

Cianciaruso, Davide; Germano, Fabrizio

Working Paper

Quotient spaces of boundedly rational types

Discussion Paper, No. 1539

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Cianciaruso, Davide; Germano, Fabrizio (2011) : Quotient spaces of boundedly rational types, Discussion Paper, No. 1539, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/59688>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1539

"Quotient Spaces of Boundedly Rational Types"

Daide Cianciaruso⁺ and Fabrizio Germano⁺⁺

September 20, 2011

JEL Classification: C72, D03, D83.

Keywords: Incomplete-information games, high-order reasoning, type space, quotient space, hierarchies of beliefs, bounded rationality.



NORTHWESTERN
UNIVERSITY



⁺ Northwestern University, daide.cianciaruso@gmail.com

⁺⁺ Universitat Pompeu Fabra and Barcelona GSE, fabrizio.germano@upf.edu

Quotient Spaces of Boundedly Rational Types*

Davide Cianciaruso[†] and Fabrizio Germano[‡]

September 20, 2011

Abstract

By identifying types whose low-order beliefs – up to level ℓ_i – about the state of nature coincide, we obtain quotient type spaces that are typically smaller than the original ones, preserve basic topological properties, and allow standard equilibrium analysis even under bounded reasoning. Our Bayesian Nash (ℓ_i, ℓ_{-i}) -equilibria capture players' inability to distinguish types belonging to the same equivalence class. The case with uncertainty about the vector of levels (ℓ_i, ℓ_{-i}) is also analyzed. Two examples illustrate the constructions.

Keywords: Incomplete-information games, high-order reasoning, type space, quotient space, hierarchies of beliefs, bounded rationality. *JEL Classification:* C72, D03, D83.

1 Introduction

Starting with Harsanyi's (1967-1968) seminal work, players' private information is conveniently represented by *types*, which correspond to infinite hierarchies of beliefs. In many applications, the predictions of standard game theory can be heavily dependent on the specification of high-order beliefs. As a matter of fact, Weinstein and Yildiz (2007a) show that any rationalizable action (Battigalli and Siniscalchi, 2003, and Dekel, Fudenberg, and Morris, 2007) of any type can be made into the *uniquely* rationalizable action for an open set of perturbed types under an appropriate class of beliefs perturbations. This generalizes the intuition provided by Carlsson and van Damme (1993) in their *global games*. Nonetheless, experiments have highlighted that high-order beliefs might have a less large relevance in practice (see the literature mentioned in the context of the examples given in Sections 3.1.1 and 4.1.1). Weinstein and Yildiz (2007b) write that bounds on rationality translate into bounds on beliefs, and consequently situations where high-order beliefs are decisive might necessitate a change of paradigm.

In this paper we provide an approach towards relaxing the (implicit) assumptions made on players' ability to keep track of high-order beliefs. The approach we propose consists of defining certain quotient type spaces,

*We thank Larbi Alaoui, Alfredo Di Tillio, Stephen Hansen, Aviad Heifetz, Frank Heinemann, Martin Meier, Rosemarie Nagel, Omiros Papaspiliopoulos, Marciano Siniscalchi, Jonathan Weinstein, and Shmuel Zamir for helpful discussions. Germano acknowledges financial support from the Spanish Ministry of Science and Technology (Grant SEJ2007-64340), from the Spanish Ministry of Education (in form of a Madariaga Fellowship), as well as from the Barcelona GSE Research Network and the Generalitat de Catalunya. Germano also thanks the CMS-EMS center at Kellogg, Northwestern University, for their generous support and hospitality. All mistakes are our own.

[†]Northwestern University, davide.cianciaruso@gmail.com

[‡]Universitat Pompeu Fabra and Barcelona GSE, fabrizio.germano@upf.edu

the elements of which gather – or “amalgamate” in the usage of Aumann and Dreze (2008, Section VI) – types that have the same low-order beliefs, where low means up to some order ℓ_i . Such quotient spaces will be referred to as ℓ_i -quotient type spaces. The underlying behavioral hypothesis is that a player acts as if he could not tell his own types apart that have identical belief hierarchies from order 1 to ℓ_i . Players hence condition their interim expected payoff on their type belonging to its *equivalence class* rather than on the type itself.

We would like to stress that our quotient type spaces embody *one* specific behavioral hypothesis. A researcher interested in modeling other aspects of human behavior can adopt different notions of similarity between types and possibly define quotient type spaces accordingly.

Standard equilibrium analysis can be performed by means of ℓ_i -quotient type spaces: in an (ℓ_i, ℓ_{-i}) -*equilibrium* each player i plays a strategy that is constant within his equivalence classes and maximizes his expected payoff. If (ℓ_i, ℓ_{-i}) is uncertain, we assume a commonly known common prior λ on levels and let players with different quotient spaces be different players that possess the same information as the original players from whom they are derived: in a λ -*equilibrium* each of these new players has a strategy that is constant within his equivalence classes and maximizes the expectation of his payoff with respect to the state of nature, the other players’ type, and the other players’ levels ℓ_{-i} as well.

The way in which we introduce uncertainty about players’ levels of sophistication deserves a few comments. In the class of games we consider, we have a common prior f on the state of nature, player i ’s type, and player $-i$ ’s type – respectively denoted by θ , t_i , and t_{-i} . We can think of f as the marginal of another distribution – denote it F – on θ , t_i , t_{-i} , ℓ_i , and ℓ_{-i} . Then, λ is the marginal of F on ℓ_i and ℓ_{-i} . We assume that F is given by the product of f and λ , which is equivalent to (ℓ_i, ℓ_{-i}) being independent of (θ, t_i, t_{-i}) .¹ In the original game, a type t_i has a posterior on θ and t_{-i} , but once we allow for the possibility of boundedly rational players, a type t_i of level ℓ_i should have a posterior on θ , t_{-i} , and ℓ_{-i} . Indeed, t_i needs to know what $-i$ is conditioning his expected payoff on in order to outguess what action $-i$ is going to take. The question is how to formalize this idea. One could let player i ’s strategy map each (t_i, ℓ_i) into an action. Otherwise, one could: (i) as usual, let player i ’s strategy be a map that associates an action to each t_i ; (ii) enlarge the set of players and treat a player i of level ℓ_i and any player i of level $\ell'_i \neq \ell_i$ as distinct players. Since our primary objective is to provide researchers with a familiar and parsimonious tool for dealing with bounded rationality in games with incomplete information, we opted for the latter solution. By means of suitable transformations of the payoff functions, we obtain that in our λ -equilibria players do not have to form beliefs about the levels of sophistication (i.e., player i does not have to think about ℓ_{-i} , about what $-i$ thinks about ℓ_i , and so forth). Importantly, this does *not* prevent us from deriving results that match behavioral patterns found in the laboratory.²

As an illustration of the theory, we consider two well-known examples where high-order beliefs play a key role, namely, Rubinstein’s (1989) electronic mail game³ and Morris and Shin’s (2002) game with private and public information. We obtain, respectively, (ℓ_i, ℓ_{-i}) -equilibria and λ -equilibria that appear to be closer to experimental evidence than usual game-theoretical predictions.

The paper is organized as follows: Section 2 sets up the notation; Section 3 introduces the ℓ_i -quotient space and computes (ℓ_i, ℓ_{-i}) -equilibria for the electronic mail game; Section 4 presents the case where players’

¹In Section 4.1 we explain the technical reasons that motivate such an assumption.

²We are unaware of whether the two modeling choices are equivalent or one is more expressive than the other.

³See also Halpern (1986) on the coordinated-attack problem in computer science.

levels of sophistication (ℓ_i, ℓ_{-i}) in distinguishing types are uncertain and computes Morris and Shin's (2002) game with private and public information; Section 5 concludes by summarizing related research.

2 Preliminaries

We consider two-player static games of incomplete information. Players are denoted by $i, -i \in I = \{1, 2\}$, where $-i = j \in I$ such that $j \neq i$. The concepts developed straightforwardly extend to the case of finitely many players $2 \leq |I| < \infty$. Our setup builds on Siniscalchi (2008). At the center of the construction is the set of all *states of nature* Θ consisting of payoff-relevant information players are uncertain about. A typical element of Θ is denoted by θ .

For an arbitrary space S , let $\Delta(S)$ denote the set of probability measures on the Borel σ -algebra of S , endowed with the weak* topology. Product spaces are always endowed with the product topology and subspaces with the subspace topology. Countable spaces are endowed with the discrete topology. Given a Polish (i.e., separable and completely metrizable) space Θ , let $X^0 = \Theta$ and for $\ell \geq 1$ define recursively $X^\ell = X^{\ell-1} \times \Delta(X^{\ell-1})$. Let $H = \prod_{\ell \geq 1} \Delta(X^{\ell-1})$ be the space of all possible *belief hierarchies* with typical element $\mu = (\mu^1, \mu^2, \dots)$. Notice that H is Polish (e.g., by Corollary 3.39 and Theorem 15.15 in Aliprantis and Border, 2006). Refining H to $H^c = \{\mu \in H : \forall \ell \geq 1, \text{marg}_{X^{\ell-1}} \mu^{\ell+1} = \mu^\ell\}$ yields the space of *coherent* belief hierarchies. Brandenburger and Dekel (1993) show the existence of a *homeomorphism* $g^c : H^c \rightarrow \Delta(\Theta \times H)$. The space H^c can be further refined in the following manner. Let $H^0 = H^c$ and for every $k \geq 1$ define $H^k = \{\mu \in H^{k-1} : g^c(\mu)(\Theta \times H^{k-1}) = 1\}$. This yields $H^{cc} = \bigcap_{k \geq 1} H^k$, the space of belief hierarchies satisfying *common certainty of coherence*. Brandenburger and Dekel (1993) also show the existence of a homeomorphism $g^{cc} : H^{cc} \rightarrow \Delta(\Theta \times H^{cc})$. We refer to $\langle \Theta, (H^{cc})_{i \in I}, (g^{cc})_{i \in I} \rangle$ as the *universal Θ -based type space*.⁴ This notion can be generalized as follows.

Definition 1 A Θ -based type space is a tuple $\langle \Theta, (T_i)_{i \in I}, (f_i)_{i \in I} \rangle$, where Θ and each T_i are Polish, and each $f_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$ is continuous.

A point in T_i is denoted by t_i , which is a type for player i . $f_i(t_i)$ is type t_i 's posterior on the state of nature and on $-i$'s type. Throughout, there is a *commonly known common prior* $f \in \Delta(\Theta \times T_i \times T_{-i})$, such that each $f_i(t_i)(\cdot) = f(\cdot | t_i)$. A point $(\theta, t_i, t_{-i}) \in \Theta \times T_i \times T_{-i}$ is a *state of the world* and is the realization of the random $(\tilde{\theta}, \tilde{t}_i, \tilde{t}_{-i})$ distributed according to f .

Each type t_i in a Θ -based type space can be mapped into an infinite hierarchy of beliefs. To see this, let $h_i^1 : T_i \rightarrow \Delta(X^0)$ be given by $h_i^1(t_i) = \text{marg}_{X^0} f_i(t_i)$. $h_i^1(t_i) \in \Delta(X^0)$ is type t_i 's first-order belief, and corresponds to the marginal on the state of nature of player i 's posterior. For $\ell \geq 2$, let $h_i^\ell : T_i \rightarrow \Delta(X^{\ell-1})$ be given by

$$h_i^\ell(t_i)(E) = f_i(t_i)(\{(\theta, t_{-i}) \in \Theta \times T_{-i} : (\theta, h_{-i}^1(t_{-i}), \dots, h_{-i}^{\ell-1}(t_{-i})) \in E\}) \quad (1)$$

for every Borel subset $E \subseteq X^{\ell-1}$. $h_i^\ell(t_i) \in \Delta(X^{\ell-1})$ is type t_i 's ℓ th-order belief. Finally, let $h_i : T_i \rightarrow H$ be given by $h_i(t_i) = (h_i^1(t_i), h_i^2(t_i), \dots)$. $h_i(t_i) \in H$ is type t_i 's entire belief hierarchy.

⁴For an analogous yet technically different construction where Θ is compact, see Mertens and Zamir (1985); Heifetz (1995) covers the case where Θ is possibly non-compact Hausdorff and beliefs are regular Borel probability measures.

We say that $\Theta \times \overline{H}_i^{cc} \times \overline{H}_{-i}^{cc}$, where $\overline{H}_i^{cc}, \overline{H}_{-i}^{cc} \subseteq H^{cc}$, is a *belief-closed subspace* of $\Theta \times H^{cc} \times H^{cc}$ if for every $i \in I$ and every $\mu_i \in \overline{H}_i^{cc}$ we have $g^{cc}(\mu_i) \left(\Theta \times \overline{H}_{-i}^{cc} \right) = 1$. It can be shown that $h_i(T_i), h_{-i}(T_{-i}) \subseteq H^{cc}$ indeed form a belief-closed subspace of $\Theta \times H^{cc} \times H^{cc}$ (e.g., by Proposition 3 of Battigalli and Siniscalchi, 1999).

Assumption 1 *Each T_i is compact and each h_i is one-to-one.*

By construction, h_i is continuous (again, by Proposition 3 of Battigalli and Siniscalchi, 1999). Assumption 1 requires it to be one-to-one. This is a *non-redundancy* assumption: $h_i(t_i) = h_i(t'_i)$ implies $t_i = t'_i$. If Assumption 1 holds, then each h_i is a continuous one-to-one map from a compact space to the Hausdorff space H (Polish spaces are metrizable, hence Hausdorff). It follows that: (i) each h_i is a closed map; (ii) each h_i is an *embedding*, i.e., $h_i : T_i \rightarrow h_i(T_i)$ is a homeomorphism (e.g., by Theorem 2.36 in Aliprantis and Border, 2006). Thus, $\Theta \times T_i \times T_{-i}$ is homeomorphic to a belief-closed subspace of $\Theta \times H^{cc} \times H^{cc}$.

3 Quotient type spaces

We can now define the ℓ_i -quotient type spaces. For $\ell_i \geq 1$, let $h_i^{1,\ell_i} : T_i \rightarrow \prod_{m=1}^{\ell_i} \Delta(X^{m-1})$ be given by $h_i^{1,\ell_i}(t_i) = (h_i^1(t_i), \dots, h_i^{\ell_i}(t_i))$. In words, $h_i^{1,\ell_i}(t_i)$ is the partial belief hierarchy of type t_i , from $h_i^1(t_i)$ to $h_i^{\ell_i}(t_i)$. It is these partial hierarchies that determine the elements of our quotient spaces.

Definition 2 *The ℓ_i -quotient type space ($i \in I, \ell_i \in \{1, 2, \dots\}$) is the space*

$$T_i^{\ell_i} = \left\{ [t_i] \subseteq T_i : t'_i \in [t_i] \iff h_i^{1,\ell_i}(t'_i) = h_i^{1,\ell_i}(t_i) \right\} \quad (2)$$

Let $\zeta^{\ell_i} : T_i \rightarrow T_i^{\ell_i}$, $\zeta^{\ell_i}(t_i) = [t_i] = \left\{ t'_i \in T_i : h_i^{1,\ell_i}(t'_i) = h_i^{1,\ell_i}(t_i) \right\}$ be the corresponding *quotient map* (e.g., see Munkres, 2000). Then we can also write the ℓ_i -quotient space as $T_i^{\ell_i} = \left\{ \zeta^{\ell_i}(t_i) : t_i \in T_i \right\}$. We endow $T_i^{\ell_i}$ with the quotient topology induced by ζ^{ℓ_i} : $E \subseteq T_i^{\ell_i}$ is open in $T_i^{\ell_i}$ if and only if $[\zeta^{\ell_i}]^{-1}(E)$ is open in T_i . An element of $T_i^{\ell_i}$ is denoted by $t_i^{\ell_i}$, which we refer to as an ℓ_i -*type* and which is an equivalence class of types $t_i \in T_i$.

In many applications, one can find $\ell_i < \infty$ such that $T_i^{\ell_i} = \left\{ \{t_i\} : t_i \in T_i \right\}$. That is, one can find a finite reasoning level ℓ_i such that a level- ℓ_i player is unboundedly rational (e.g., in the game of Section 4.1.1). Sometimes, however, such a finite ℓ_i does not exist (e.g., in the game of Section 3.1.1). It is therefore notationally convenient to identify the unbounded-rationality benchmark with $\ell_i = \infty$ and let T_i^∞ be the quotient space induced by the quotient map $\zeta^\infty(t_i) = \{t'_i \in T_i : h_i(t'_i) = h_i(t_i)\}$. For every $i \in I$, we write $L_i = \left\{ \ell_i \in \{1, 2, \dots\} : T_i^{\ell_i} \neq T_i^\infty \right\} \cup \{\infty\}$ for the set of all possible levels.

Proposition 1 *If Assumption 1 holds, then there exists a homeomorphism $\varphi^{\ell_i} : T_i^{\ell_i} \rightarrow h_i^{1,\ell_i}(T_i)$ ($i \in I, \ell_i \in L_i \setminus \{\infty\}$).*

Proof. Let $\varphi^{\ell_i}(t_i^{\ell_i}) = h_i^{1,\ell_i}([\zeta^{\ell_i}]^{-1}(t_i^{\ell_i}))$. Notice that φ^{ℓ_i} is one-to-one since by definition $\varphi^{\ell_i}(t_i^{\ell_i}) = \varphi^{\ell_i}(t_i^{\ell_i'})$ implies $t_i^{\ell_i} = t_i^{\ell_i'}$. It is continuous since ζ^{ℓ_i} is a quotient map and $h_i^{1,\ell_i} = \varphi^{\ell_i} \circ \zeta^{\ell_i}$ is continuous

(e.g., by Theorem 22.2 in Munkres, 2000). $T_i^{\ell_i}$ is compact, as it is the image under the continuous ζ^{ℓ_i} of the compact T_i (e.g., by Theorem 2.34 in Aliprantis and Border, 2006). Finally, $\prod_{m=1}^{\ell_i} \Delta(X^{m-1})$ is Hausdorff, so $\varphi^{\ell_i} : T_i^{\ell_i} \rightarrow \prod_{m=1}^{\ell_i} \Delta(X^{m-1})$ is a closed map and $\varphi^{\ell_i} : T_i^{\ell_i} \rightarrow \varphi^{\ell_i}(T_i^{\ell_i})$, where $\varphi^{\ell_i}(T_i^{\ell_i}) = h_i^{1, \ell_i}(T_i)$, is a homeomorphism. ■

Corollary 1 *If Assumption 1 holds, then $T_i^{\ell_i}$ is Polish ($i \in I, \ell_i \in L_i \setminus \{\infty\}$).*

Proof. A subset of a Polish space is Polish if and only if it is a G_δ set (e.g., by Corollary 3.5, Lemma 3.33, and Alexandroff's Lemma 3.34 in Aliprantis and Border, 2006). It then suffices to show that $h_i^{1, \ell_i}(T_i)$ is closed in $\prod_{m=1}^{\ell_i} \Delta(X^{m-1})$, as in metrizable spaces every closed set is G_δ (e.g., by Corollary 3.19 in Aliprantis and Border, 2006). This follows from $\varphi^{\ell_i} : T_i^{\ell_i} \rightarrow \prod_{m=1}^{\ell_i} \Delta(X^{m-1})$ being a closed map. ■

3.1 Equilibrium analysis

The defined quotient spaces are compatible with standard equilibrium analysis. Let the space A_i contain player i 's pure actions, which are denoted by a_i . A strategy for player i is a measurable map $\sigma_i : T_i \rightarrow A_i$. The set of all strategies is denoted by Σ_i .⁵

Definition 3 *A strategy $\sigma_i \in \Sigma_i$ is $T_i^{\ell_i}$ -measurable if there exists a measurable $\tilde{\sigma}_i^{\ell_i} : T_i^{\ell_i} \rightarrow A_i$ such that $\sigma_i = \tilde{\sigma}_i^{\ell_i} \circ \zeta^{\ell_i}$.⁶*

In particular, for a $T_i^{\ell_i}$ -measurable strategy σ_i we have that $\zeta^{\ell_i}(t_i) = \zeta^{\ell_i}(t'_i)$ implies $\sigma_i(t_i) = \sigma_i(t'_i)$ for all $t_i, t'_i \in T_i$. If player i is of level ℓ_i , then σ_i must be $T_i^{\ell_i}$ -measurable. This way, the corresponding distribution of actions can be derived from the distribution of ℓ_i -types, as if player i 's type space were $T_i^{\ell_i}$ instead of T_i . Player i 's payoff function is given by the measurable $u_i : \Theta \times A_i \times A_{-i} \rightarrow \mathbb{R}$. A Bayesian game is a tuple $\Gamma = \langle I, (A_i)_{i \in I}, \Theta, (T_i)_{i \in I}, (f_i)_{i \in I}, (u_i)_{i \in I} \rangle$.

Definition 4 *The strategy profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ is a Bayesian Nash ℓ -equilibrium ($\ell = (\ell_i, \ell_{-i}) \in L_i \times L_{-i}$) of Γ if each σ_i^* is $T_i^{\ell_i}$ -measurable and, for each $i \in I$ and each $t_i \in T_i$,*

$$\sigma_i^*(t_i) \in \arg \max_{a_i \in A_i} \mathbb{E} \left[u_i \left(\tilde{\theta}, a_i, \sigma_{-i}^*(\tilde{t}_{-i}) \right) \mid \tilde{t}_i \in \zeta^{\ell_i}(t_i) \right] \quad (3)$$

3.1.1 Example: The electronic mail game

Consider the game described in Rubinstein (1989), where $I = \{1, 2\}$ and $\Theta = \{\theta_a, \theta_b\}$ is given by

$$\theta_a = \begin{array}{c|cc} 1 \setminus 2 & a & b \\ \hline a & M, M & 1, -L \\ \hline b & -L, 1 & 0, 0 \end{array}, \quad \theta_b = \begin{array}{c|cc} 1 \setminus 2 & a & b \\ \hline a & 0, 0 & 1, -L \\ \hline b & -L, 1 & M, M \end{array}$$

⁵We leave out mixed strategies because when T_i is uncountable they involve measurability issues (e.g., see Aumann, 1964, and Milgrom and Weber, 1985) that are beyond the scope of our paper. Indeed, all our examples feature equilibria in pure strategies.

⁶Existence of such a $\tilde{\sigma}_i^{\ell_i}$ implies that σ_i is measurable with respect to the σ -algebra given by the sets $[\zeta^{\ell_i}]^{-1}(E)$ such that E is Borel in $T_i^{\ell_i}$.

with $1 < M < L$. The state of nature is θ_a with probability $\rho \in (0, 1)$. Only player 1 is informed about whether θ_a or θ_b is the realization of $\tilde{\theta}$. If $\tilde{\theta} = \theta_b$, then an automatic communication protocol has players sequentially sending e-mails to each other - starting from player 1 sending an e-mail to player 2 - that fail to be delivered with probability $\epsilon \in (0, 1)$. Let $T_1 = T_2 = \{0, 1, 2, \dots\}$ count the number of e-mails *sent* by a given player. By defining $n(t_1, t_2) = |\{t'_1 \in T_1 : t'_1 < t_1\}| + |\{t'_2 \in T_2 : t'_2 < t_2\}| + 1$, we can write the underlying common prior on $\Theta \times T_1 \times T_2$ as

$$f(\theta, t_1, t_2) = \rho \cdot \mathbf{1}_{\{\theta=\theta_a\}} + (1 - \rho) (1 - \epsilon)^{n(t_1, t_2)} \epsilon \cdot \mathbf{1}_{\{\theta=\theta_b\}} \quad (4)$$

for $t_1 - 1 \leq t_2 \leq t_1$ and 0 otherwise. Player $i \in I$ of level $\ell_i \in \{1, 2, \dots, \infty\}$ maximizes his ex-ante payoff $\mathbb{E} \left[u_i \left(\tilde{\theta}, \sigma_i(\tilde{t}_i), \sigma_{-i}(\tilde{t}_{-i}) \right) \right]$ by choosing $\sigma_i : T_i \rightarrow A_i$ subject to σ_i being $T_i^{\ell_i}$ -measurable.⁷ In contrast with the unique Bayesian Nash (∞, ∞) -equilibrium, action b may be chosen in a boundedly rational ℓ -equilibrium.

Proposition 2 *For any $k \in \{1, 2, \dots\}$, there exists $\ell = (\ell_1, \ell_2) \in \{1, 2, \dots, \infty\}^2$ and an associated Bayesian Nash ℓ -equilibrium of the electronic mail game in which players that have sent k or more messages play action b .*

Proof. Let $\ell_1 < \ell_2 \leq \infty$.⁸ Notice that $T_1^1 = T_2^1 = \{\{0\}, \{1, 2, \dots\}\}$ because: (i) only the type that has sent no e-mails believes $\tilde{\theta} = \theta_a$ with positive probability; (ii) all types that have sent at least one e-mail believe $\tilde{\theta} = \theta_b$. Suppose player 1 has sent exactly $\tilde{t}_1 = t_1 \geq 2$ messages. In terms of lower bounds on the other player's type, player 1 then believes $\tilde{t}_2 \geq t_1 - 1$, that player 2 believes $\tilde{t}_1 \geq t_1 - 1$, that 2 believes 1 believes $\tilde{t}_2 \geq t_1 - 2$, that 2 believes 1 believes 2 believes $\tilde{t}_1 \geq t_1 - 2$, and so forth. Thus, by letting $e(\ell_1) = \max\{t_1 \leq \ell_1 : t_1 \equiv 0 \pmod{2}\}$ we have $T_1^{\ell_1} = \{\{0\}, \{1\}, \dots, \{\frac{1}{2}e(\ell_1)\}, \{\frac{1}{2}e(\ell_1) + 1, \dots\}\}$. Suppose $\tilde{t}_2 = t_2 \geq 1$. As lower bounds, player 2 then believes $\tilde{t}_1 \geq t_2$, that 1 believes $\tilde{t}_2 \geq t_2 - 1$, that 1 believes 2 believes $\tilde{t}_1 \geq t_2 - 1$, that 1 believes 2 believes 1 believes $\tilde{t}_2 \geq t_2 - 2$, and so on. Thus, by letting $o(\ell_2) = \max\{t_2 \leq \ell_2 : t_2 \equiv 1 \pmod{2}\} + 1$ we have $T_2^{\ell_2} = \{\{0\}, \{1\}, \dots, \{\frac{1}{2}o(\ell_2) - 1\}, \{\frac{1}{2}o(\ell_2), \dots\}\}$. It follows that, with $\ell_1 < \ell_2$, $T_2^{\ell_2}$ is no coarser than $T_1^{\ell_1}$.

By iterated elimination of strictly dominated strategies at an interim level, $\sigma_1(0) = a$, which implies $\sigma_2(0) = a$, which implies $\sigma_1(1) = a, \dots$, which implies $\sigma_1(\frac{1}{2}e(\ell_1)) = a$, which implies $\sigma_2(\frac{1}{2}e(\ell_1)) = a$. Now suppose $\sigma_1(\frac{1}{2}e(\ell_1) + 1) = \sigma_1(\frac{1}{2}e(\ell_1) + 2) = \dots = b$. Then $\sigma_2(\frac{1}{2}e(\ell_1) + 1) = \sigma_2(\frac{1}{2}e(\ell_1) + 2) = \dots = b$. There remains to check that 1's ex-ante payoff under this strategy is maximal. One shows that the equilibrium is sustained under the condition $\epsilon \leq \epsilon^* = \frac{M-1}{M+L-1}$.

To prove the claim, let ℓ_1 be such that $\frac{1}{2}e(\ell_1) + 1 = k$. ■

This is consistent with experimental evidence by Camerer (2003): subjects tend to play action b already after few messages exchanged.

3.2 Countable Θ -based Type Spaces

We say that a Θ -based type space $\langle \Theta, (T_i)_{i \in I}, (f_i)_{i \in I} \rangle$ is *countable* if the following Assumption 2 is verified.

⁷If T_i is countable, then the ex-ante perspective is equivalent to the interim perspective of Definition 4. If instead T_i is uncountable, a strategy σ_i that maximizes the ex-ante payoff need not maximize the interim payoff on a subset of types t_i of null measure.

⁸Equilibria for $\ell_2 < \infty$ and $\ell_2 \leq \ell_1 \leq \infty$ can be found with analogous computations.

Assumption 2 $\Theta \times T_i \times T_{-i}$ is countable.

In this case, $T_i^{\ell_i}$ and $T_{-i}^{\ell_{-i}}$ can be employed to construct another Θ -based type space. Let $f_i^{\ell_i} : T_i^{\ell_i} \rightarrow \Delta(\Theta \times T_{-i})$ be given by $f_i^{\ell_i}(t_i^{\ell_i})(\cdot) = f(\cdot \mid [\zeta^{\ell_i}]^{-1}(t_i^{\ell_i}))$. The product of the identity map on Θ and the quotient map $\zeta^{\ell_{-i}}$ defines a map $\zeta_{\Theta}^{\ell_i} = id_{\Theta} \times \zeta^{\ell_{-i}} : \Theta \times T_{-i} \rightarrow \Theta \times T_{-i}^{\ell_{-i}}$, which in turn induces a $\widehat{\zeta}_{\Theta}^{\ell_i} : \Delta(\Theta \times T_{-i}) \rightarrow \Delta(\Theta \times T_{-i}^{\ell_{-i}})$ given by $\widehat{\zeta}_{\Theta}^{\ell_i}(\varrho)(E) = \varrho\left([\zeta_{\Theta}^{\ell_i}]^{-1}(E)\right)$ for $\varrho \in \Delta(\Theta \times T_{-i})$ and $E \subseteq \Theta \times T_{-i}^{\ell_{-i}}$. Finally, let $f_i^{\ell_i, \ell_{-i}} = \widehat{\zeta}_{\Theta}^{\ell_i} \circ f_i^{\ell_i}$.

Proposition 3 If Assumption 2 holds, then $\left\langle \Theta, \left(T_i^{\ell_i}\right)_{i \in I}, \left(f_i^{\ell_i, \ell_{-i}}\right)_{i \in I} \right\rangle$ is a countable Θ -based type space ($\ell = (\ell_i, \ell_{-i}) \in L_i \times L_{-i}$).

Proof. $T_i^{\ell_i}$ is Polish for $i \in I$, since each of these ℓ_i -quotient type spaces are countable - possibly finite, as in the e-mail game of Section 3.1.1 - and can be shown to inherit the discrete topology from T_i (i.e., every subset of each $T_i^{\ell_i}$ is open). By the discrete topology, each $f_i^{\ell_i}$ is continuous. $\zeta_{\Theta}^{\ell_i}$ is a continuous map between metrizable spaces, thus $\widehat{\zeta}_{\Theta}^{\ell_i}$ is continuous (e.g., by Theorem 15.14 in Aliprantis and Border, 2006). It follows that each composite map $f_i^{\ell_i, \ell_{-i}}$ is continuous. Countability follows from $|\Theta \times T_i^{\ell_i} \times T_{-i}^{\ell_{-i}}| \leq |\Theta \times T_i \times T_{-i}|$. ■

Let $\Gamma = \langle I, (A_i)_{i \in I}, \Theta, (T_i)_{i \in I}, (f_i)_{i \in I}, (u_i)_{i \in I} \rangle$ and $\Gamma^{\ell} = \left\langle I, (A_i)_{i \in I}, \Theta, \left(T_i^{\ell_i}\right)_{i \in I}, \left(f_i^{\ell_i, \ell_{-i}}\right)_{i \in I}, (u_i)_{i \in I} \right\rangle$.

Proposition 4 If Assumption 2 holds, then the strategy profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ is a Bayesian Nash ℓ -equilibrium of Γ if and only if the induced strategy profile $\widehat{\sigma}^{\ell*} = (\widehat{\sigma}_i^{\ell_i*}, \widehat{\sigma}_{-i}^{\ell_{-i}*})$ is a Bayesian Nash (∞, ∞) -equilibrium of Γ^{ℓ} ($\ell = (\ell_i, \ell_{-i}) \in L_i \times L_{-i}$).

Proof. First, note that $\widehat{\sigma}_i^{\ell_i*}(t_i^{\ell_i}) = \sigma_i^*\left([\zeta^{\ell_i}]^{-1}(t_i^{\ell_i})\right) = \sigma_i^*(t_i)$ for all $t_i \in [\zeta^{\ell_i}]^{-1}(t_i^{\ell_i})$ because a $T_i^{\ell_i}$ -measurable strategy is constant over equivalence classes. Take any such t_i .

$$\begin{aligned} \mathbb{E}\left[u_i\left(\widetilde{\theta}, a_i, \sigma_{-i}^*(\widetilde{t}_{-i})\right) \mid \widetilde{t}_i \in \zeta^{\ell_i}(t_i)\right] &= \sum_{(\theta, t_{-i}) \in \Theta \times T_{-i}} u_i(\theta, a_i, \sigma_{-i}^*(t_{-i})) f(\theta, t_{-i} \mid [\zeta^{\ell_i}]^{-1}(t_i^{\ell_i})) \\ &= \sum_{(\theta, t_{-i}^{\ell_{-i}}) \in \Theta \times T_{-i}^{\ell_{-i}}} u_i(\theta, a_i, \widehat{\sigma}_{-i}^{\ell_{-i}*}(t_{-i}^{\ell_{-i}})) f\left([\zeta_{\Theta}^{\ell_i}]^{-1}(\theta, t_{-i}^{\ell_{-i}}) \mid [\zeta^{\ell_i}]^{-1}(t_i^{\ell_i})\right) \\ &= \mathbb{E}\left[u_i\left(\widetilde{\theta}, a_i, \sigma_{-i}^*(\widetilde{t}_{-i}^{\ell_{-i}})\right) \mid \widetilde{t}_i \in \{t_i^{\ell_i}\}\right] \end{aligned} \quad (5)$$

Therefore, an action a_i satisfies (3) for game Γ with players' levels (ℓ_i, ℓ_{-i}) if and only if it satisfies (3) for Γ^{ℓ} with players' levels (∞, ∞) . ■

4 Uncertainty about players' levels

More realistically, players' levels $\ell = (\ell_i, \ell_{-i})$ may be private information. To address this, we treat players of different levels as different players. Such new players have the same information as the original players from whom they are derived (see Definition 5 in Section 4.1 below). As before, let $L_i =$

$\{\ell_i \in \{1, 2, \dots\} : T_i^{\ell_i} \neq T_i^\infty\} \cup \{\infty\}$, $i \in I$. Let $L = L_i \times L_{-i}$ and assume there is a *common prior* $\lambda \in \Delta(L)$ that is common knowledge between the players. Let $\bar{L}_i = \{\ell_i \in L_i : \ell_i \in \text{supp}(\text{marg}_{L_i} \lambda)\}$. The *expanded set of players* is $I^e = \bar{L}_i \cup \bar{L}_{-i}$. $|I^e| = \sum_{i \in I} |\text{supp}(\text{marg}_{L_i}(\lambda))|$ is the total number of players. The set of players other than ℓ_i is $I_{-\ell_i}^e = \bar{L}_{-i}$ so that, in particular, $I_{-\ell_i}^e \cap \bar{L}_i = \emptyset$ for any $\ell_i \in I^e$. That goes to say that player ℓ_i does not face any ℓ'_i , but only some ℓ_{-i} .

4.1 Equilibrium analysis

Let player ℓ_i 's set of *pure actions* be $A_i^{\ell_i} = A_i$, with typical element $a_i^{\ell_i}$. A *strategy* for player ℓ_i is a $T_i^{\ell_i}$ -measurable map $\sigma_i^{\ell_i} \in \Sigma_i$. Player ℓ_i 's posterior on the other player's level ℓ_{-i} is given by the map $\lambda_i : \bar{L}_i \rightarrow \Delta(L_{-i})$, such that $\lambda_i(\ell_i)(\cdot) = \lambda(\cdot | \ell_i)$. Player ℓ_i 's *payoff function* is $u_i^{\ell_i} : \Theta \times A_i \times \left(\prod_{\ell_{-i} \in \bar{L}_{-i}} A_{-i}\right) \rightarrow \mathbb{R}$, where

$$u_i^{\ell_i} \left(\theta, a_i^{\ell_i}, \left(a_{-i}^{\ell_{-i}} \right)_{\ell_{-i} \in \bar{L}_{-i}} \right) = \sum_{\ell_{-i} \in \bar{L}_{-i}} u_i \left(\theta, a_i^{\ell_i}, a_{-i}^{\ell_{-i}} \right) \lambda_i(\ell_i)(\ell_{-i}) \quad (6)$$

In words, player $i \in I$ of level ℓ_i sees player $-i$'s ($-i \in I$) action as a random \tilde{a}_{-i} with realizations $\left(a_{-i}^{\ell_{-i}} \right)_{\ell_{-i} \in \bar{L}_{-i}}$ that occur with probabilities $(\lambda(\ell_i)(\ell_{-i}))_{\ell_{-i} \in \bar{L}_{-i}}$. The game without uncertainty of Section 3.1.1 is a special case where λ is degenerate at some $\ell = (\ell_i, \ell_{-i})$.

It is essential that $\lambda_i(\ell_i)(\cdot)$ only depends on ℓ_i . If $\lambda_i(\ell_i)(\cdot)$ were to depend on either t_i or t_{-i} , then $u_i^{\ell_i}$ would depend on them as well, so its argument would not be a point in the product space $\Theta \times A_i \times \left(\prod_{\ell_{-i} \in \bar{L}_{-i}} A_{-i}\right)$. If $\lambda_i(\ell_i)(\cdot)$ depended on θ , then player i could infer the state of nature by knowing his level. He would therefore be conditioning his expected payoff both on his type belonging to the respective equivalence class *and* on his level being ℓ_i . If this were true, then $-i$ would need to form beliefs about i 's level of sophistication, but this is precisely what we want to avoid.

The function $u_i^{\ell_i}$ has to be measurable. A sufficient condition for this is its continuity - which is satisfied in our applications.

Definition 5 *The strategy profile $\sigma^* = \left(\left(\sigma_i^{\ell_i^*} \right)_{\ell_i \in \bar{L}_i}, \left(\sigma_{-i}^{\ell_{-i}^*} \right)_{\ell_{-i} \in \bar{L}_{-i}} \right)$ is a Bayesian Nash λ -equilibrium ($\lambda \in \Delta(L)$) of Γ if each $\sigma_i^{\ell_i^*}$ is $T_i^{\ell_i}$ -measurable and, for each $\ell_i \in I^e$ and each $t_i \in T_i$,*

$$\sigma_i^{\ell_i^*}(t_i) \in \arg \max_{a_i^{\ell_i} \in A_i} \mathbb{E} \left[u_i \left(\tilde{\theta}, a_i^{\ell_i}, \left(\tilde{a}_{-i}^{\ell_{-i}} \right)_{\ell_{-i} \in \bar{L}_{-i}} \right) \mid \tilde{t}_i \in \zeta^{\ell_i}(t_i) \right] \quad (7)$$

4.1.1 Example: A game with private and public information

Consider a version of the game described in Morris and Shin (2002), where $I = \{1, 2\}$ and the state of nature θ follows an improper uniform distribution on $\Theta = \mathbb{R}$. Both players receive a public signal $y = \theta + \epsilon_y$, where $\epsilon_y \sim U[-c, c]$, $c > 0$. In addition, each of them privately observes a signal $x_i = \theta + \epsilon_i$, where $\epsilon_i \sim U[-c, c]$, $i \in I$. The random variables θ , ϵ_y , ϵ_1 , and ϵ_2 are all independent of each other. Payoffs are given by

$$u_i(\theta, a_i, a_{-i}) = -(1 - \nu)(a_i - \theta)^2 - \nu(a_i - a_{-i})^2 \quad (8)$$

where $\nu \in (0, 1)$ is a parameter. We have $T_i = \{t_i = (x_i, y) : (x_i, y) \in \mathbb{R}^2\}$ and $T_i^{\ell_i} = \{\zeta_i^{\ell_i}(t_i) : t_i \in T_i\}$. In particular: (i) ζ_i^1 is such that $\zeta_i^1(t_i) = \zeta_i^1(t'_i)$ if and only if $[x_i = x'_i \text{ and } y = y']$ or $[x_i = y' \text{ and } y = x'_i]$;

(ii) $\zeta_i^2 = \zeta_i^3 = \dots = \zeta_i^\infty$.⁹ To interpret, a player of level 1 does not distinguish private and public signal. Each player i can be either of level 1 or ∞ . The common prior on levels $\lambda \in \Delta(\{1, \infty\}^2)$ is given by $\lambda(1, 1) = p^2$, $\lambda(1, \infty) = \lambda(\infty, 1) = p(1-p)$, and $\lambda(\infty, \infty) = (1-p)^2$. The payoff function for player i of level $\ell_i \in \{1, \infty\}$ is

$$u_i^{\ell_i}(\theta, a_i^{\ell_i}, a_{-i}^1, a_{-i}^\infty) = -(1-\nu)(a_i^{\ell_i} - \theta)^2 - \nu p (a_i^{\ell_i} - a_{-i}^1)^2 - \nu(1-p)(a_i^{\ell_i} - a_{-i}^\infty)^2 \quad (9)$$

We look for symmetric linear equilibria of the form $a_i^\ell = \kappa_\ell x_i + (1-\kappa_\ell)y$, $\kappa_\ell \in (0, 1)$, $\ell \in \{1, \infty\}$, $i \in I$.

Proposition 5 *For each $p \in (0, 1)$ there exists a Bayesian Nash λ -equilibrium such that the mean weight put on the private signal (i.e., $p\kappa_1 + (1-p)\kappa_\infty$) is larger than when there is common certainty that both players are unboundedly rational (i.e., when $p = 0$).*

Proof. Let $\underline{x}_i = \min\{x_i, y\}$ and $\bar{x}_i = \max\{x_i, y\}$. The first-order condition for a level- ℓ_i player i is $a_i^{\ell_i} = (1-\nu)\mathbb{E}_i^{\ell_i}\theta + \nu p \mathbb{E}_i^{\ell_i}a_{-i}^1 + \nu(1-p)\mathbb{E}_i^{\ell_i}a_{-i}^\infty$, where for any random variable z , $\mathbb{E}_i^{\ell_i}z = \mathbb{E}\left[z \mid \zeta_i^{\ell_i}(t_i)\right]$. Notice that $\mathbb{E}_i^1 z = \mathbb{E}\left[z \mid \underline{x}_i, \bar{x}_i\right]$ and $\mathbb{E}_i^\infty z = \mathbb{E}[z \mid x_i, y]$. If $\kappa_1 = \frac{1}{2}$, then $a_i^1 = (1-\nu)\mathbb{E}_i^1\theta + \nu p \mathbb{E}_i^1\left[\frac{1}{2}x_{-i} + \frac{1}{2}y\right] + \nu(1-p)\mathbb{E}_i^1[\kappa_\infty x_{-i} + (1-\kappa_\infty)y]$. From $\mathbb{P}[\underline{x}_i = y] = 1 - \mathbb{P}[\bar{x}_i = y] = \frac{1}{2}$ and $\mathbb{E}_i^1 x_{-i} = \mathbb{E}_i^1\theta = \frac{1}{2}\underline{x}_i + \frac{1}{2}\bar{x}_i$, we have that for any $\alpha \in (0, 1)$

$$\mathbb{E}_i^1[\alpha y + (1-\alpha)x_{-i}] = \frac{1}{2}\left[\alpha\underline{x}_i + (1-\alpha)\frac{\underline{x}_i + \bar{x}_i}{2}\right] + \frac{1}{2}\left[\alpha\bar{x}_i + (1-\alpha)\frac{\underline{x}_i + \bar{x}_i}{2}\right] = \frac{1}{2}\underline{x}_i + \frac{1}{2}\bar{x}_i \quad (10)$$

The best-response to $\kappa_1 = \frac{1}{2}$ and any $\kappa_\infty \in (0, 1)$ is therefore $\kappa_1 = \frac{1}{2}$. As to κ_∞ , we have $a_i^\infty = (1-\nu)\mathbb{E}_i^\infty\theta + \nu p \mathbb{E}_i^\infty\left[\frac{1}{2}x_{-i} + \frac{1}{2}y\right] + \nu(1-p)\mathbb{E}_i^\infty[\kappa_\infty x_{-i} + (1-\kappa_\infty)y]$, from which $\kappa_\infty = \frac{2-2\nu+\nu p}{2(2-\nu+\nu p)}$.

Since $p \in (0, 1)$, it is $p\kappa_1 + (1-p)\kappa_\infty \in \left(\frac{1-\nu}{2-\nu}, \frac{1}{2}\right)$, which is larger than the standard (∞, ∞) -equilibrium weight $\frac{1-\nu}{2-\nu}$. ■

These results are consistent with experimental evidence by Cornand and Heinemann (2009). Also, notice that $\lim_{p \rightarrow 1} \kappa_\infty = \frac{2-\nu}{4}$, which is the best-response of a level-2 player to a level-1 player who equally randomizes between signals in their *level- k /cognitive hierarchy* model (Camerer, Ho, and Chong, 2004, Nagel, 1995, Stahl and Wilson, 1994, 1995).

5 Related Literature

Kets (2010) characterizes *extended type structures*, which feature Harsanyi type structures as special cases. As in the Harsanyi framework, each type in an extended type structure is associated with beliefs about the state of nature and the other player's types, yet the beliefs of different types are allowed to be defined on different σ -algebras. These σ -algebras reflect types' coarseness of perception of their opponent's high-order beliefs. The author shows that standard Harsanyi type structures are characterized by common certainty that each player has an infinite depth of reasoning. The difference with our model is twofold: (i) we do not impose cognitive limitations on player i 's capacity to perceive player $-i$'s high-order beliefs, but rather on

⁹If $y \neq y'$, then types (x_i, y) and (x'_i, y') such that $[x_i = y' \text{ and } y = x'_i]$ have different second-order beliefs.

i 's capacity to perceive *his own* high-order beliefs; (ii) we place emphasis on being in a position to retain the standard game-theoretical tools for equilibrium analysis of incomplete-information games. By contrast, in Kets' (2010) analysis of the e-mail game, player i with depth of reasoning ℓ_i plays strategies that are rationalizable for him given his coarse perception of player $-i$'s high-order beliefs. This amounts to each player i perceiving a different game of incomplete information in which: (i) his set of types is the same as in the original game; (ii) player $-i$'s set of types is changed according to i 's own perception of player $-i$'s high-order beliefs; (iii) player i plays strategies that are rationalizable in the modified game.¹⁰

Jehiel (2005) proposes a solution concept for multi-stage games with perfect information, the *analogy-based expectation equilibrium*. Jehiel and Koessler (2008) make use of analogy-based expectations in static two-player games of incomplete information. The basic idea behind their analogy-based expectation equilibrium is that player i plays best-responses to player $-i$'s average strategy within analogy classes, which are bundles of states of the world. The authors interpret the analogy-based expectation equilibrium as "the limiting outcome of a learning process involving populations of players $i = 1, 2$, who would get a coarse feedback about the past behavior of players in population $j \neq i$ and no feedback on their own past performance until they exit the system" (p. 534). They then apply their theory to the e-mail game and find the same threshold for the probability with which each message gets lost as we do. However, this should be seen as a mere coincidence, as generally analogy-based expectation equilibria need not coincide with Bayesian Nash equilibria of a game with modified information partitions.

Dulleck (2007) resolves the e-mail game paradox by assuming that a player loses track of the number of e-mails received in a given interval and that the corresponding information structure is common knowledge between players. While this is close in spirit to our approach, the analysis falls short of addressing formally where the confusion about the e-mail count originates, and the model is specific to the e-mail game, leaving open how it can be adapted to other situations of strategic interaction.

Finally, Strzalecki (2010) develops a level- k /cognitive hierarchy model to show that coordination on the Pareto-efficient outcome is possible in the e-mail game, provided sophisticated players put enough weight on the other being less sophisticated. The model builds on a non-equilibrium approach according to which players are boundedly rational in terms of reasoning about actions (i.e., what i thinks $-i$ thinks ... will play) but unboundedly rational in terms of epistemic reasoning (i.e., what i thinks $-i$ thinks ... about the state of nature). On the contrary, we do not depart from equilibrium analysis and our quotient type spaces have clear-cut interpretations of players' ability – or inability – to precisely perceive high-order beliefs about states of nature.

References

- [1] ALIPRANTIS, C. D. AND K. C. BORDER (2006), *Infinite Dimensional Analysis: A hitchhiker's guide, 3rd Edition*. Berlin: Springer.

¹⁰In Kets' (2010) paper, iterated elimination of strictly dominated strategies at an interim level also takes into account uncertainty about players' depth of reasoning. This is to say that since an action for type t_i may or may not be a best-response depending on how sophisticated i is, an action for type t_{-i} may or may not be eliminated depending on how sophisticated $-i$ thinks i is.

- [2] AUMANN, R. J. (1964), “Mixed and behavior strategies in infinite extensive games.” In Dresher, M., L. S. Shapley, and A. W. Tucker (Eds.), *Advances in Game Theory, Annals of Mathematics Studies, 52*. Princeton, NJ: Princeton University Press.
- [3] AUMANN, R. J., AND J. H. DREZE (2008), “Rational expectations in games.” *American Economic Review*, 98, 72-86.
- [4] BATTIGALLI, P. AND M. SINISCALCHI (1999), “Hierarchies of conditional beliefs and interactive epistemology in dynamic games.” *Journal of Economic Theory*, 88, 188-230.
- [5] BATTIGALLI, P. AND M. SINISCALCHI (2003), “Rationalization and incomplete information,” *B. E. Journal of Theoretical Economics*, 3, article 3.
- [6] BRANDENBURGER, A. AND E. DEKEL (1993), “Hierarchies of beliefs and common knowledge.” *Journal of Economic Theory*, 59, 189-198.
- [7] CAMERER, C. F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- [8] CAMERER, C. F., T.-H. HO, AND J.-K. CHONG (2004), “A cognitive hierarchy model of games.” *Quarterly Journal of Economics*, 119, 861-898.
- [9] CARLSSON, H. AND E. VAN DAMME (1993), “Global games and equilibrium selection.” *Econometrica*, 61, 989-1018.
- [10] CORNAND, C. AND F. HEINEMANN (2009), “Measuring agents’ reaction to private and public information in games with strategic complementarities.” CESifo Working Paper 2947.
- [11] DEKEL, E., D. FUDENBERG, AND S. MORRIS (2007), “Interim correlated rationalizability.” *Theoretical Economics*, 2, 15-40.
- [12] DULLECK, U. (2007), “The e-mail game revisited - modeling rough inductive reasoning.” *International Game Theory Review*, 9, 323-339.
- [13] HALPERN, J. Y. (1986), “Reasoning about knowledge: An overview.” In Halpern, J. Y. (Eds.), *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning about Knowledge*. San Francisco, CA: Morgan Kaufmann Publishers Incorporated.
- [14] HARSANYI, J. C. (1967–1968), “Games with incomplete information played by “Bayesian” players. I–III.” *Management Science*, 14, 159–182, 320–334, 486–502.
- [15] HEIFETZ, A. (1995), “The Bayesian formulation of incomplete information: The non-compact case.” *International Journal of Game Theory*, 21, 329-338.
- [16] JEHIEL, P. (2005), “Analogy-based expectation equilibrium.” *Journal of Economic Theory*, 123, 81-104.
- [17] JEHIEL, P. AND F. KOESSLER (2008), “Revisiting games of incomplete information with analogy-based expectations.” *Games and Economic Behavior*, 62, 533-557.
- [18] KETS, W. (2010), “Bounded reasoning and higher-order uncertainty.” Job market paper.

- [19] MERTENS, J.-F. AND S. ZAMIR (1985), "Formulation of Bayesian analysis for games with incomplete information." *International Journal of Game Theory*, 14, 1-29.
- [20] MILGROM, P. R. AND R. J. WEBER (1985), "Distributional strategies for games with incomplete information." *Mathematics of Operations Research*, 10, 619-632.
- [21] MORRIS, S. AND Y. S. SHIN (2002), "Social value of public information." *American Economic Review*, 92, 1521-1534.
- [22] MUNKRES, J. R. (2000), *Topology, 2nd edition*. Upper Saddle River, NJ: Prentice Hall.
- [23] NAGEL, R. (1995), "Unraveling in guessing games: An experimental study." *American Economic Review*, 85, 1313-1326.
- [24] RUBINSTEIN, A. (1989), "The electronic mail game: Strategic behavior under 'almost common knowledge'." *American Economic Review*, 79, 385-391.
- [25] SINISCALCHI, M. (2008), "Epistemic game theory: Beliefs and types." In Durlauf, S. N. and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics, 2nd Edition*. Palgrave Macmillan.
- [26] STAHL, D. O. AND P. W. WILSON (1994), "Experimental evidence on players' models of other players." *Journal of Economic Behavior and Organization*, 25, 309-327.
- [27] STAHL, D. O. AND P. W. WILSON (1995), "On players' models of other players: Theory and experimental evidence." *Games and Economic Behavior*, 10, 218-254.
- [28] STRZALECKI, T. (2010), "Depth of reasoning and higher order beliefs." Working paper.
- [29] WEINSTEIN, J. AND M. YILDIZ (2007a), "A structure theorem for rationalizability with application to robust predictions of refinements." *Econometrica*, 75, 365-400.
- [30] WEINSTEIN, J. AND M. YILDIZ (2007b), "Impact of higher-order uncertainty." *Games and Economic Behavior*, 60, 200-212.