\Omega, \mathcal{F}) : P = P_i \text{ on } \mathcal{G}\Big\} \bigcap \Big\{P \in \Delta(\Omega, \mathcal{F}) : P \gg \mathbb{P} \text{ with } \frac{\mathrm{d}P}{\mathrm{d}\mathbb{P}} \in [[0, \sigma]] \subset L^2(\mathbb{P})\Big\}.$$

   The latter set is closed, since

$$\Big\{P : P \gg \mathbb{P} \text{ with } \frac{\mathrm{d}P}{\mathrm{d}\mathbb{P}} \in [[0, \sigma]]\Big\} = [[0, \sigma]] \bigcap \Big\{\frac{\mathrm{d}P}{\mathrm{d}\mathbb{P}} \in L^2(\mathbb{P}) : E^{\mathbb{P}}\Big[\frac{\mathrm{d}P}{\mathrm{d}\mathbb{P}}\Big] = 1\Big\}$$

   is the intersection of L^2-norm closed and convex sets, and hence weakly closed, since the weak and norm topology coincide on convex sets. But this means that $\mathcal{P}(\mathcal{G})$ is also closed since any converging sequence with $P^n \in \mathcal{P}(\mathcal{G})$ satisfies $P^n = P_i$ on $\mathcal{G}$. Hence, we have $\lim_n P^n = P_i$ on $\mathcal{G}$. The boundedness of $\mathcal{P}(\mathcal{G})$, by Assumption 1(iv), in combination with Alaoglu's Theorem gives the result.

2. The fact that any extension is atomless follows directly from the absolute continuity of the atomless $\mathbb{P}$. For a proof we refer to Theorem 3.1, p.265, of Belley, Dubois, and Morales (1983). Moreover, note that by Assumption 1(i) any player is aware of the product structure of the objective probability measure $\mathbb{P}$. Consequently, any player only considers those extensions that satisfy the atomless product structure.

$\square$

Proof of Lemma 1: By Proposition 1.1, $\mathcal{P}(\mathcal{F}_i)$ is convex and $\sigma(L^2, L^2)$–compact. Hence we can derive the α-maxmin payoff as the expected payoff under some specific prior $\overline{\underline{P}}^{\alpha_i,s}$, as in (4) and depending on the particular strategy profile (s_i, s_{-i}):

$$\begin{aligned} W_i(s_i, s_{-i}) &= \alpha_i \min_{P\in\mathcal{P}(\mathcal{F}_i)} \mathbb{E}^P[u_i(s_i, s_{-i})] + (1-\alpha_i) \max_{P\in\mathcal{P}(\mathcal{F}_i)} \mathbb{E}^P[u_i(s_i, s_{-i})] \\ &= \alpha_i\, \mathbb{E}^{\underline{P}_i^s}[u_i(s_i, s_{-i})] + (1-\alpha_i)\mathbb{E}^{\overline{P}_i^s}[u_i(s_i, s_{-i})] \\ &= \mathbb{E}^{\alpha_i \underline{P}_i^s + (1-\alpha_i)\overline{P}_i^s}[u_i(s_i, s_{-i})] = \mathbb{E}^{\overline{\underline{P}}^{\alpha_i,s}}[u_i(s_i, s_{-i})] = \mathbb{E}^{\mathbb{P}}\left[\frac{\mathrm{d}\overline{\underline{P}}^{\alpha_i,s}}{\mathrm{d}\mathbb{P}} u_i(s_i, s_{-i})\right]. \end{aligned}$$

The existence of a minimizer $\underline{P}_i$ and a maximizer $\overline{P}_i$ follows from the weak compactness of $\mathcal{P}(\mathcal{F}_i)$ (by Proposition 1.1) and the linearity of $P \mapsto \mathbb{E}^P[u_i(s_i, s_{-i})]$, which implies weak continuity.

The concatenated prior $\overline{\underline{P}}^{\alpha_i,s}$ lies in the set $\mathcal{P}(\mathcal{F}_i)$, which is convex by Proposition 1.1, and depends on the strategy profile $s = (s_i, s_{-i})$. The last equality in the former derivation follows from Assumption 1(ii).

In order to check that W_i is finite, we apply the Cauchy–Schwarz inequality and get

$$W_i(s_i, s_{-1}) \leq \mathbb{E}^{\mathbb{P}}\left[\left(\frac{\mathrm{d}\overline{\underline{P}}^{\alpha_i,s}}{\mathrm{d}\mathbb{P}}\right)^2\right]^{1/2} \mathbb{E}^{\mathbb{P}}\left[(u_i(s_i, s_{-i}))^2\right]^{1/2} < \infty. \tag{10}$$

For the finiteness of the first factor after the first inequality of (10), note that

$$\frac{\mathrm{d}\overline{\underline{P}}^{\alpha_i,s}}{\mathrm{d}\mathbb{P}} = \alpha_i \frac{\mathrm{d}\underline{P}_i}{\mathrm{d}\mathbb{P}} + (1-\alpha_i)\frac{\mathrm{d}\overline{P}_i}{\mathrm{d}\mathbb{P}}$$

is $\mathbb{P}$–square integrable, since $\frac{\mathrm{d}\underline{P}_i}{\mathrm{d}\mathbb{P}}, \frac{\mathrm{d}\overline{P}_i}{\mathrm{d}\mathbb{P}} \in L^2(\mathbb{P})$ by the construction of $\mathcal{P}_i(\mathcal{F}_i)$.

The finiteness of the second factor follows from the square integrability of u_i guaranteed by Assumption 1(ii). $\square$

References

Ahn, D. (2007): "Hierarchies of ambiguous beliefs," *Journal of Economic Theory*, 136(1), 286–301.

Al-Najjar, N. I., and J. Weinstein (2015): "A Bayesian model of Knightian uncertainty," *Theory and Decision*, 78(1), 1–22.

Aliprantis, C., and K. Border (2006): *Infinite dimensional analysis: a hitchhiker's guide*. Berlin: Springer Verlag.

Aumann, R. J. (1974): "Subjectivity and correlation in randomized strategies," *Journal of Mathematical Economics*, 1(1), 67–96.

Azrieli, Y., and R. Teper (2011): "Subjectivity and correlation in randomized strategies," *Games and Economic Behavior*, 73(2), 310–317.

Bade, S. (2011): "Ambiguous act equilibria," *Games and Economic Behavior*, 2(71), 246–260.

Belley, J. M., J. Dubois, and P. Morales (1983): *Measure theory and its applications: Proceedings of a conference held at Sherbrooke*. Berlin: Springer Verlag.

Bewley, T. F. (1986): "Knightian decision theory. Part I," *Cowles Foundation Discussion Paper No. 807*.

Bělohlávek, R., B. Dauben, Joseph, and G. J. Klir (2017): *Fuzzy logic and mathematics: A historical perspective*. Oxford: Oxford University Press.

Couso, I., and S. Moral (2010): "Independence concepts in evidence theory," *International Journal of Approximte Reason*, 51(7), 748–758.

Couso, I., S. Moral, and P. Walley (1999): "Examples of independence for imprecise probabilities," in *Proceedings of the First Symposium on Imprecise Probabilities and Their Applications*, ed. by G. de Cooman. Ghent: ISIPTA.

Eichberger, J., S. Grant, and D. Kelsey (2008): "Differentiating ambiguity: An expository note," *Economic Theory*, 36(2), 327–336.

Friedenberg, A., and M. Meier (2015): "Differentiating ambiguity: An expository note," *mimeo, Arizona State University.*

Fu, H., Y. Sun, N. C. Yannelis, and Z. Zhang (2007): "Pure strategy equilibria in games with private and public information," *Journal of Mathematical Economics*, 43(5), 523–531.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004): "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 118(2), 133–173.

Gilboa, I. (1987): "Expected utility with purely subjective non-additive probabilities," *Journal of Mathematical Economics*, 16(1), 65–88.

——— (2009): *Theory of decision under uncertainty*. Cambridge: Cambridge University Press.

Gilboa, I., and M. Marinacci (2016): "Ambiguity and the Bayesian paradigm," in *Readings in Formal Epistemology*, pp. 385–439. Berlin: Springer.

Grant, S., I. Meneghel, and R. Tourky (2016): "Savage games," *Theoretical Economics*, 11(2), 641–682.

Gul, F., and W. Pesendorfer (2014): "Expected uncertain utility theory," *Econometrica*, 82(1), 1–39.

——— (2015): "Hurwicz expected utility and subjective sources," *Journal of Economic Theory*, 159(1), 465–488.

Hansen, L., and T. Sargent (2001): "Robust control and model uncertainty," *American Economic Review*, 91(2), 60–66.

Harsanyi, J. (1967/68): "Games with incomplete information played by Bayesian players, I-III," *Management Science*, 14(3/5/7), 59–182, 320–334, 486–502.

He, W., X. Sun, and Y. Sun (2017): "Modeling infinitely many agents," *Theoretical Economics*, 12(2), 771–815.

He, W., X. Sun, Y. Sun, and Y. Zeng (2015): "Pure strategy equilibrium in general Bayesian games," *Preliminary draft.*

Heifetz, A., and P. Mongin (2001): "Probability logic for type spaces," *Games and Economic Behavior*, 35(1), 31–53.

Hurwicz, L. (1979): "On the interaction between information and incentives in organizations," in *Communication and control in society*, ed. by K. Krippendorff, pp. 123–147. New York: Gordon and Breach.

——— (2007): "But who will guard the guardians?," *Nobel lecture, available at www.nobelprize.org*.

Kajii, A., and T. Ui (2005): "Incomplete information games with multiple priors," *Japanese Economic Review*, 56(3), 332–351.

Kamenca, E. (2017): "Information economics," *Journal of Political Economy*, forthcoming.

Keynes, J. (1937): "The general theory of employment," *The Quarterly Journal of Economics*, 51(2), 209–223.

Khan, M. A. (2008): "Perfect foresight," in *International Encyclopedia of the Social Sciences, 2nd Edition*, ed. by W. A. Darity Jr. New york:Macmillan.

Khan, M. A., and K. P. Rath (2009): "On games with incomplete information and the Dvoretsky–Wald–Wolfowitz theorem with countable partitions," *Journal of Mathematical Economics*, 45(12), 830–837.

Khan, M. A., K. P. Rath, and Y. Sun (2006): "The Dvoretzky-Wald-Wolfowitz theorem and purification in atomless finite-action games," *International Journal of Game Theory*, 34(1), 91–104.

Khan, M. A., and Y. Sun (1995): "Pure strategies in games with private information," *Journal of Mathematical Economics*, 24(7), 633–653.

——— (1996): "Nonatomic games on Loeb spaces," *Proceedings of the National Academy of Sciences, USA*, 93(26), 15518–15521.

——— (1999): "Non-cooperative games on hyperfinite Loeb spaces," *Journal of Mathematical Economics*, 31(4), 455–492.

Khan, M. A., Y. Sun, R. Tourky, and Z. Zhang (2008): "Similarity of differential information with subjective prior beliefs," *Journal of Mathematical Economics*, 44(9), 1024–1039.

Khan, M. A., and Y. Zhang (2014): "On the existence of pure-strategy equilibria in games with private information: A complete characterization," *Journal of Mathematical Economics*, 50(1), 197–202.

——— (2017): "On pure-strategy equilibria in games with correlated information," *Games and Economic Behavior, forthcoming.*

Knight, F. (1921): *Risk, uncertainty and profit.* New York: Beard Books.

Kyberg Jr., H. (1977): "Randomness and the right reference class," *Journal of Philosophy*, 74(9), 501–521.

Leroy, S. F., and L. D. Singell (1987): "Knight on risk and uncertainty," *Journal of Political Economy*, 95(2), 394–406.

Levi, I. (1977): "Direct inference," *Journal of Philosophy*, 74(1), 5–29.

——— (1987): "Conflict and social agency," *Journal of Philosophy*, 79(5), 231–247.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006): "Ambiguity aversion, robustness, and the variational representation of preferences," *Econometrica*, 74(6), 1447–1498.

Machina, M. J. (1982): "'Expected Utility' analysis without the independence axiom," *Econometrica*, 50(2), 277–324.

Marinacci, M. (2000): "Ambiguous games," *Games and Economic Behavior*, 31(2), 191–219.

Maschler, M., E. Solan, and S. Zamir (2013): *Game theory: An introduction.* Cambridge: Cambridge University Press.

Milgrom, P. R., and R. J. Weber (1985): "Distributional strategies for games with incomplete information," *Mathematics of Operations Research*, 10(4), 619–632.

Mirowski, P., and E. Nik-Khah (2017): *The knowledge we have lost in information.* Oxford: Oxford University Press.

Morris, S. (1995): “The common prior assumption in economic theory,” *Economics and Philosophy*, 11(02), 227–253.

Pedersen, P., and G. Wheeler (2014): “Demystifying dilations,” *Erkenntnis*, 79(06), 1305–1342.

Pellesoni, R., and P. Vicig (2016): “2-Coherent and 2-convex conditional lower previsions,” *International Journal of Approximate Reasoning*, 77(1), 66–86.

Radner, R. (1968): “Competitive equilibrium under uncertainty,” *Econometrica*, 36(1), 31–58.

Radner, R., and R. W. Rosenthal (1982): “Private information and pure-strategy equilibria,” *Mathematics of Operations Research*, 7(3), 401–409.

Reiter, S. (2009): “Two topics in Leo Hurwicz’s research,” *Review of Economic Design*, 13(1–2), 3–6.

Riedel, F., and L. Sass (2014): “Ellsberg games,” *Theory and Decision*, 76(4), 469–509.

Schmeidler, D. (1982): “Subjective probability without additivity,” *Foerder Institute for Economic Research Working Paper, Tel-Aviv University*, 57.

——— (1989): “Subjective probability and expected utility without additivity,” *Econometrica*, 57(3), 571–587.

Seidenfeld, T., and L. Wasserman (1993): “Dilations for sets of probabilities,” *The Annals of Statistics*, 21(3), 1139–1154.

Stinchcombe, M. (2008): “A consequentilaist approach to ambiguity,” *mimeo, University of Texas.*

Vicig, P., and T. Sidenfeld (2012): “Bruno de Finetti and imprecision: Imprecise probability does not exist!,” *International Journal of Approximate Reasoning*, 53(8), 1115–1123.

Von Hayek, F. A. (1937): “Economics and knowledge,” *Economica*, 4(13), 33–54.

Walley, J. (1991): *Statistical reasoning with imprecise probabilities*. Chapman and Hall, New York.

Yager, R. R., M. Fedrizzi, and J. e. Kacprzyk (1994): *Advances in the Dempster-Shafer theory of evidence*. John Wiley and Sons, New York.

Zhang, Y., X. Luo, W. Ma, and H. Leung (2014): “Ambiguous Bayesian games,” *International Journal of Intelligent Systems*, 29(12), 1138–1172.