Title
An Explicit Approach to Modeling Finite-Order Type Spaces and Applications

Authors
Qin, Cheng-Zhong
Yang, Chun-Lei

Publication Date
2009-12-03

An Explicit Approach to Modeling Finite-Order Type Spaces and Applications *

Cheng-Zhong Qin and Chun-Lei Yang†

December 3, 2009

Abstract

Every abstract type of a belief-closed type space corresponds to an infinite belief hierarchy. But only finite order of beliefs is necessary for most applications. As we demonstrate, many important insights from recent development in the theory of Bayesian games with higher-order uncertainty involve belief hierarchies of order 2.

We start with characterizing order-2 "consistent priors" and show that they form a convex set and contain the convex hull of both the *naive* and complete-information type spaces. We establish conditions for private-value heterogeneous naive priors to be embedded in order-2 consistent priors, so as to retro-fit the Harsanyi doctrine of having nature generate all fundamental uncertainties in a game at the very beginning.

We then extend the notion of consistent priors to arbitrary finite order k. We define an abstract belief-closed type space to be of order k if it can be mapped via a type morphism into the "canonical representation" of an order-k consistent prior. We show that order-k type spaces are those in which any two types of each player must be either identical implying one of them is redundant or separable by their order-(k-1) belief hierarchies. Finite type spaces are always of finite orders. We consider "finite-order projections" of a type space and show that they are finite-order type spaces themselves. The condition of global stability

*We thank Yi-Chun Chen, Kim-Sau Chung, Qingmin Liu, Dov Monderer, and seminar participants at UCR for helpful discussions, comments, and references during the development of this paper.

†Contact: Cheng-Zhong Qin, Department of Economics, University of California at Santa Barbara; qin@econ.ucsb.edu; Chun-Lei Yang, Research Center for Humanity and Social Sciences, Academia Sinica, Taipei, Taiwan; cly@gate.sinica.edu.tw

under uncertainty ensures the convergence of the Bayesian-Nash equilibria with the projection type spaces to those with the original type space.

By defining a total variation norm based on finite-order projections, we generalize Kajii and Morris's (1997) idea of equilibrium robustness to Bayesian games. We then establish the robustness of Bayesian-Nash equilibria that generalizes the robustness results of Monderer and Samet (1989) for complete-information games.

We apply our framework of finite-order type spaces or consistent priors to review several important models in the literature and illustrate some new insights.

KEYWORDS: Bayesian game, consistent prior, embedding, finite-order type space, heterogeneous priors, projection type space, robustness, total variation norm. (*JEL* C72, D82)

1 Introduction

In most Bayesian game applications, information incompleteness is typically captured by a naive prior for simplicity, where each type of a player is identified with his privately informed *payoff type* who then forms his posterior belief by updating the prior. Despite its virtue of simplicity, naive type spaces may be too restrictive for certain applications.

For example, consider a sealed-bid, first-price auction of a high-tech patent between two potential buyers. In addition to the uncertainty about each other's reservation prices, one bidder might have reason to suspect that the other bidder may be able to acquire information about his reservation price with some positive probability, say, by hiring a super-talented computer hacker to steal his sensitive information. Naive type spaces cannot deal with such information settings.

Indeed, several recent studies show that deviations from naive type spaces may lead to qualitatively different conclusions in well-known models (see for example, Neeman 2004, Feinberg and Skrzypacz, 2005). There is, however, a great deal of additional complexity attached to type spaces beyond the naive ones. Though they may look simple in the implicit model *a la* Harsanyi (1967), without further classification, their correspondences in the Mertens and Zamir's (1985) explicit model of *universal* belief hierarchies can be too involved for intuitive comprehension. Furthermore, given a game of incomplete information, neither the implicit nor the explicit model specifies the maximum order of interactive beliefs that are strategically relevant. As we discuss later, however, characterizing type spaces in terms of belief orders that are strategically relevant turns out to be crucial in many applications, such as in the analysis of robustness of a Bayesian-Nash equilibrium (BNE).

In this paper, we show that the notions of a "consistent finite-order prior" and its "canonical type space" help to explicitly model type spaces with finite-order belief hierarchies, to establish a simple criterion for determining the order of an abstract belief-closed type space, and to analyze robustness of BNEs, among other issues.

We begin with the order-2 case. In the simplest version of an order-2 consistent prior, for each player nature draws according to the prior a payoff type and a first-order belief about the others' payoff types. Unlike with naive types, nature may draw multiple beliefs for a player's payoff type. The consistency of the prior requires that for each player, all his beliefs drawn by nature coincide with the posterior he can obtain by updating the prior from his private information.

We show that the set of consistent order-2 priors is convex, and it includes all naive type spaces. The complete-information type spaces turn out to be special cases of consistent order-2 priors. Consequently, type spaces with consistent order-2 common priors consist of the smallest class for robustness analysis of complete information games. Since this class contains the convex hull of the naive and the complete information type spaces, there is sufficient richness in structure to allow for new and often diametrical results in contrast to games with either naive or complete information type spaces. We defer further discussions to sections 1.1 and 4.

Models with heterogeneous priors suffer the disadvantage for not being explicit about nature's move to generate the interim private information, which is possible only with a common prior. For this reason, heterogeneous priors are often criticized for being too ad hoc, making trivial the explanation for how information differences result in behavior differences. As Myerson (2004) argues, there are strong modeling methodological reasons why the consistent common-prior assumption (CPA) has become standard in economic analysis. However, heterogeneous beliefs without support of an explicit common prior are often encountered in applications.[1] We show that with two players, heterogeneous naive priors with a common support can always be embedded in a consistent order-2 prior, in the sense that any posterior belief a player obtains by updating his naive prior can be obtained by updating the (embedding) consistent order-2 prior from his private information. For the n-player case, we characterize a suf-

[1] See Morris (1994) for related literature and Aumann (1998), Gul (1998), Morris (1995), among others for controversies about CPA.

ficient condition for players' naive heterogeneous priors to be embedded in an order-2 consistent prior.

This result reconciles the CPA controversy to the extent that players' heterogeneous priors could be "justified" by their induced posteriors being simultaneously generated from a common prior of exactly one-order higher. It could be imagined as if nature made its move according to the higher-order common prior, but systematically failed to inform, as private information, certain realizations. With this failure being common knowledge, the resulting information structure behaves as if every player is supposed to simply update his own heterogeneous prior to get the appropriate posteriors.

The model of private-value order-2 consistent-priors generalizes to models of private- or common-value consistent priors of arbitrary finite orders: An order-k consistent prior jointly draws a payoff-relevant parameter profile and a coherent belief hierarchy of order $(k-1)$ for each player, which matches the player's order-(k-1) belief hierarchy generated from updating the common prior using private information. Re-interpret such order-(k-1) belief hierarchies in the support of the prior as abstract types, the prior then induces a canonical common-prior belief-closed type space, such that the order-(k-1) projection of the associated universal belief hierarchies coincide with those of the abstract types. Private-value consistent finite-order priors are special common-value consistent finite-order priors, with additional information on private values.

We define a common-prior belief-closed type space to be of order-k if it can be mapped by some *type morphism* into a canonical representation of an order-k consistent prior. We show that a belief-closed type space is of order-k *if and only if* any two types of each player must be either identical, implying one of them is redundant, or separable by their order-(k-1) belief hierarchies. In this sense, an order-k consistent prior is equivalent to an order-k type space with a common prior. For this reason, we use consistent prior and common-prior type space interchangeably.

We consider several applications of our model of order-k consistent priors or order-k common-prior belief-closed type spaces. First, we consider "finite-order projections" of an arbitrary belief-closed type space. Such projections can be interpreted as *bounded* rational realizations of the original type space. With these projections, Weinstein and Yildiz's (2007b) *global stability under uncertainty* condition can be applied to ensure the convergence of the BNEs of the projection type spaces to those of the original type space. Second, we introduce an "order-k total-variation norm" for priors as an alternative to the standard product weak topology. Given any finite integer k, within the space of all order-k common-prior type spaces, the product topology is equivalent to both strategic topology (Dekel, Fudenberg, and Morris, 2006) and uniform weak topology (Chen et al, 2009). Applying the order-k total variation norm, we extend Kajii and Morris (1997) notion of robustness to general incomplete information games. We show that all BNEs of a Bayesian game with an order-k type space are "strong-k robust" against "order-k perturbations". This result extends the robustness result by Monderer and Samet (1989) for Nash equilibria of complete information games.

1.1 Literature Review of Order-2 Bayesian Games

Popular in applications of Bayesian games are the identification of the players' types with their payoff types and the assumption that it is common knowledge among the players that their types are jointly drawn by nature according to a common prior. These have an extreme implication: if a player's payoff is known, then so is his belief about the other players' types.

In recent literature, we have witnessed two strands of work that used information structure beyond the naive type spaces to arrive at some refreshing and striking results. One strand starts with allowing one payoff type to have more than one first-order beliefs, which results in less efficient outcomes. Indeed, Neeman (2004) demonstrates that the well-known *full-surplus-extraction* (FSE) results in mechanism design literature crucially hinge on the property that the agents' beliefs uniquely determine their payoff types (the BDP property). By implicitly constructing an order-2 consistent prior that easily violates the BDP property, he shows that FSE is not valid. Moreover, Heifetz and Neeman (2006) show that BDP is non-generic in general type spaces, in contrast to the claim in Cremer and McClean (1985, 1988) that BDP, and hence FSE, is generic in the class of naive type spaces.[2]

Bergemann and Valimaki (2008, section 3.2) has an example that demonstrates the strategic implication of introducing non-naive order-2 priors into First-Price Auction. Once again, allowing a payoff type to have two different order-1 beliefs leads to altered bidding strategies so that the outcome in the first-price auction is no longer efficient, and revenue equivalence to the Second-Price Auction is not valid.

Feinberg and Skrzypacz (2005) show that the *Coase conjecture* fails and delay occurs

[2]Given a convex family of priors containing at least one non-BDP prior, Heifetz and Neeman (2006) show that the subset of BDP priors is contained in a *proper face* of and *finitely shy* in the convex family of priors. Thus, the subset of BDP priors is small in both geometric and measure-theoretic senses. In relation to our discussion of Feinberg and Skrzypacz (2005), we will see that it is quite common that the same property is generic at the lower order space but is non-generic at the higher order ones.

with positive probability in bilateral trading, once the underlying information structure allows for uncertainty about seller's first-order beliefs and there is positive probability that the seller has a certainty belief. We demonstrate in details later in this paper that their prior turns out to be the convex combination of two naive ones; hence, it is a consistent order-2 prior. As an implication of the convexity of consistent priors and the insights from Heifetz and Neeman (2006), we confirm their conjecture that the failure of Coase conjecture is generic within the class of order-2 priors.

The other strand of the new literature is aimed at robustness of solutions around the condition of complete information. Bergemann and Morris (2005) show that, under some conditions, a social choice function being Bayesian implementable for all naive and complete information type spaces implies its implementability for all universal type spaces. The latter is equivalent to ex post implementability. As complete-information type spaces turn out to be order-2 consistent priors, their proof implies that for robust implementation, it is sufficient to have Bayesian implementability on order-2 consistent priors.

In a common-value environment, Carlsson and van Damme (1993) introduce a class of perturbations of the complete information to argue in favor of risk dominance in equilibrium selection. Morris and Shin (2000) modify their model with additional source of uncertainty in the form of a public signal and argued that rationalizability can be perturbation-specific.[3] As we argue in Section 4, the perturbed type spaces in these papers are all of order-2. It is also interesting to notice that Chung and Ely (2003) apply order-2 perturbations of a complete-information type space in their proof of the necessity of the monotonicity for undominated Nash equilibrium-closure implementability.

Another critical issue currently discussed in Ely and Peski (2006) and Liu (2009) concerns strategic relevance of redundancy in belief-closed type spaces. We show that our framework can be formally extended to modeling strategic redundancy using additional signal spaces as equilibrium correlation devices. In other words, even though our initial model has the feature of being free of redundancy, this feature can be easily bypassed to allow for richer strategic environment based on the essential principle of finite-order spaces.

The rest of the paper is organized as follows. In Section 2, we introduce the notion of consistent order-2 priors for a private-value payoff environment and discuss its

[3] Weinberg and Yildiz (2007) show in general universal space that ambiguity like this is the rule with product weak topology.

properties. Complete-information type spaces are shown to belong to this class. We also extend the definition to common-value environments. In Section 3, we extend our notion to general consistent finite-order priors, link them with finite-order universal type spaces, and extend results in Section 2. Moreover, we define projections of type spaces and finite-order total variation norms of priors. We apply them to the analysis of robustness of BNEs. In Section 4, we apply our framework of finite-order consistent priors to problems analyzed in the recent literature. Section 5 concludes.

2 Order-2 Consistent Priors

We adopt the following notation in what follows. Given a subset C of a topological space, $\Delta(C)$ denotes the set of probability measures on the σ-algebra of Borel sets of C. Given $P \in \Delta(C)$, supp$P$ denotes the support of P. Given an element $a \in C$, δ_a denotes the indicator function of the set $\{a\}$. For $C = X \times Y$, we use P_X or $\text{marg}_X P$ to denote the marginal probability measure of $P \in \Delta(C)$ over X and B_X or $\text{proj}_X B$ the projection of B onto X for $B \subseteq C$.

A game with incomplete information is a collection $\Gamma = \{(A_i, u_i, T_i, \pi_i)_{i\in N}, \Theta\}$, where $N = \{1, 2, \cdots, n\}$ is the player set, Θ is the space of payoff-relevant parameters, and for each $i \in N$, A_i is the action set of player i, $u_i : A \times \Theta \longrightarrow \Re$ his payoff function, T_i his type set, and $\pi_i : T_i \longrightarrow \Delta(\Theta \times T_{-i})$ is his belief function. Given $t_i \in T_i$, $\pi_i(t_i)$ specifies player i's belief about the others' types $t_{-i} \in T_{-i}$ and the payoff-relevant parameter $\theta \in \Theta$. A probability measure $P \in \Delta(\Theta \times T)$ is a *common prior* for Γ if $\pi_i(t_i) = P(\cdot|t_i)$ for all $t_i \in T_i$ and for all $i \in N$. With a common prior, the timing of the game can be perceived as nature moving first to generate all relevant information at the beginning. For technical tractability, we focus on priors with *finite support* for our theoretical analysis throughout the paper.

A strategy for player i in game Γ is a mapping $\sigma_i : T_i \longrightarrow \Delta(A_i)$. A strategy profile $\sigma = (\sigma_i)_i$ is a BNE for Γ if for all $i \in N, t_i \in T_i$, and $a_i \in A_i$,

$$E_{t_i}[u_i((\sigma_i(t_i), \sigma_{-i}(t_{-i})), \theta)] \geq E_{t_i}[u_i((a_i, \sigma_{-i}(t_{-i})), \theta)]$$

where the expectation is taken with respect to $\pi_i(t_i)$ on both sides.

The collection $\mathcal{T} = \{\Theta, (T_i, \pi_i)_{\in N}\}$ is the *belief-closed type space* based on payoff environment Θ for game Γ. It is *finite* if each T_i is finite; it has *full support* if $\sum_\theta \pi_i(t_i)[\theta, t_{-i}] > 0$ for all i and all $t \in T$. For each player $i \in N$, each of his type $t_i \in T_i$ can be associated with an infinite belief hierarchy as in Mertens and Zamir (1985) and Brandenburger and Dekel (1993), which explicitly describes his belief about payoff-relevant parameters in Θ, his belief about the others' beliefs about elements in Θ, and so on. In this paper, we analyze conditions on a belief-closed type space $\mathcal{T}$ under which only beliefs up to some finite order are strategically relevant. We study special features of such belief-closed type spaces. Since most applications of Bayesian games (e.g. in auction and mechanism design literature) involve private values of the players, we begin our analysis with the private-value case and then extend the analysis to the common-value case.

2.1 The Private Value Case

2.1.1 Consistency and Convexity

Assume $\Theta = \prod_{i \in N} \Theta_i$, where Θ_i is a finite set of payoff-relevant parameter values for each player $i \in N$. Elements in Θ_i are payoff-relevant parameter values and are called *payoff types* of player i. A *posterior* belief of player i about other players' payoff types is a probability measure on $\Theta_{-i} = \Pi_{j \neq i} \Theta_j$. Such a belief is known as player i's *first-order* or *order-1* belief.

Now imagine that nature first draws a profile $(\theta, \phi) \in \Theta \times \Pi_{i \in N} \Delta(\Theta_{-i})$ according to a prior $F \in \Delta(\Theta \times \Pi_{i \in N} \Delta(\Theta_{-i}))$ and then informs each player i of both his payoff type θ_i and order-1 belief ϕ_i. Since player i can use F in combination with his private information (θ_i, ϕ_i) to derive a posterior belief about θ_{-i} via Bayes rule, this derived posterior belief must coincide with ϕ_i to avoid any contradiction for considering the latter as his belief.

Definition 1 *We say that* $F \in \Delta(\Theta \times \Pi_{i=1}^{n} \Delta(\Theta_{-i}))$ *is a* consistent order-2 prior *if for all* i *and* $(\theta_i, \phi_i) \in \operatorname{supp} \operatorname{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F$,

$$\phi_i(\tilde{\theta}_{-i}) = F(\tilde{\theta}_{-i} \mid \theta_i, \phi_i), \; \forall \tilde{\theta}_{-i} \in \Theta_{-i}. \tag{1}$$

The following proposition establishes the convexity of the class of consistent order-2 priors. The convexity of the space of consistent order-2 priors is an important property. For example, it enables us to show that two players' heterogenous priors can always be embedded in a consistent order-2 prior.

Proposition 1 *Any convex combination of two consistent order-2 priors is again a consistent order-2 prior.*

Proof. Let F' and F'' be any two consistent order-2 priors (i.e., in $\Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})))$ and let λ be any weight in (0, 1). Then, $F = \lambda F' + (1-\lambda)F''$ is an order-2 prior and $\mathrm{supp} F = \mathrm{supp} F' \cup \mathrm{supp} F''$.

Fix player i and consider $(\theta_i, \phi_i) \in \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F$. If

$$(\theta_i, \phi_i) \in \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F' \setminus \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F'',$$

then $F(\theta_i, \theta_{-i}, \phi_i) = \lambda F'(\theta_i, \theta_{-i}, \phi_i)$ and

$$\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F(\theta_i, \phi_i) = \lambda \mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F'(\theta_i, \phi_i)$$

for all $\theta_{-i} \in \Theta_{-i}$. It follows $F(\theta_{-i}|\theta_i, \phi_i) = F'(\theta_{-i}|\theta_i, \phi_i)$. In this case, the consistency of F' implies $\phi_i(\theta_{-i}) = F(\theta_{-i}|\theta_i, \phi_i)$. Similarly, if

$$(\theta_i, \phi_i) \in \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F'' \setminus \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F',$$

then

$$\phi_i(\theta_{-i}) = F''(\theta_{-i}|\theta_i, \phi_i) = F(\theta_{-i}|\theta_i, \phi_i)$$

for all $\theta_{-i} \in \Theta_{-i}$. Finally, if $(\theta_i, \phi_i) \in \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F' \cap \mathrm{supp}\,\mathrm{marg}_{\Theta_i \times \Delta(\Theta_{-i})} F''$, we have

$$\phi_i(\theta_{-i}) = F'(\theta_{-i}|\theta_i, \phi_i) = F''(\theta_{-i}|\theta_i, \phi_i).$$

Applying the fact of $a/b = c/d$ implying $[\lambda a + (1-\lambda)c]/[\lambda b + (1-\lambda)d] = a/b$, we get $\phi_i(\theta_{-i}) = F(\theta_{-i}|\theta_i, \phi_i)$. ■

Note that a convex combination of two priors can be interpreted as a result of random choice by nature. As Example 1 below illustrates, richer information structures can be generated by having nature randomly select among simpler ones.

Example 1: Think of a bilateral trading game with an indivisible good between a buyer (player 1) and a seller (player 2). Assume the buyer's reservation value is either high or low, denoted θ_{1h} and θ_{1l}, while the seller has a single reservation value θ_2. Assume further there are two possible naive priors $P^A, P^B \in \Delta(\Theta)$ with $\Theta = \{(\theta_{1h}, \theta_2), (\theta_{1l}, \theta_2)\}$, due to different product origins, say place A and place B, initially unknown. It is commonly known that the probability for A (resp. B) being the true origin is λ (resp. $1-\lambda$). As private information at the subsequent stage, the buyer is informed of his realized reservation value and the seller of the product origin. What would be an appropriate type space to model this information incompleteness?

To answer the question, consider the following order-2 priors F^A and F^B with

$$\text{supp}F^A = \{(\theta_{1h}, \theta_2, \phi_1^A, \phi_2^A), (\theta_{1l}, \theta_2, \phi_1^B, \phi_2^B\}, \ F^A(\theta_{1h}, \theta_2, \phi_1^A, \phi_2^A) = P^A(\theta_{1h}, \theta_2)$$

and

$$\text{supp}F^B = \{(\theta_{1h}, \theta_2, \phi_1^B, \phi_2^B), (\theta_{1l}, \theta_2, \phi_1^B, \phi_2^B)\}, \ F^B(\theta_{1h}, \theta_2, \phi_1^B, \phi_2^B) = P^B(\theta_{1h}, \theta_2),$$

where $\phi_1^z = \delta_{\theta_2}, \phi_2^z(\theta_1) = P^z(\theta_1, \theta_2)$ for $z = A, B, \theta_1 \in \Theta_1$. By construction, F^A and F^B generate exactly the same first-order beliefs for the agents as the naive priors P^A and P^B do. Further, they both satisfy the consistency condition (1).

By Proposition 1, $F^\lambda = \lambda F^A + (1-\lambda)F^B$ is consistent with $\text{supp}F^\lambda = \text{supp}F^A \cup \text{supp}F^B$ for all $\lambda \in (0,1)$. It has the following explicit form:

$$F^\lambda(\theta, \phi) = \begin{cases} \lambda P^A(\theta_{1h}, \theta_2) & \text{if } \theta = (\theta_{1h}, \theta_2), \ \phi = (\phi_1^A, \phi_2^A), \\ (1-\lambda)P^B(\theta_{1h}, \theta_2) & \text{if } \theta = (\theta_{1h}, \theta_2), \ \phi = (\phi_1^B, \phi_2^B), \\ \lambda P^A(\theta_{1l}, \theta_2) & \text{if } \theta = (\theta_{1l}, \theta_2), \ \phi = (\phi_1^A, \phi_2^A), \\ (1-\lambda)P^B(\theta_{1l}, \theta_2) & \text{if } \theta = (\theta_{1l}, \theta_2), \ \phi = (\phi_1^B, \phi_2^B), \\ 0 & \text{otherwise.} \end{cases}$$

It follows that random drawing by nature according to F^λ would yield the information structure in this example. Thus, F^λ can be used as the common prior to model the information structure. Note that F^λ is not naive because the same payoff type θ_2 of the seller is now associated with *two* different first-order beliefs ϕ_2^A and ϕ_2^B. Thus, interpreting a convex combination as nature's move to choose among naive priors enriches the information structure by *one* belief order higher. It turns out, as we discuss in section 4, that Feinberg and Skrzypacz (2005) exploited an information structure that can be generated by having F^λ for a suitable choice of λ as the common prior in their reconsideration of the Coase conjecture with higher-order uncertainty.

2.1.2 Embedding Naive Priors

We are interested in embedding players' naive heterogenous priors in a consistent order-2 common prior.

Definition 2 *We say that a profile $\mathcal{P} = (P_1, P_2, \ldots, P_n)$ of individual naive priors in $\Delta(\Theta)$ can be* embedded *in a consistent order-2 prior F if for all $\theta \in \Theta$ such that*

$\theta_i \in \operatorname{supp} \operatorname{marg}_{\Theta_i} P_i$ *for all* i,

$$(\theta, P(\cdot|\theta)) = (\theta, (P_1(\cdot|\theta_1), \cdots, P_n(\cdot|\theta_n))) \in \operatorname{supp} F. \tag{2}$$

Definition 2 requires that players' posterior beliefs determined by their naive priors be jointly drawn with positive probability by a consistent order-2 prior. Notice that the support of the common prior may contain elements that are not possible under the given heterogeneous priors. Consequently, the embedding common prior may add beliefs not relevant in the original heterogeneous type space.[4] Notice also that (2) implies that the individual naive priors must have the same support.

We now show that with $n = 2$, the players' naive heterogenous priors can always be embedded in a consistent order-2 nature, provided they have a common support.

Proposition 2 *Suppose $n = 2$. Then, the players' naive priors can always be embedded in an order-2 prior whenever they have a common support.*

Proof. Let P_1 and P_2 be the naive priors for player 1 and player 2 such that $\operatorname{supp} P_1 = \operatorname{supp} P_2$. For each $\bar{\theta} \in \operatorname{supp} P_i$, let $P^{\bar{\theta}} \in \Delta(\Theta)$ be a prior given by

$$P^{\bar{\theta}}(\theta) = P_1(\theta_2|\bar{\theta}_1) P_2(\theta_1|\bar{\theta}_2), \ \theta \in \Theta.$$

It is clear that $P^{\bar{\theta}}$ is well-defined and θ_1 and θ_2 are independently distributed under $P^{\bar{\theta}}$. Thus, for $i \neq j$, $(\theta_i, \bar{\theta}_j) \in \operatorname{supp} P_j$ implies

$$P^{\bar{\theta}}(\theta_j|\theta_i) = P_i(\theta_j|\bar{\theta}_i), \ \theta_j \in \Theta_j, \tag{3}$$

which does not vary with θ_i. By construction,

$$\theta \in \operatorname{supp} P^{\bar{\theta}} \Longleftrightarrow (\theta_i, \bar{\theta}_j) \in \operatorname{supp} P_j.$$

Thus, putting this together with (3), every $\theta \in \operatorname{supp} P^{\bar{\theta}}$ results in the same order-1 beliefs $P_1(\cdot|\bar{\theta}_1)$ for player 1 and $P_2(\cdot|\bar{\theta}_2)$ for player 2. It follows that

$$F^{\bar{\theta}}(K, \Phi) = P^{\bar{\theta}}(K \cap \operatorname{supp} P^{\bar{\theta}}) \delta_{P(\cdot|\bar{\theta})}(\Phi), \ K \times \Phi \subseteq \Theta \times \Pi_{i=1}^2 \Delta(\Theta_{-i}), \tag{4}$$

[4] This is similar to the *matching* of finite-order beliefs by common priors as in Lipman (2003). The differences between embedding as defined above and finite-order belief matching will be discussed later in this paper.

defines a probability measure on the algebra of Borel sets $K \times \Phi$ of the cartesian product $\Theta \times \Pi_{i=1}^{2} \Delta(\Theta_{-i})$. By the Extension Theorem (see Theorem 3.1 in Billingsley 1995), $F^{\bar{\theta}}$ can be extended to a unique probability measure on the σ-algebra generated by Borel sets $K \times \Phi \subseteq \Theta \times \Pi_{i=1}^{2} \Delta(\Theta_{-i})$. We use $F^{\bar{\theta}}$ itself to denote the extension. By (3) and (4), $F^{\bar{\theta}}$ is consistent and

$$\text{supp} F^{\bar{\theta}} = \text{supp} P_i \cup \{P(\cdot|\bar{\theta})\}. \tag{5}$$

Now choose $\alpha_{\bar{\theta}} > 0$ for $\bar{\theta} \in \text{supp} P_i$ such that $\sum_{\bar{\theta} \in \text{supp} P_i} \alpha_{\bar{\theta}} = 1$. Set

$$F = \sum_{\bar{\theta} \in \text{supp} P_i} \alpha_{\bar{\theta}} F^{\bar{\theta}}.$$

Then, F is a convex combination of priors $F^{\bar{\theta}}, \bar{\theta} \in \text{supp} P_i$. By Proposition 1, F is a consistent prior in $\Delta(\Theta \times \Pi_{i=1}^{2} \Delta(\Theta_{-i}))$. Furthermore, for all $\bar{\theta} \in \text{supp} P_i$, since $\text{supp} P_1 = \text{supp} P_2$ and $\text{supp} F^{\bar{\theta}} \subseteq \text{supp} F$, it follows from (5) that $(\bar{\theta}, P^{\bar{\theta}}(\cdot|\bar{\theta})) \in \text{supp} F$. This shows that F also satisfies (2). ■

The following proposition provides a generalization of Proposition 2 to the n-person case.

Proposition 3 *Let $\mathcal{P} = (P_1, \cdots, P_n)$ be a profile of personalized naive priors in $\Delta(\Theta)$. Suppose (i) $\text{supp} P_1 = \cdots = \text{supp} P_n$ and (ii) for all $\bar{\theta} \in \text{supp} P_1$, there exists $Q^{\bar{\theta}} \in \Delta(\Theta)$ such that for all $\theta \in \text{supp} P_1$ and $i \neq j$,*

$$P_i(\theta_{-i}|\bar{\theta}_i) Q^{\bar{\theta}}(\theta_i) = P_j(\theta_{-j}|\bar{\theta}_j) Q^{\bar{\theta}}(\theta_j).$$

Then, $\mathcal{P} = (P_1, \cdots, P_n)$ can be embedded in an order-2 consistent prior.

Proof. Fix $\bar{\theta} \in \Delta$ and define $P^{\bar{\theta}} \in \Delta(\Theta)$ by

$$P^{\bar{\theta}}(\theta) = P_i(\theta_{-i}|\bar{\theta}_i) Q^{\bar{\theta}}(\theta_i), \ \theta \in \Theta.$$

By conditions (i) and (ii), $P^{\bar{\theta}}$ is well-defined (i.e., belonging to $\Delta(\Theta)$ and independent of $i \in N$). Notice $P^{\bar{\theta}}(\theta_i) = Q^{\bar{\theta}}(\theta_i)$ for all $\theta_i \in \Theta_i$ and for all i. Hence,

$$P^{\bar{\theta}}(\theta_{-i}|\theta_i) = P_i(\theta_{-i}|\bar{\theta}_i), \ \forall \theta \in \Theta, \forall i.$$

The rest of the proof repeats the proof of Proposition 2. ■

Notice that condition (ii) is automatic when $n = 2$. In that case, we can simply take $Q^{\bar{\theta}}(\theta) = P_1(\theta_2|\bar{\theta}_1)P_2(\theta_1|\bar{\theta}_2)$ for all $\bar{\theta}, \theta \in \text{supp}P_1$. Notice also that condition (ii) implies that any two different players have the same belief about a third player's individual payoff types. That is, $P_j(\theta_i|\bar{\theta}_j) = P_k(\theta_i|\bar{\theta}_k)$, for all $\theta, \bar{\theta} \in \Theta$ and three different players i, j, k. This is a strong condition. We characterize necessary and sufficient conditions for $\mathcal{P}$ to be embedded in an order-2 consistent prior in an appendix to this paper.

2.1.3 Complete Information as Consistent Order-2 Priors

Games of complete information are often benchmarks for modeling economic problems. For example, Bergemann and Morris (2005) show that a class of social choice functions is *ex post* implementable if they are Bayesian-Nash implementable for all complete-information type spaces. Therefore, it is important to properly sort them out within the universal type space. Intuitively, players have complete information about $\bar{\theta} \in \Theta$ if every player i believes with probability 1 that the others' payoff types are as in $\bar{\theta}_{-i}$.

Definition 3 *Given an order-2 prior $F \in \Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i}))$ and given $\bar{\theta} \in \Theta$, we say that information about payoff-type profile $\bar{\theta}$ is* complete *if*

$$\phi_i(\cdot) = \delta_{\bar{\theta}_{-i}}(\cdot)$$

for all i and for all ϕ such that $(\theta, \phi) \in \text{supp}F$; we say that information is complete or F is a complete-information prior *if the above condition holds for all $\bar{\theta} \in \text{supp}F$.*

Observe that the definition implies that F is consistent whenever F is a complete-information prior. To explore further special properties, let $F^{\bar{\theta}}$ be the order-2 prior defined by

$$F^{\bar{\theta}}(\theta, \phi) = \delta_{\bar{\theta}}(\theta)\delta_{(\delta_{\bar{\theta}_{-i}})_{i=1}^n}(\phi), \;\; \bar{\theta} \in \Delta. \qquad (6)$$

Then, $F^{\bar{\theta}}$ is consistent and $(\theta, \phi) \in \text{supp}F^{\bar{\theta}}$ if and only if $\theta = \bar{\theta}$ and $\phi = (\delta_{\bar{\theta}_{-i}})_{i=1}^n$. That is, the one-point prior $F^{\bar{\theta}}$ in (6) is a complete-information prior. Furthermore, we have

Proposition 4 *An order-2 prior $F \in \Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i}))$ is a complete-information prior if and only if*

$$F \in \text{convex hull of } \; \{F^{\bar{\theta}}\}_{\bar{\theta}\in\Theta}.$$

Proof. Suppose that $F \in \Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i}))$ is a complete information prior. From Definition 3,

$$\sum_{\bar{\theta}\in\Theta} F(\bar{\theta}, (\delta_{\bar{\theta}_{-i}})_{i=1}^n) = 1.$$

Thus, from (6),

$$F(\theta,\phi) = \sum_{\bar\theta\in\Theta} F(\bar\theta,(\delta_{\bar\theta_{-i}})_{i=1}^n)F^{\bar\theta}(\theta,\phi).$$

This shows that F is in the convex hull of $\{F^{\bar\theta}\}_{\bar\theta\in\Theta}$.

Conversely, given two complete-information order-2 priors F', F'' and given a number $\alpha\in[0,1]$, it follows from Proposition 1 that the convex combination $F=\alpha F' + (1-\alpha)F''$ is consistent. Since

$$\text{supp}F = \text{supp}F' \cup \text{supp}F'',$$

$(\bar\theta,\bar\phi)\in\text{supp}F$ implies $(\bar\theta,\bar\phi)\in\text{supp}F'$ or $(\bar\theta,\bar\phi)\in\text{supp}F''$. In either case, since F' and F'' are both complete-information priors, $\bar\phi=(\delta_{\bar\theta_{-i}})_{i=1}^n$. Hence, F is also a complete-information prior. This shows that priors in the convex hull of $\{F^{\bar\theta}\}_{\bar\theta\in\Theta}$ are all complete-information priors. ■

Proposition 4 shows in particular that complete-information priors form a convex subspace of the space of consistent order-2 priors. We show in the next section that a consistent order-2 prior corresponds to an order-2 universal type sub-space.

2.2 The Common Value Case

Let Θ be the space of payoff-relevant parameter values not necessarily decomposable into individual elements. In this case, the information setting is such that nature first draws a belief in $\Delta(\Theta)$ for each player and an element in Θ according to a prior $F\in\Delta(\Theta\times\Pi_{i=1}^n\Delta(\Theta))$. Then, player i is informed of his order-1 belief over Θ. Unlike the private-value case, no part of θ is directly revealed to him as part of private information. Thus, prior F is *consistent* if for all $(\theta,\phi)\in\text{supp}F$,

$$\phi_i(\tilde\theta) = F(\tilde\theta|\phi_i),\ \tilde\theta\in\Theta,\ i\in N. \tag{7}$$

With (7) replacing (1), a parallel extension of our previous analysis and results can be naturally made. In particular, any common-value complete information type space can be written as an order-2 consistent prior F such that $(\theta,\phi)\in\text{supp}F$ if and only if $\phi_i=\delta_\theta$ for all i.

A private-value prior can be viewed as a special form of common-value prior. To see this, let $\Theta=\Pi_{i=1}^n\Theta_i$ and let $F\in\Delta(\Theta\times\Pi_{i=1}^n\Delta(\Theta_{-i}))$ be a consistent prior satisfying

(1). Then,

$$F'(\theta, \psi) := F(\theta, (\text{marg}_{\Theta_{-i}} \psi_i)_{i=1}^n) \cdot \delta_{(\delta_{\theta_i})_{i=1}^n}((\text{marg}_{\Theta_i} \psi_i)_{i=1}^n)$$

for all $(\theta, \psi) \in (\Theta, \Pi_{i=1}^n \Delta(\Theta))$ is a consistent common-value order-2 prior with the additional information on private values. It is clear that F and F' yield the same information structure. Unless needed, our subsequent discussion will focus on the common-value case.

3 Finite-Order Belief Hierarchies, Priors, and Type Spaces

In this section we extend results in section 2 to belief hierarchies of finite order $k-1$ for any arbitrary integer $k \geq 2$. A consistent prior on players' belief hierarchies of order $k-1$ is naturally of order k. When viewing such an order k consistent prior as specifying an abstract belief-closed type space, an order-k "canonical" type space is induced which has the same order-(k-1) belief hierarchies for each player.

An abstract belief-closed type space is of order k if it has a type morphism with the canonical type space of some consistent order-k prior. We establish necessary and sufficient conditions for an abstract belief-closed type space to be of order k.

By truncating the belief hierarchies of a type space at a certain finite order, a finite-order projection of a type space is obtained. We introduce an order-k total variation norm to the projection priors to refine a notion of robustness of solutions. We fix Θ as the given common-value payoff environment throughout the rest of this section.

3.1 Order-(k-1) Belief Hierarchies and Order-k Priors

Let $k \geq 2$ be an integer. For each player i, the following recursive construction of the spaces of belief hierarchies of orders $h = 1, \cdots, k-1$ is standard:

$$\begin{aligned} T_i^1 &= \Delta(\Theta), \\ T_i^2 &= T_i^1 \times \Delta(\Theta \times T_{-i}^1), \\ &\vdots \\ T_i^{k-1} &= T_i^{k-2} \times \Delta(\Theta \times T_{-i}^{k-2}),\ k > 2. \end{aligned} \tag{8}$$

An element $t_i^{k-1} \in T_i^{k-1}$ is identified with player i's order-1 belief $\phi_i^1 \in \Delta(\Theta)$, order-2 belief $\phi_i^2 \in \Delta(\Theta \times T_{-i}^1)$, $\cdots$, and order-(k-1) belief $\phi_i^{k-1} \in \Delta(\Theta \times T_{-i}^{k-2})$. We write it as $t_i^{k-1} = (\phi_i^1, \phi_i^2, \cdots, \phi_i^{k-1})$ and refer to it as an (k-1)-belief hierarchy of player i. As usual, t_i^{k-1} is *coherent* if

$$\text{marg}_\Theta \phi^2 = \phi^1 \text{ and } \text{marg}_{\Theta \times T_{-i}^{h-2}} \phi_i^h = \phi_i^{h-1},\ h = 3, \cdots, k-1. \tag{9}$$

Set $T^{k-1} = \Pi_{i=1}^n T_i^{k-1}$, where T_i^{k-1} is given in (8). Elements in $\Delta(\Theta \times T^{k-1})$ are order-k priors. The next definition extends our notion of consistency to order-k priors.

Definition 4 *For $k \geq 2$, an order-k prior $F \in \Delta(\Theta \times T^{k-1})$ is* consistent *if for all $t^{k-1} \in \text{supp}\,\text{marg}_{T^{k-1}} F$ and for all i, $t_i^{k-1} = (\phi_i^1, \cdots, \phi_i^{k-1})$ is coherent and*

$$\phi_i^{k-1}(\theta, \tilde{t}_{-i}^{k-2}) = F(\theta, \tilde{t}_{-i}^{k-2} | t_i^{k-1}),\ \forall (\theta, \tilde{t}_{-i}^{k-2}) \in T_{-i}^{k-2}. \tag{10}$$

Notice that the consistency (10) of F together with the coherency (9) of $t_i^{k-1} \in \text{supp}\,\text{marg}_{T_i^{k-1}} F$ implies consistency at lower orders. That is,

$$\phi_i^1(\theta) = F(\theta | t_i^{k-1})$$

and

$$\phi_i^h(\theta, \tilde{t}_{-i}^{h-1}) = F(\theta, \tilde{t}_{-i}^{h-1} | t_i^{k-1}),\ \tilde{t}_{-i}^{h-1} \in T_{-i}^{h-1},\ 2 \leq h \leq k-1. \tag{11}$$

By sorting terms carefully, a direct application of Proposition 1 establishes the next result.

Corollary 1 *Consistent order-k priors consist of a convex subspace of* $\Delta(\Theta \times T^{k-1})$.

3.2 Order-(k-1) Belief Hierarchies in Belief-Closed Type Spaces

Recall that an abstract belief-closed type space is a collection $\mathcal{T} = \{T_i, \pi_i\}_{i \in N}$, where $\pi_i : T_i \longrightarrow \Delta(\Theta \times T_{-i})$ is the belief mapping for player i. For any $k \geq 2$, each of player i's type $t_i \in T_i$ generates recursively via the standard method an order-(k-1) belief hierarchy $\tau_i^k(t_i|\mathcal{T}) = (\psi_i^1(t_i|\mathcal{T}), \cdots, \psi_i^{k-1}(t_i|\mathcal{T}))$ as follows:

$$\psi_i^1(t_i|\mathcal{T})[\theta] = \sum_{t_{-i} \in T_{-i}} \pi_i(t_i)[\theta, t_{-i}],\ \forall\, \theta \in \Theta, \tag{12}$$

and for $2 \leq h \leq k-1$ and for all $(\theta, t_{-i}^{h-1}) \in \Theta \times T_{-i}^{h-1}$,

$$\psi_i^h(t_i|\mathcal{T})[\theta, t_{-i}^{h-1}] = \sum_{\substack{t_{-i} \in T_{-i}:\ t_{-i}^{h-1} = \\ (\psi_{-i}^1(t_{-i}|\mathcal{T}), \psi_{-i}^2(t_{-i}|\mathcal{T}), \cdots, \psi_{-i}^{h-1}(t_{-i}|\mathcal{T}))}} \pi_i(t_i)[\theta, t_{-i}]. \quad (13)$$

Here, $\psi_i^h(t_i|\mathcal{T})[\theta, t_{-i}^{h-1}]$ is the probability with which player i believes that the payoff-relevant parameter is θ and the other players' order-$(h-1)$ belief hierarchies are prescribed by $t_{-i}^{h-1} \in T_{-i}^{h-1}$, given that his own type is $t_i \in T_i$. By construction, $\tau_i^k(t_i|\mathcal{T})$ is coherent. Furthermore, it is a measurable transformation from T to T^{k-1}

Let ρ_i^k denote the *inverse* of τ_i^k. Then,

$$\rho_i^k(t_i^{k-1}|\mathcal{T}) = \{t_i \in T_i | \tau_i^k(t_i|\mathcal{T}) = t_i^{k-1}\} \quad (14)$$

for all $t_i^{k-1} \in T_i^{k-1}$.[5] Notice that

$$\tau_i^\infty(t_i|\mathcal{T}) = \lim_{k \longrightarrow \infty} \tau_i^k(t_i|\mathcal{T})$$

is the belief hierarchy in the *universal type space* based on Θ that corresponds to player i's type $t_i \in T_i$. Subsequently, we will suppress $\mathcal{T}$ from τ_i^k and ρ_i^k to save space whenever there is no danger of confusion.

3.3 The Canonical Belief-Closed Type Space of an Order-k Prior

An order-k prior $F \in \Delta(\Theta \times T^{k-1})$ uniquely determines a full-support common-prior belief-closed type space, $\mathcal{T}^F = \{T_i^F, \pi_i^F\}_{i=1}^n$, where

$$T_i^F = \text{supp}\,\text{marg}_{T_i^{k-1}} F \quad (15)$$

is the set of types for player i and

$$\pi_i^F(t_i^F)[\theta, t_{-i}^F] = F(\theta, t_{-i}^F | t_i^F),\ \forall\ (\theta, t_{-i}^F) \in \Theta \times T_{-i}^F,\ t_i^F \in T_i^F. \quad (16)$$

By viewing $\mathcal{T}^F$ as an abstract belief-closed type space, (12) and (13) can be applied to derive players finite-order belief hierarchies $\tau_i^h(t_i^F)$ for all $t_i^F \in T_i^F$ and all i. The

[5]See Billingsly (1995, pp. 182, 537) for formal definitions of measurable mappings and properties of inverse images of measurable mappings.

following lemma shows that players' order-(k-1) belief hierarchies in $\mathcal{T}^F$ coincide with those in supp$F$, provided F is consistent.

Lemma 1 *Let $F \in \Delta(\Theta \times T^{k-1})$ be consistent for $k \geq 2$. Then, $\tau_i^k(t_i^F) = t_i^F$ for all $t_i^F \in T_i^F$ and for all i.*

Proof. By (12), (15), and (16), for all $\theta \in \Theta$ and for all $t_i^F = (\phi_i^1, \cdots, \phi_i^{k-1}) \in T_i^F$,

$$\psi_i^1(t_i^F)[\theta] = \sum_{t_{-i}^F \in T_{-i}^F} \pi_i(t_i^F)[\theta, t_{-i}^F] = \sum_{t_{-i}^F \in T_{-i}^F} F(\theta, t_{-i}^F | t_i^F) = F(\theta | t_i^F).$$

Thus, by (11), $\psi_i^1(t_i^F) = \phi_i^1$. Consequently, by the definition of τ_i^h and (15),

$$\tau_i^2(t_i^F) = (\psi_i^1(t_i^F)) = \text{proj}_{T_i^1} t_i^F, \ \forall t_i^F \in T_i^F, \ \forall i.$$

Proceeding now with induction, suppose for all i and $2 \leq h \leq k-1$,

$$\tau_i^h(t_i^F) = (\phi_i^1, \phi_i^2, \cdots, \phi_i^{h-1}) = \text{proj}_{T_i^{h-1}} t_i^F. \tag{17}$$

Then, by (13)-(16), for all $(\theta, t_{-i}^{h-1}) \in \Theta \times T_{-i}^{h-1}$,

$$\begin{aligned} \psi_i^h(t_i^F)[\theta, t_{-i}^{h-1}] &= \sum_{t_{-i}^F \in T_{-i}^F:\ t_{-i}^F \in \rho_{-i}^h(t_{-i}^{h-1})} \pi_i^F(t_i)[\theta, t_{-i}^F] \\ &= \sum_{t_{-i}^F \in T_{-i}^F:\ t_{-i}^F \in \rho_{-i}^h(t_{-i}^{h-1})} F(\theta, t_{-i}^F | t_i^F) \end{aligned}$$

By (17), $\tau_{-i}^h(t_{-i}^F) = \text{proj}_{T_{-i}^{h-1}} t_{-i}^F$ for all $t_{-i}^F \in T_{-i}^F$. It follows that

$$\sum_{t_{-i}^F \in T_{-i}^F:\ t_{-i}^F \in \rho_{-i}^h(t_{-i}^{h-1})} F(\theta, t_{-i}^F | t_i^F) = F(\theta, t_{-i}^{h-1} | t_i^F).$$

Putting the preceding equations together with consistency (11) shows

$$\psi_i^h(t_i^F)[\theta, t_{-i}^{h-1}] = \phi_i^h(\theta, t_{-i}^{h-1}) \Rightarrow \tau_i^{h+1}(t_i^F) = \text{proj}_{T_i^h} t_i^F.$$

Thus, by induction, $\tau_i^k(t_i^F) = \text{proj}_{T_i^{k-1}} t_i^F$ holds. Since $\text{proj}_{T_i^{k-1}} t_i^F = t_i^F$, the proof is completed. ■

3.4 The Order of a Belief-Closed Type Space

Recall that a belief-closed type space satisfies the *common prior assumption* if there is a prior $P \in \Delta(\Theta \times T)$ such that for all i and for all $t_i \in T_i$, $\text{marg}_{T_i} P(t_i) > 0$ and $\pi_i(t_i)[\theta, t_{-i}] = P(\theta, t_{-i}|t_i)$ for all $(\theta, t_{-i}) \in \Theta \times T_{-i}$. Now consider two belief-closed type spaces $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ and $\mathcal{T}' = \{T_i', \pi_i'\}_{i=1}^n$ based on Θ. Following Heifetz and Samet (1998), a mapping $f = (f_0, f_1, \cdots, f_n)$ from $\Theta \times T$ to $\Theta \times T'$, with $f_0 : \Theta \mapsto \Theta$ and $f_i : T_i \mapsto T_i'$, is a *type morphism* from $\mathcal{T}$ to $\mathcal{T}'$ if f_0 is the identity mapping and

$$\pi_i'(f_i(t_i))[E_{-i}'] = \pi_i(t_i)[f_{-i}^{-1}(E_{-i}')]. \tag{18}$$

for all $i = 1, 2, \cdots, n$, $t_i \in T_i$, and $E_{-i}' \subseteq \Theta \times T_{-i}'$. A type morphism between $\mathcal{T}$ and $\mathcal{T}'$ preserves universal types (see Heifetz and Samet, 1998, for details).

Definition 5 *We say that a belief-closed type space $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ is of* order-k *for some integer $k \geq 1$, if there is a type morphism from $\mathcal{T}$ to the canonical type space $\mathcal{T}^F$ of some consistent prior $F \in \Delta(\Theta \times T^{k-1})$.*

The following theorem establishes necessary and sufficient conditions for a common-prior belief-closed type space to be of order k.

Theorem 1 *A full support common-prior belief-closed type space $\mathcal{T} = (T_i, \pi_i)_{i=1}^n$ is of order-k for some $k \geq 1$ if and only if for all i and $t_i, t_i' \in T_i$,*

$$\tau_i^k(t_i) = \tau_i^k(t_i') \Rightarrow \tau^\infty(t_i) = \tau^\infty(t_i'). \tag{19}$$

Proof. Consider first the case with $k \geq 2$. Suppose that $\mathcal{T}$ is of order-k. Then, by Definition 5, there exist a consistent order-k prior $F \in \Delta(\Theta \times T^{k-1})$ and a type morphism $f = (f_0, f_1, \cdots, f_n)$ from $\mathcal{T}$ to $\mathcal{T}^F$. By (15), (16), and (18),

$$\tau_i^k(T_i) = \tau_i^k(T_i^F). \tag{20}$$

Since $t_i^F \neq t_i'^F$ for all $t_i^F, t_i'^F \in T_i^F$ by construction of T_i^F, (20) together with Lemma 1 implies either $\tau^k(t_i) \neq \tau^k(t_i')$ or $\tau^\infty(t_i) = \tau^\infty(t_i')$ for all $t_i, t_i' \in T_i$ and all i. This establishes the necessity of (19).

To prove the sufficiency, let $P \in \Delta(\Theta \times T)$ be the common prior for $\mathcal{T} = (T_i, \pi_i)_{i=1}^n$. Define $F \in \Delta(\Theta \times T^{k-1})$ by

$$F(\theta, t^{k-1}) = P(\theta, \rho^k(t^{k-1})) \tag{21}$$

for $(\theta, t^{k-1}) \in \Theta \times T^{k-1}$, where ρ^k is the inverse of τ^k as in (14). Then, since $P \in \Delta(\Theta \times T)$, it follows from (19) and (21) that $F \in \Delta(\Theta \times T^{k-1})$. By the construction of τ_i^k, F is consistent. Finally, by (19), the mapping $f = (f_0, f_1, \cdots, f_n)$ with f_0 the identity mapping on Θ and $f_i(t_i) = \tau_i^k(t_i)$ for all $t_i \in T_i$ and all i is a type morphism from $\mathcal{T}$ to $\mathcal{T}^F$.

Consider now the case with $k = 1$. Then, $\Delta(\Theta \times T^{k-1})$ reduces to $\Delta(\Theta)$. In this case, the only order-1 belief of each player is the prior F itself. A similar proof as for the previous case establishes the sufficiency and necessity of (19). ■

By (19), all non-redundant types of an order-k type space have different order-(k-1) belief hierarchies. Notice that for any two types t_i and t_i' of player i in a type space, we have $\tau_i^k(t_i) \neq \tau_i^k(t_i')$ for some $k < \infty$ or $\tau^\infty(t_i) = \tau^\infty(t_i')$. Thus, when the type space is finite, there exists a finite integer k such that any two types of each player can be separated at the (k-1)-th order beliefs or have the same infinite belief hierarchy (i.e. one of them is redundant). Thus, a direct application of Theorem 1 establishes the following corollary.

Corollary 2 *Any finite common-prior belief-closed type space is of finite order.*

In an order-k type space, the universal hierarchy $\tau_i^\infty(t_i)$ contains no additional information not already contained in $\tau_i^k(t_i)$. That is, the tails of the belief hierarchies in the type space can be ignored for strategic analysis. The central message of Theorem 1 is that any order-k type space has the canonical representation of an order-k consistent prior as its non-redundant equivalence.

An handy implication of our approach is that convergence of finite-order priors can be handled using topologies of finite product spaces instead of infinite ones. Further discussion along this line is provided after the introduction of finite order projections. Notice that if $\mathcal{T}$ is of order-k, then it is also of order-l for all $l > k$. Naturally, for the sake of minimal complexity of strategically relevant hierarchies, we are only interested in the smallest such number.

Remark 1 *In a heterogenous prior type space, the consistency can be applied to the personalized priors. Thus, with heterogeneity, we define the order of an heterogenous-prior type space to be the maximum order of the individual priors.*

Remark 2 *In the private-value case, substituting T^{k-2} for the term Θ in the statements, Proposition 2 and Proposition 3 can be generalized to cover heterogeneous type* *spaces of arbitrary finite order $k \geq 2$. That is, order-(k-1) heterogeneous priors can be embedded in an order-k consistent prior.*[6]

Lipman (2003) shows that given any common-value finite type space and given any order $k \geq 1$, there exists a finite common-prior type space such that any k-belief hierarchies of the former can find its *matching* in the latter. However, the *matching* of common-prior type space constructed in his analysis is not exactly one-order higher than those of the heterogeneous priors as we do.[7] On the other hand, our tighter embedding result can only be proved for private-value case and that, aside from $n = 2$, we cannot embed all finite order type spaces.

The following is an example of an order-3 private-value common-prior type space. Complete-information type spaces can be constructed from this example as those priors that assign non-zero probabilities only to the diagonal. This example illustrates that perturbations to complete-information type spaces may easily lead to higher-order type spaces.[8]

Example 2: Consider a private-value belief-closed type space $\mathcal{T} = (T_i, \hat{\theta}_i, \hat{\pi}_i)_{i=1}^2$ where for $i = 1, 2$, $\Theta_i = \{\theta_{i1}, \theta_{i2}\}$, $T_i = \{t_{i1}, t_{i2}, t_{i3}, t_{i4}\}$,

$$\hat{\theta}_i(t_i) = \begin{cases} \theta_{i1} & \text{if } t_i = t_{i1}, t_{i2}; \\ \theta_{i2} & \text{if } t_i = t_{i3}, t_{i4}, \end{cases}$$

[6] In a private-value type space, the payoff environment Θ is the cartesian product of personalized payoff-relevant parameters Θ_i for all i. In such a private-value type space, $\theta_i \in \Theta_i$ becomes part of player i's private information. That is, each type t_i also specifies a payoff type $\hat{\theta}_i(t_i)$ for player i as well as a belief type $\hat{\pi}_i(t_i)$. Thus, in the description of the belief hierarchy of player i's type t_i, it is also required that player i believe his payoff type is $\hat{\theta}_i(t_i)$ with probability 1:

$$\text{marg}_{\Theta_i} \phi_i^h(t_i) = \delta_{\hat{\theta}_i(t_i)}, \ \forall h.$$

A private-value type space is customarily summarized as $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i \in N}$ and with private-value environment $\Theta = \Pi_{i=1}^n \Theta_i$, the spaces of finite-order belief hierarchies are recursively constructed as: $T_i^0 = \Theta_i, T_i^1 = T_i^0 \times \Delta(T_{-i}^0), \cdots, T_i^{k-1} = T_i^{k-2} \times \Delta(T_{-i}^{k-2})$. An element in T_i^{k-1} is then a (k-1)-belief hierarchy $t_i^{k-1} = (\theta_i, \phi_i^1, \cdots, \phi_i^{k-1})$ with $\theta_i \in T_i^0, \phi_i^1 \in \Delta(T_{-i}^0), \cdots, \phi_i^{k-1} \in \Delta(T_{-i}^{k-2})$. With these in place, concepts and results for common-value cases can be established in a parallel fashion for private-value cases.

[7] For example, in his illustration of matching belief hierarchies of order k, the common-prior type space is constructed by adding 2^k copies of the players' types of the original type space. This increases the order of the type space by more than one order (Lipman, 2003, p. 1259). An important implication of Lipman's (2003) results is that finite common-prior type spaces are dense in the universal space. Note that in general it is not easy to translate results obtained in partitions models into type space models, and vice versa.

[8] Unfamiliar users of belief-closed type spaces may make the mistake to take the abstract symbols for types as having fixed meanings while manipulating the priors. This is avoided in our consistent-prior approach as we only track the "belief-relevant meaning" instead of the abstract symbols.

and $\hat{\pi}_i$ is given by a common prior over $T_1 \times T_2$ in the following table:

	t_{21}	t_{22}	t_{23}	t_{24}
t_{11}	a	0	b	0
t_{12}	$\frac{a}{2}$	$\frac{a}{2}$	b	0
t_{13}	0	$\frac{1-2(a+b)}{2}$	0	0
t_{14}	0	0	0	$\frac{1-2(a+b)}{2}$

Table 1: An order-3 belief-closed type space

By (12), (13), and (19), simple calculations show that $\mathcal{T}$ is not of order 1 or order 2 but is of order 3.

3.5 Order-k Projection Type Spaces and Imperfect Nature

For any $k \geq 2$, the mappings $\tau^k(\cdot)$ in (12) and (13) can be applied to project a belief-closed type space $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ with common prior P to T^{k-1}as follows. Let $\tilde{T}_i^{k-1} = \tau_i^k(T_i) \subset T_i^{k-1}$ be the set of player i's *projection* types. Set $\tilde{F}^k = \text{marg}_{\Theta \times T^{k-1}} P$ so that

$$\tilde{F}^k(\theta, t^{k-1}) = P(\theta, \rho^k(t^{k-1})) \tag{22}$$

for $(\theta, t^{k-1}) \in \Theta \times \tilde{T}^{k-1}$. Next, let $\tilde{\pi}_i^k : \tilde{T}_i^{k-1} \longrightarrow \Delta(\Theta \times \tilde{T}_{-i}^{k-1})$ be given by

$$\tilde{\pi}_i^k(t_i^{k-1})[\theta, t_{-i}^{k-1}] = \tilde{F}^k(\theta, t_{-i}^{k-1} | t_i^{k-1}) \tag{23}$$

for $(\theta, t_{-i}^{k-1}) \in \Theta \times \tilde{T}_{-i}^{k-1}$ and $t_i^{k-1} \in T_i^{k-1}$. It follows from (22) and (23) that $\tilde{\mathcal{T}}^k = \{\tilde{T}_i^{k-1}, \tilde{\pi}_i^k\}_{i \in N}$ is itself a belief-closed type space of order-k with common prior $\tilde{F}^k$. We refer to $\tilde{\mathcal{T}}^k$ as the order-k projection of $\mathcal{T}$.

With these notations in place, an iterated application of the mappings τ^h establishes the following result.

Lemma 2 *Let $\mathcal{T}$ be a belief-closed type space. Then, $\tau^k(\tau^h(t)) = \tau^k(t)$ for all type profile $t \in \mathcal{T}$ and $h \geq k \geq 2$.*[9]

With our characterization of order-k type spaces, projections of orders $h > k$ do not carry more information than order-k projection. That is, as a corollary of Lemma 2 and Theorem 1, we have

[9]To derive the order-(k-1) belief hierarchy profile $\tau^k(\tau^h(t))$, we view $\tau^h(t)$ as an abstract type profile with the order-h projection prior.

Corollary 3 *Let $\mathcal{T}$ be an order-k type space. Then, for all $h > k$ and $t^{h-1} \in \tilde{T}^{h-1}$,*

$$\tilde{\pi}_i^h(t_i^{h-1})[\theta, t_{-i}^{h-1}] = \tilde{\pi}_i^k(\tau_i^k(t_i^{h-1}))[\theta, \tau_{-i}^k(t_{-i}^{h-1})].$$

Suppose that it is common knowledge that the true model of information structure is captured by the type space $\mathcal{T}$, which *cerebrally* can be of infinite order. One way to model *bounded rationality* is perhaps to have an "imperfect nature", which draws a combination (θ, t) and informs the players of their corresponding beliefs of some (randomly determined) finite order-k instead of the entire belief hierarchies. Our order-k consistent prior approach reflects realization of such imperfect nature.

Note that $\tilde{\mathcal{T}}^k$ converges to $\mathcal{T}$ in the weak product topology. A natural question is whether the BNEs of a game with the projection type spaces converge to those of the original one as the order k approaches to infinity. Rubinstein's (1989) email game shows that this is not the case in general. However, Weinstein and Yildiz (2007b) showed that under a *global stability condition*, which implies dominance solvability of the game, the sensitivity of BNEs with respect to variations of higher order beliefs vanishes as the order of beliefs approaching to infinity. Incorporating their arguments into our framework of projection type spaces, this condition ensures convergence of BNEs with the projection type spaces to those with the original type space. We omit a formal discussion here to avoid complexity of too many new definitions. But, we illustrate the convergence with the following example taken from Weinstein and Yildiz (2007b).

Example 3: Consider a Cournot duopoly with the demand function $p(Q) = \theta - Q$, where $\theta \in \Theta$ is an unknown demand parameter. Assuming zero production costs, firm i's profit at quantities q_i, q_j and parameter θ is given by

$$v_i(\theta, q_i, q_j) = q_i \cdot (\theta - q_i - q_j).$$

Let $\mathcal{T} = \{T_i, \pi_i\}_{i \in N}$ denote a type space with common prior P based on Θ. A strategy profile (q_1^*, q_2^*) with $q_i^* : T_i \mapsto A$,where A is a compact action space of real numbers, is a BNE if for each type t_i, $q_i^*(t_i)$ maximizes $q_i \cdot (E[\theta|t_i] - q_i - E[q_j^*|t_i])$. This yields the following first-order condition:

$$q_i^*(t_i) = 2^{-1}(E[\theta|t_i] - E[q_j^*|t_i])$$

For a random variable X with appropriate measurability, let $E_{t_i^h}[X] = E[X|t_i^h]$,

$E_i[X] = E[X|\mathcal{B}(T_i)]$, and $E_i^h[X] = E[X|\mathcal{B}(T_i^h)]$, where $\mathcal{B}(T_i)$ and $\mathcal{B}(T_i^h)$ denote the sigma-algebras of Borel sets in T_i and T_i^h with product topology, respectively. By iterated substitution using both players' first-order conditions, we have

$$q_i^*(t_i) = 2^{-1}E_{t_i}[\theta] - 2^{-2}E_{t_i}E_j[\theta] + 2^{-3}E_{t_i}E_jE_i[\theta] - \cdots .$$

Notice that $E[\theta|t_i] = E_{\tau_i^1(t_i)}[\theta]$. By repeated use of the law of iterated expectation,

$$\begin{aligned} q_i^*(t_i) &= 2^{-1}E_{\tau^1(t_i)}[\theta] - 2^{-2}E_{\tau^2(t_i)}E_j^1[\theta] + 2^{-3}E_{\tau^3(t_i)}E_j^2E_i^1[\theta] - \cdots \\ &= \sum_{m=1}^{\infty}(-1)^{m+1}2^{-m}E_{\tau_i^m(t_i)}E_j^{m-1}E_i^{m-2}\cdots E_{i_m}^1[\theta] \end{aligned}$$

where $i_m = i$ whenever m is odd and $i_m = j$ otherwise. This is i's explicit strategy in the unique BNE. If $\mathcal{T}$ is of order-k, then, by Corollary 3, $E_i^h = E_i^k$ for all $h > k$ and thus the equilibrium strategy depends only on belief hierarchies up to order k. In general, if $q^*(k)$ denotes the BNE with the order-k projection of $\mathcal{T}$. then, its first k terms in the explicit solution are identical with those with the original type space. By the boundedness of the action space, the terms in the tail converge to zero as fast as 2^{-k}. It follows that $q^*(k) \to q^*$.

3.6 Total Variation Norm, Solution continuity, and Robustness

Continuity of solution with varying type spaces has been a focus of recent research in the literature. For the interim correlated rationalizability, Chen et al (2009) showed that continuity is always preserved for uniform weak convergence. Given any finite integer $k \geq 1$, uniform weak convergence is equivalent to weak convergence in T^k. However, as we illustrate in the next section, weak convergence may be too weak to guarantee continuity of BNE even for $k = 2$.

We consider a stronger notion of convergence for finite order applications. Recall that the *total variation norm* of a probability measure on a measurable space $(\Omega, \mathcal{B})$ is given by

$$||P||_{TV} = \sup \sum_m P(\Omega_m)$$

where the supremum is taken over all partitions $\{\Omega_m\}_m$ of Ω into measurable subsets $\Omega_m \in \mathcal{B}$. Given two measures P and Q on $(\Omega, \mathcal{B})$, the *total variation distance* between them is given by

$$||P - Q||_{TV} = 2 \sup_{B \in \mathcal{B}} |P(B) - Q(B)|.$$

The reader is referred to Stokey and Lucas (1989, Section 11.3) for an introduction and applications of this distance in economics. We apply the total variation norm to finite order projections of type spaces.

Definition 6 *Let P be a common prior for a belief-closed type space. We call $||P||_k = ||\tilde{F}^k||_{TV}$ the* order-k total variation norm *(henceforth, the TV-k norm) of P, where $\tilde{F}^k \in \Delta(\Theta \times T^{k-1})$ is the order-k projection of P given in (22).*

Observe that, although the common priors P' and P'' of two belief-closed type spaces are not necessarily defined on the same σ-algebra, their order-k projections, $\tilde{F}'^k$ and $\tilde{F}''^k$, are both in $\Delta(\Theta \times T^{k-1})$. Thus, the *TV-$k$ distance* $||P' - P''||_k = ||\tilde{F}'^k - \tilde{F}''^k||_{TV}$ is well-defined.

We say that a sequence of priors *converges* to a prior in TV-k norm if the convergence is in TV-k distance. The next two lemmas establish properties on the convergence of priors in TV-k norm that are useful later.

Lemma 3 *Fix $k \geq 2$. If a sequence $\{P^m\}_m$ of priors converges to prior P in TV-k norm, then $\text{supp}\tilde{F}^k \subseteq \text{supp}\tilde{F}^{m,k}$ for large enough m, where $\tilde{F}^k$ and $\tilde{F}^{m,k}$ are the order-k projections of P and P^m, respectively.*

Proof. Suppose not. Then, $\text{supp}\tilde{F}^k \setminus \text{supp}\tilde{F}^{m,k} \neq \emptyset$ for infinitely many integers m. Without loss of generality, assume it is true for all m. It follows

$$||\tilde{F}^{m,k} - \tilde{F}^k||_{TV} \geq \min_{(\theta, t^{k-1}) \in \text{supp}\tilde{F}^k} \tilde{F}^k(\theta, t^{k-1}), \ \forall m.$$

Since $\text{supp}\tilde{F}^k$ is finite in number, the right-hand-side of the preceding inequality is well-defined and positive. Consequently, $\{P^m\}_m$ does not converge to P in TV-k norm, which is a contradiction. ■

Lemma 4 *Fix $k \geq 2$. If a sequence $\{P^m\}_m$ of priors converges to prior P in TV-k norm, then for any $\gamma > 0$, there exists an integer m_γ such that*

$$\tilde{F}^{m,k}(\text{supp}\tilde{F}^{m,k} \setminus \text{supp}\tilde{F}^k) < \gamma$$

for all $m > m_\gamma$.

Proof. Since P^m converges to P^k in TV-k norm, for any $\gamma > 0$ there exists an integer m_γ such that

$$|\tilde{F}^{m,k}(B) - \tilde{F}^k(B)| < \gamma$$

for all Borel sets $B \subseteq \Theta \times T^{k-1}$ and for all $m > m_\gamma$. Thus, the proof the lemma is completed by considering $B^m = \text{supp}\tilde{F}^{m,k} \setminus \text{supp}\tilde{F}^k$ for each $m > m_\gamma$. ■

An equilibrium of a game specifies a *type-specific* distribution over the action profiles for each type profile. Given an equilibrium, Kajii and Morris (1997) consider the *equilibrium action distribution* as the average of the type-specific action distributions across all type profiles as specified by the equilibrium (Kajii and Morris, 1997, Definition 2.3, p. 1289). They define a (correlated) Nash equilibrium for a game with complete information to be *robust* against incomplete perturbations, if for every sequence of *elaborations* there is a sequence of BNEs whose equilibrium action distributions converge to the Nash equilibrium action distributions (Kajii and Morris, 1997, Definitions 2.4, 2.5, p. 1289).

It turns out that if a sequence of type spaces converges in TV-2 norm, then it qualifies as a sequence of elaborations in the sense of Kajii and Morris, while convergence in their sense implies convergence in TV-1 norm. Below we extend the robustness notion of Kajii and Morris (1997) to general Bayesian games based on the TV-k norm. First, given a strategy profile σ for a game with type space $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ and common prior P, we use $\tilde{\sigma}^k$ to denote the order-k projection of σ given by

$$\tilde{\sigma}_i^k(t_i^{k-1}) = \frac{\sum_{t_i \in \rho^k(t_i^{k-1})} \text{marg}_{T_i} P(t_i)\sigma_i(t_i)}{\sum_{t_i \in \rho^k(t_i^{k-1})} \text{marg}_{T_i} P(t_i)} \tag{24}$$

for all $t_i^{k-1} \in T_i^{k-1}$ such that $\rho^k(t_i^{k-1}) \neq \emptyset$ and for all i.

Definition 7 *Let $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ be a type space with common prior P and $\mathcal{C}$ a collection of common-prior type spaces with $\mathcal{T} \in \mathcal{C}$. Let $\Gamma = \{\{A_i, u_i\}_{i=1}^n, \mathcal{T}\}$ be a Bayesian game. We say that a BNE σ^* for Γ is* $\mathcal{C}$-strong-k robust *if for any sequence of type spaces $\mathcal{T}^m \in \mathcal{C}$ with common prior P^m converging to P in* TV-k norm*, there is a sequence of BNEs σ^{*m} for $\Gamma^m = \{\{A_i, u_i\}_{i=1}^n, \mathcal{T}^m\}$ such that*

$$\lim_{m \longrightarrow \infty} \Bigg[\sum_{t^{k-1} \in \text{supp}\tilde{F}^k} ||\tilde{\sigma}^{*m,k}(t^{k-1}) - \tilde{\sigma}^{*k}(t^{k-1})|| \text{marg}_{T^{k-1}} \tilde{F}^{m,k}(t^{k-1})$$

$$+\sum_{t^{k-1}\in \mathrm{supp}\tilde{F}^{m,k}\setminus \mathrm{supp}\tilde{F}^{k}} ||\tilde{\sigma}^{*m,k}(t^{k-1})|| \mathrm{marg}_{T^{k-1}}\tilde{F}^{m,k}(t^{k-1})\Bigg] = 0,$$

where $||\cdot||$ *denotes the Euclidean norm on* $\Re^A$ *and* $\tilde{F}^{m,k}$ *the order-k projection of* P^m*.*

In the preceding definition, the closeness between a BNE of the target prior and that of a perturbed prior is measured by the average of the *type-specific* Euclidean distances between the order-k projections of these two equilibrium strategies, with the average taken with respect to projections of the perturbed priors. By Lemmas 3 and 4, the support of each perturbed prior is larger than that of the target under the TV-k norm convergence, but the difference vanishes as the perturbation reduces. Thus, weighting the type-specific distance by the perturbed priors is appropriate. In comparison, Kajii and Morris (1997) measure the closeness by the distance in the *max norm* between the equilibrium action distributions of the equilibrium strategies. For non-degenerate incomplete information games, however, two different strategy profiles may specify the same action distribution. Thus, the limit of a sequence under their convergence may be consistent with a class of strategy profiles, many of which are not BNEs. Of course, in the degenerate case of complete information, such limit leads to a unique strategy profile as there is only type for each player.

Given a Bayesian game Γ with an order-k non-redundant type space $\mathcal{T}$, it follows from (19) and (24) that for any strategy profile for Γ,

$$\tilde{\sigma}^k(\tau^h(t)) = \sigma(t),\ \ t \in T, h \geq k.$$

Consequently, σ is a BNE for Γ if and only if $\tilde{\sigma}^h$ is a BNE for $\tilde{\Gamma}^h = \{(A_i, u_i)_{i=1}^n, \tilde{\mathcal{T}}^h\}$ for all $h \geq k$.

Definition 8 *We say that type space* $\mathcal{T} = \{T_i, \pi_i\}_{i\in N}$ *is* self-closed *in another type space* $\mathcal{T}' = \{T_i', \pi_i'\}_{i\in N}$ *if* $T_i \subseteq T_i'$ *and* $\pi_i(t_i) = \pi_i'(t_i)$ *for all* $t_i \in T_i$ *and for all* i*.*

The self-closedness of type space $\mathcal{T}$ in $\mathcal{T}'$ means that the players' beliefs are preserved as their type sets are enlarged. Thus, an implication of the self-closedness is that, fixing the players' action sets, any BNE with type space $\mathcal{T}$ remains to be a component of a BNE as the players' type sets are enlarged. This allows us to expand a BNE from a smaller type space to a BNE for a larger one, as shown in the following lemma.

Lemma 5 *Let $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ be a non-redundant type space, $\Gamma = \{(A_i, u_i)_{i=1}^n, \mathcal{T}\}$ a Bayesian game, and σ^* a BNE for Γ. Suppose $\mathcal{T}' = \{T'_i, \pi'_i\}_{i=1}^n$ is another type space such that $\mathcal{T}$ is self-closed in $\mathcal{T}'$. Then, there is a BNE σ'^* for $\Gamma' = \{(A_i, u_i)_{i=1}^n, \mathcal{T}'\}$ such that*

$$\sigma'^*(t) = \sigma^*(t), \ \forall t \in T.$$

Proof. For $i \in N$ and $t_i \in T'_i$, set

$$X_i = \prod_{t_i \in T'_i} \Delta(A_i) \text{ and } \hat{X}^i = \prod_{t_i \in T'_i \setminus T_i} \Delta(A_i)$$

and for $t_i \in T'_i \setminus T_i$ and $\hat{\sigma}_{-i} \in \hat{X}_{-i}$, define $\hat{BR}_i(t_i)[\hat{\sigma}_{-i}]$ to be the set of $\hat{\sigma}_i(t_i) \in \Delta(A_i)$ such that

$$E_{t_i}[u_i((\hat{\sigma}_i(t_i), f_{-i}(\hat{\sigma}_{-i}, \sigma^*_{-i}), \theta)] = \max_{a_i \in A_i} E_{t_i}[u_i((a_i, f_{-i}(\hat{\sigma}_{-i}, \sigma^*_{-i})), \theta)],$$

where $f_{-i}(\hat{\sigma}_{-i}, \sigma^*_{-i}) = (f_j(\hat{\sigma}_j, \sigma^*_j))_{j \neq i} \in X_{-i}$ such that

$$f_j(\hat{\sigma}_j, \sigma^*_j)(t_j) = \begin{cases} \sigma^*_j(t_j) & if \ t_j \in T_j; \\ \hat{\sigma}_j(t_j) & if \ t_j \in T'_j \setminus T_j. \end{cases}$$

Define $\hat{BR} : \hat{X} \longrightarrow \hat{X}$ by

$$\hat{BR}(\hat{\sigma}) = \prod_{i \in N} \prod_{t_i \in T'_i \setminus T_i} \hat{BR}_i(t_i)[\hat{\sigma}_{-i}].$$

By construction, $\hat{BR}$ together with $\hat{X}$ satisfies all the conditions for Kakutani's fixed point theorem. Thus, $\hat{BR}$ has a fixed point which we denote by $\hat{\sigma}^*$. Now define $\sigma'^* \in X$ by

$$\sigma'^*_i(t_i) = \begin{cases} \sigma^*_i(t_i) & if \ t_i \in T_i, \\ \hat{\sigma}^*_i(t_i) & if \ t_i \in T'_i \setminus T_i, \end{cases}$$

for all i. By construction, $\sigma'^*_i(t_i)$ is a best response against σ'^*_{-i} for all $t'_i \in T'_i \setminus T_i$. On the other hand, by the self-closedness and the assumption that σ^* is a BNE for Γ, $\sigma'^*(t_i)$ is also a best response against σ'_{-i} for $t_i \in T_i$. ■

We now apply Lemma 5 to establish the robustness of BNEs.

Theorem 2 *Let $\mathcal{T} = \{T_i, \pi_i\}_{i=1}^n$ be an order-k non-redundant type space with common* *prior P and $\Gamma = \{(A_i, u_i)_{i=1}^n, \mathcal{T}\}$ be a Bayesian game. Then, all BNEs for Γ are $\mathcal{C}$-strong-(k+1) robust for any collection $\mathcal{C}$ of order-k common-prior type spaces.*

Proof. Let σ^* be a BNE for Γ. Let $\{P^m\}_m \subseteq \mathcal{C}$ be a sequence of common priors converging to P in TV-(k+1) norm. By Lemma 3, we may assume that for all m,

$$\text{supp}\tilde{F}^{k+1} \subseteq \text{supp}\tilde{F}^{m,k+1}, \tag{25}$$

where supp$\tilde{F}^{k+1}$ and $\tilde{F}^{m,k+1}$ are the order-(k+1) projections of P and P^m, respectively. Since both $\mathcal{T}$ and $\mathcal{T}^m$ are of order k, it follows from Theorem 1 that $\tilde{F}^{k+1}$ and $\tilde{F}^{m,k+1}$ are also of order k.

Now putting the consistency together with (25), we have

$$\tilde{F}^{k+1}(\theta, t_{-i}^{k-1} | t_i^k) = \phi^k(\theta, t_{-i}^{k-1}) = \tilde{F}^{m,k+1}(\theta, t_{-i}^{k-1} | t_i^k), \forall t^k \in \text{supp}\tilde{F}^{k+1}, \forall i,$$

where $\phi_i^k(\cdot)$ is the k-th order belief of player i in t_i^k about θ and the other players' order-(k-1) belief hierarchies. Since $\tilde{F}^{k+1}$ and $\tilde{F}^{m,k+1}$ are of order k, by (25) and the preceding equation,

$$\text{supp}\tilde{F}^k \subseteq \text{supp}\tilde{F}^{m,k} \text{ and } \tilde{F}^k(\theta, t_{-i}^{k-1} | t_i^{k-1}) = \tilde{F}^{m,k}(\theta, t_{-i}^{k-1} | t_i^{k-1}), \forall (\theta, t^{k-1}) \in \text{supp}\tilde{F}^k, \forall i,$$

where $\tilde{F}^k$ and $\tilde{F}^{m,k}$ are the order-k projections P and P^m, respectively. This shows that $\tilde{\mathcal{T}}^k$ (the order-k projection of $\mathcal{T}$) is self-closed in $\tilde{\mathcal{T}}^{m,k}$ (the order-k projection of $\mathcal{T}^m$).

Since $\mathcal{T}$ is non-redundant and of order-k, $\rho_i^k(t_i^{k-1})$ is singleton for all $t_i^{k-1} \in \tau_i^k(T_i)$ and for all i. Hence, by (24), $\tilde{\sigma}^{*k}(\tau^k(t)) = \sigma^*(t)$ for all $t \in T$ which implies that $\tilde{\sigma}^{*k}$ is a BNE for $\tilde{\Gamma}^k = \{(A_i, u_i)_{i \in N}, \tilde{\mathcal{T}}^k\}$. Next, by Lemma 5, there exists a BNE $\sigma^{*m,k}$ for $\tilde{\Gamma}^{m,k} = \{\{A_i, u_i\}_{i \in N}, \tilde{\mathcal{T}}^{m,k}\}$ such that

$$\sigma^{*m,k}(\tau^k(t)) = \tilde{\sigma}^{*k}(\tau^k(t)), \forall t \in T. \tag{26}$$

Now set $\sigma^{*m} = \sigma^{*m,k}(\tau^k(t))$ for $t \in T^m$. Then, since $\mathcal{T}^m$ is of order-k and $\sigma^{*m,k}$ is a BNE for $\tilde{\Gamma}^{m,k}$, σ^{*m} is a BNE for Γ^m and $\tilde{\sigma}^{*m,k} = \sigma^{*m,k}$. By (26),

$$\sum_{t^{k-1} \in \text{supp}\tilde{F}^k} ||\tilde{\sigma}^{*m,k}(t^{k-1}) - \tilde{\sigma}^{*k}(t^{k-1})|| \text{marg}_{T^{k-1}} \tilde{F}^{m,k}(t^{k-1})$$

$$+ \sum_{t^{k-1} \in \text{supp}\tilde{F}^{m,k} \setminus \text{supp}\tilde{F}^{k}} ||\tilde{\sigma}^{*m,k}(t^{k-1})|| \text{marg}_{T^{k-1}} \tilde{F}^{m,k}(t^{k-1})$$

$$= \sum_{t^{k-1} \in \text{supp}\tilde{F}^{m,k} \setminus \text{supp}\tilde{F}^{k}} ||\tilde{\sigma}^{*m,k}(t^{k-1})|| \text{marg}_{T^{k-1}} \tilde{F}^{m,k}(t^{k-1}).$$

By Lemma 4,

$$\lim_{m \longrightarrow \infty} \sum_{t^{k-1} \in \text{supp}\tilde{F}^{m,k} \setminus \text{supp}\tilde{F}^{k}} ||\tilde{\sigma}^{*m,k}(t^{k-1})|| \text{marg}_{T^{k-1}} \tilde{F}^{m,k}(t^{k-1}) = 0.$$

This completes the proof. ■

A special interesting case is that of complete information.

Corollary 4 *Every Nash equilibrium of a complete information game is $\mathcal{C}$-strong-3 robust for any collection $\mathcal{C}$ of order-2 common-prior type spaces.*

Theorem 2 offers an interesting extension of the result by Monderer and Samet (1989) that all NEs for a complete-information game are robust, if some common-1 belief condition is satisfied.[10] It also implies that a genuine selection among BNEs with the robustness approach is only possible if the required strong convergence is weaker than full-order.

Another potential source for genuine selection is by allowing for higher-order perturbations under TV norm convergence. Because of non-trivial merging of higher-order types, best responses based on lower-order projections are not necessarily best responses in the original game, making a similar construction as in Theorem 2 impossible.

4 Applications

4.1 Generic Failure of Coase Cojecture with Order-2 Consistent Priors

Consider an one-sided repeated offers bargaining game between the buyer and seller of an indivisible good. At every period the seller makes an offer that the buyer can either accept or reject. There is a common discount factor. One-sided private information with common knowledge of positive gains from trade is known to yield no *delay* for

[10]See Kajii and Morris, 1997, p. 1287 and footnote 4 for further discussions.

this simple bargaining procedure, when the information structure is modeled by a naive type space. Here, delay means that as offers are increasingly frequent, the actual time for reaching agreement does not goe to zero (see Fudenberg, Levine, and Tirole (1985) and Gul, Sonnenschein, and Wilson (1986)). This convergence is known as the *Coase conjecture or property.* The underlying information structures are naive private-value type spaces, or order-1 private-value consistent priors.[11]

Recently, Feinberg and Skrzypacz (2005) showed that introducing higher-order uncertainty dramatically changes the outcomes. Specifically, the presence of the buyer's (second-order) uncertainties about seller's (first-order) beliefs can lead to the failure of the Coase property, if the seller is considered by the buyer to entertain a *certainty belief.* Below we demonstrate that the information structures used in their paper can be modeled as order-2 consistent priors.

In their setting, the buyer's value for the indivisible good is either $\theta_1 = \theta_{1h}$ or θ_{1l} with $\theta_{1h} > \theta_{1l} > 0$, but the value of the good to the seller is $\theta_2 = 0$. The buyer's private information is his value of the good. Though the seller's value is commonly known, he has as private information two order-1 beliefs, ϕ_2^A and ϕ_2^B, that assign probabilities $0 < p_h^B < p_h^A$ to the buyer having high value as in our Example 1 ($P^A(\theta_{1h}, \theta_2) = p_h^A$ and $P^B(\theta_{1h}, \theta_2) = p_h^B$).[12] Consequently, the buyer's types can be identified by his payoff types θ_{1h} and θ_{1l} of the good, while the seller's are by his two order-1 beliefs. With these notation, the information structure (6) in their paper, which we denote by F_6, is the same as F^λ in Example 1 with

$$\lambda = \frac{p_h^B(1-\beta)}{p_h^A\beta + p_h^B(1-\beta)}. \tag{27}$$

Feinberg and Skrzypacz also consider an information structure, in which the seller has an *informed* type, in the sense that he is sure that the buyer has the high value. This is information structure (1) in their paper modeled as having $p_h^A = 1$, which we denote by F_1. It is the same as F^λ with

$$\lambda = \frac{p_h^B(1-\beta)}{\beta + p_h^B(1-\beta)}. \tag{28}$$

Feinberg and Skrzypacz show that, imposing the *intuitive criterion*, the Coase property

[11]See Feinberg and Skrzypacz (2005) for further discussions of models based on naive type spaces for studying the Coase property.

[12]These two probabilities are denoted by α^0 and α^N with $\alpha^N < \alpha^0$ in Feinberg and Skrzypacz (2005).

is valid with information structure F_6, but fails with F_1 in that delay in reaching agreement occurs with positive probability.

We say that an order-2 prior F has "certainty belief" for player i if there is a pair $(\theta, \phi_i) \in \operatorname{supp} \operatorname{marg}_{\Theta \times \Delta(\Theta_{-i})} F$ such that $\phi_i = \delta_{\theta_{-i}}$. It can be verified that a convex combination of consistent priors is "free of certainty belief types" if and only if both priors are free of certainty belief types. Thus, following Heifetz and Neeman (2006), having a certainty belief may be regarded as a *generic* property.[13] On the other hand, Feinberg and Skrzypacz (2005) state in their conclusion that delay occurs with an order-2 consistent prior based on their chosen private-value payoff environment if and only if the prior has a certainty belief type for the seller. It follows that the Coase property fails genericly with respect to order-2 consistent priors. This provides an answer to the question, asked at the end of their paper, of whether delay is generic in a general type space.

Note that, by (27) and (28), $F_6 \longrightarrow F_1$ as $p_h^A \longrightarrow 1$. Thus with respect to the weak topology, the Coase property is discontinuous in consistent priors. However, considering that F_1 is generic while F_6 is not, within order-2 consistent priors, this discontinuity indicates that weak convergence is too weak. On the other hand, for F^λ to converge in total-variation norm to a common prior with certainty belief, say, F_1, it is necessary that $\operatorname{supp} F_1 \subset \operatorname{supp} F^\lambda$ for λ close to the number in (28), as shown in Lemma 3. Hence, F_6 does not converge to F_1 in our TV-2 norm and by Feinberg and Skrzypacz's (2005) results, this implies the failure of the Coase property in any 'strong' neighborhood of F_1. Consequently, there is no evidence that the Coase property is not continuous while the failure of the Coase property is continuous with strong convergence.

4.2 The Type Spaces of Global Games Are Order 2

Perturbations of a complete information type space associated with a game can result in an equilibrium selection as forcefully demonstrated by Carlsson and van Damme

[13] By our Proposition 1, the family of consistent order-2 priors is a convex family. Since a convex combination of two priors is free of certainty beliefs for a player if and only if both of these two priors are free of certain beliefs for the player, it follows that the subset of consistent order-2 priors free of certainty beliefs for a player is contained in a proper face of the family of all consistent order-2 priors. In this geometric sense, the subset of consistent order-2 priors free of certainty beliefs (resp. containing a certainty belief) for a player is regarded non-generic (resp. generic) as in Heifetz and Neeman (2006).

(1993).[14] They consider a class of perturbations of the following form:

$$x_i = \theta + \epsilon\eta_i, \tag{29}$$

where θ is the payoff-relevant parameter, $\epsilon > 0$ is a constant, and η_1 and η_2 are independent random variables. From (29), it follows that as ϵ approaches to zero, x_i converges to θ, meaning that information about θ is complete in the limit. Since η_1 and η_2 are independent and θ is constant, so are x_1 and x_2. Consequently, the distributions of x_1 and x_2 are the marginal distributions of a common prior for (θ, x_1, x_2). In addition, the posterior distribution of θ given x_i is uniquely identifiable. Equivalently, player i's order-1 beliefs are completely separable. Thus, given $\epsilon > 0$, a similar reasoning as in the proof of our Theorem 1 shows that the corresponding incomplete-information type spaces as implied by (29) are of order 2. Carlsson and van Damme show that in the stag-hunt game only the risk-dominant Nash equilibrium is robust against this class of perturbations.

Morris and Shin (2000) consider the same form of perturbations as (29), except that they assume that θ itself is also uncertain. A specification of their perturbations as in Weinstein and Yildiz (2007) requires that θ, η_1, η_2 be independently and normally distributed with $\theta \sim N(\mu_\theta, \sigma_\theta^2)$ and $\eta_i \sim N(0, 1)$ for $i = 1, 2$. The normality and independence together with (29) implies

$$x_i \sim N(\mu_\theta, \sigma_\theta^2 + \epsilon^2), \ i = 1, 2. \tag{30}$$

Furthermore, (θ, x_1, x_2) jointly have a multivariate normal distribution:

$$(\theta, x_1, x_2) \sim N_3(\mu, \Sigma), \ \mu = \begin{pmatrix} \mu_\theta \\ 0 \\ 0 \end{pmatrix}, \ \Sigma = \begin{pmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 \\ \sigma_\theta^2 & \sigma_\theta^2 + \epsilon^2 & \sigma_\theta^2 \\ \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 + \epsilon^2 \end{pmatrix}. \tag{31}$$

By (30) and (31), the distributions of x_1 and x_2 are marginal distributions of $N_3(\mu, \Sigma)$. The probability measure P over the Borel sets of $\Re^3$ as determined by this multivariate normal distribution is the common prior for (θ, x_1, x_2) (see Weinstein and Yildiz, 2007, p. 392). Furthermore, the posterior distribution of θ given x_i is normal (see DeGroot,

[14]Their work promoted fruitful research on *global games*. See Morris and Shin (2003) for a survey on global games and applications.

1975, p. 269):

$$\theta|x_i \sim N\left(\frac{\epsilon^2\mu_\theta + \sigma_\theta^2 x_i}{\sigma_\theta^2 + \epsilon^2}, \frac{\sigma_\theta^2\epsilon^2}{\sigma_\theta^2 + \epsilon^2}\right). \tag{32}$$

For $i = 1, 2$, let $\phi_i(x_i)$ denote the probability measure about θ with density function (32). Now, define $f : \Re^3 \longrightarrow \Re \times \Pi_{i=1}^2 \Delta(\Re)$ by

$$f(\theta, x_1, x_2) = (f_1(\theta), f_1(x_1), f_2(x_2)) = (\theta, \phi_1(x_1), \phi_2(x_2)),\ \forall(\theta, x_1, x_2) \in \Re^3. \tag{33}$$

By (32), f is 1-1. Denote by f^{-1} the inverse of f. Next, set

$$F(A) = P \circ f^{-1}(A \cap f(\Re^3)),\ A \subseteq \Re \times \Pi_{i=1}^2 \Delta(\Re). \tag{34}$$

Then, since f is measurable with the weak topology on $\Delta(\Re)$ and product topology on $\Re \times \Pi_{i=1}^2 \Delta(\Re)$, F is a probability measure over $\Re \times \Pi_{i=1}^2 \Delta(\Re)$ with $(\theta, \phi_1(x_1), \phi_2(x_2)) \in \text{supp}F$, for all $(\theta, x_1, x_2) \in \Re^3$. By (30)-(34) and the 1-1 property of f, for all $(\theta, x_1, x_2) \in \Re^3$,

$$\begin{aligned} F(\theta|\phi_i(x_i)) &= P(\theta|f_i^{-1}(x_i)) \\ &= P(\theta|x_i) \\ &= \phi_i(x_i)[\theta] \end{aligned}$$

This shows that F is a consistent order-2 prior. Since f is a type isomorphism, P is of order 2 as well.[15] It follows that the type spaces associated with the proto-type models of Global Games are of order 2.

4.3 Correlation Devices and Strategic Relevance of Redundancy

Ely and Peski (2006) point out that 'redundancy' may have strong strategic implications. In particular, they show that strategic relevant primitives ought to include information about correlations between beliefs across players not captured in the formulation of the universal type space. Liu (2005, 2009) investigates the relationships between redundant type spaces and their non-reducible forms. His analysis shows that, with respect to BNE, the added redundancy can be interpreted as correlation devices to generate correlated equilibrium with incomplete information.

[15]Morris and Shin (2000) show that there is a subclass of perturbations within this family of order-2 priors that selects the Pareto-dominant NE. Weinstein and Yildiz (2007) have a sharp result that shows that every (interim) rationalizable strategy can be sustained with weak convergence in the universal type space.

Our finite-order consistent prior approach is by construction redundancy-free. However, the model can be extended to deal with the issue of "strategic redundancy" in finite-order type spaces. For example, in an order-2 environment, let us enlarge the set of fundamental uncertainty Θ by adding a space $S = \Pi S_i$ of profiles of private signals $s_i \in S_i$. Assume nature jointly draws according to common prior $F \in \Delta(\Theta \times S \times \Pi_{i=1}^n \Delta(\Theta \times S_{-i}))$ an element $\theta \in \Theta$, a profile $s \in S = \Pi_{i=1}^n S_i$ of private signals, and a profile $\psi = (\psi_1, \psi_2, \cdots, \psi_n)$ of beliefs in $\Pi_{i=1}^n \Delta(\Theta \times S_{-i})$. Next, nature informs player i of belief $\psi_i \in \Delta(\Theta \times S_{-i})$ as well as signal s_i. The pair (s_i, ψ_i) summarizes player i's private information. Player i can derive a posterior belief by updating prior F using his private information (s_i, ψ_i) via the Bayes Rule. Accordingly, F is consistent if for all i

$$\psi_i(\theta, s_{-i}) = F(\theta, s_{-i}|s_i, \psi_i), \forall (s_i, \psi_i) \in \operatorname{supp} \operatorname{marg}_{S_i \times \Delta(S_{-i})} F \tag{35}$$

Given a prior $F \in \Delta(\Theta \times S \times \Pi_{i=1}^n \Delta(\Theta \times S_{-i}))$, we can derive an order-2 prior $\operatorname{marg}_{\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})} F$ based on Θ only by (canonical) projection:

$$\operatorname{marg}_{\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})} F(\theta, \phi) = \sum_{\substack{(\theta,\psi):(\theta,s_i,\psi_i) \in \operatorname{supp} \operatorname{marg}_{S_i \times \Delta(S_{-i})} F, \\ \operatorname{marg}_\Theta \psi_i = \phi_i, \forall i}} F(\theta, s, \psi), \tag{36}$$

for all $(\theta, \phi) \in \Theta \times \Pi_{i=1}^n \Delta(\Theta)$. It turns out that such a projection preserves consistency.

Lemma 6 *If* $F \in \Delta(\Theta \times S \times \Pi_{i=1}^n \Delta(\Theta \times S_{-i}))$ *is consistent, then* $\operatorname{marg}_{\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})} F$ *is a consistent prior in* $\Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta))$.

Proof. Fix i. For any $(\theta, \phi_i) \in \Theta \times \Delta(\Theta)$ and for any $\psi_i \in \Delta(\Theta \times S_{-i})$ with $(\theta, s_i, \psi_i) \in \operatorname{supp} \operatorname{marg}_{\Theta \times S_i \times \Delta(\Theta \times S_{-i})} F$, it follows from (35) that $\phi_i = \operatorname{marg}_\Theta \psi_i$ implies

$$\phi_i(\tilde{\theta}) = \sum_{s_{-i} \in S_{-i}} \psi_i(\tilde{\theta}, s_{-i}) = \frac{F(\tilde{\theta}, s_i, \psi_i)}{F(s_i, \psi_i)}.$$

Thus, exploiting the fact that $\frac{a}{b} = \frac{c}{d}$ implies $\frac{a+c}{b+d}$,

$$\phi_i(\tilde{\theta}) = \frac{\sum_{\substack{\psi:\ (\tilde{\theta}, s_i, \psi_i) \in \operatorname{supp} \operatorname{marg}_{S_i \times \Delta(S_{-i})} F, \\ \operatorname{marg}_\Theta \psi_i = \phi_i}} F(\tilde{\theta}, s_i, \psi_i)}{\sum_{\substack{\psi:\ (\tilde{\theta}, s_i, \psi_i) \in \operatorname{supp} \operatorname{marg}_{S_i \times \Delta(S_{-i})} F, \\ \operatorname{marg}_\Theta \psi_i = \phi_i}} F(s_i, \psi_i)}.$$

The proof is completed by combining the above equation with (36). ■

Consequently, the signal-enriched extension of consistent order-2 prior serves as an implicit device to build in (potentially strategically relevant) redundancy into Bayesian games, preserving the same projection image. In what follows we demonstrate, using the basic motivating example in Ely and Peski (2006) and Liu (2009), that this extension can indeed accommodate correlation of equilibrium strategies not feasible in redundancy-free models.

Assume $\Theta = \{\theta_1, \theta_2\}$, where θ_1 and θ_2 are associated with the payoff matrices in a 2-person game as in Table 2 below.

	L	R			L	R
U	1, 1	0, 0		U	0, 0	1, 1
D	0, 0	1, 1		D	1,1	0, 0
	(θ_1)				(θ_2)	

Table 2: A Coordination Game with Common-Value Uncertainty.

Assume there is a common naive prior with which the probability for having θ_1 is $p \geq 1/2$. Then, the type space has one type for each player given by the first-order belief that assigns probability p to θ_1 and $1-p$ to θ_2. It is straightforward to show that the only BNE outcome has each player receive payoff p. Let $S_1 = \{s^D, s^U\}$ and $S_2 = \{s^L, s^R\}$ be the spaces of private signals for player 1 and player 2. Set $S = S_1 \times S_2$ and define $\Psi \in \Delta(\Theta \times S)$ by

$$\Psi(\theta, s) = \begin{cases} \frac{p}{2}, & (\theta, s) = (\theta_1, s^U, s^L) \text{ or } (\theta_1, s^D, s^R), \\ \\ \frac{1-p}{2}, & (\theta, s) = (\theta_2, s^U, s^R) \text{ or } (\theta_2, s^D, s^L). \end{cases}$$

Next, for $s_i \in S_i$ and $i = 1, 2$, define $\psi^{s_i} \in \Delta(\Theta \times S_{-i})$ by

$$\psi_i^{s_i}(\theta, s_{-i}) = \Psi(\theta, s_{-i}|s_i).$$

Finally, define $F \in \Delta(\Theta \times S \times \Pi_{i=1}^2 \Delta(\Theta \times S_{-i}))$ by

$$F(\theta, s, \psi) = \Psi(\theta, s)\delta_{\psi^s}(\psi),$$

where $\psi^s = (\psi_1^{s_1}, \psi_2^{s_2})$ for $s = (s_1, s_2) \in S$. By construction, F satisfies consistency (35). Furthermore, $\psi_i^{s_i}(\theta_1, S_{-i}) = p$ for all $s_i \in S_i$ and for all i, implying that the two different types of each player share the same first-order belief about $\theta \in \Theta$. By the construction of F, it is also straightforward to show that all their higher-order beliefs are identical. In fact, the canonical projection yields the prior with a single-element support. However, given F, the strategies $a_1 = (a_1(s^U), a_1(s^D)) = (U, D)$ $a_2 = (a_2(s^L), a_2(s^R)) = (L, R)$ and constitute a BNE that yields payoff 1 for each player as a result of perfect correlation via private signals $s \in S$. More generally, Liu (2009) shows that given a redundant type space, a"signal-enriched" non-redundant type space can be constructed such that a type morphism between the two exits. Our construction can be considered as a special case of his.[16]

5 Conclusion

Belief-closed subsets of the universal type space do not necessarily need the whole extent of the associated infinite belief hierarchies for Bayesian game applications. Naive type spaces need only the first order of the belief hierarchies. Complete information type spaces turn out to require exactly the first two orders of the belief hierarchies. As demonstrated in the present paper, order-2 type spaces seem to possess sufficiently rich structures, in order to fundamentally challenge many stunning theoretical insights gained using naive type spaces.

Walking up the order from the ground, we interpreted finite-order projections as imperfect realizations of an infinite order type space and characterize the impact on the solutions of the projected Bayesian games. We also uncovered the embedding relationship between finite-order heterogeneous-prior type spaces and some one-order higher common-prior ones.

The finite-order approach introduced here also sheds lights on the recent discussion of strategic topology. Restricted to order-k type spaces, uniform-weak, strategic, and weak topologies are all equivalent. This suggests that the discussion of critical types by Ely and Peski (2007) necessarily implicates infinite orders. As we observed, discontinuity of NEs already occur in order-2 type spaces both in global games and the model by Feinberg and Skrzypacz (2005), strategic topology seems not to have much impact if the solution concept changes from interim correlated rationalizability to BNE.

With the finite-order total-variation norm defined, we offered a new approach to analyzing the robustness of BNE. The initial insight is that all BNEs for a game with

[16]For a partition model, Liu (2005) provides an abstract condition to generate redundant equivalents to a non-redundant universal type space.

an order-k type space is robust against order-k perturbations under the TV-k norm. Allowing for more perturbations by either relaxing the convergence requirement or the order requirement for the perturbed spaces, thus a harder test, may yield meaningful stricter selections.

References

[1] Aumann, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1-18.

[2] Aumann, R. (1998): "Common priors: A reply to Gul," *Econometrica*, 66, 929-938.

[3] Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771-1813.

[4] Bergemann, D. and J. Välimäki (2007): "Information in Mechanism Design," in R. Blundell, W. Newey, and T. Persson (eds), *Proceedings of the 9th World Congress of Econometric Society*, Cambridge University Press, pp. 186-221.

[5] Billingley, P. (1995): *Probability and Measure*. John Wiley & Sons.

[6] Brandenburg, A. and E. Dekel (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189-198.

[7] Chen, Y-C., A. D. Tillio, E. Faingold, and S. Xiong (2009): "Uniform Topologies on Types," mimeo.

[8] Chung, K. and J. Ely (2003): "Implementation with Near-Complete Information," *Econometrica*, 71, 857-871.

[9] Crémer, J. and R. McLean (1985): "Optimal selling strategies under Uncertainty for a Discriminating Monopolist when Demands Are Interdependent," *Econometrica*, 53, 345-361.

[10] Crémer, J. and R. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247-1258.

[11] DeGroot, M. H. (1975): *Probability and Statistics,* Addison-Wesley Publishing Company, Inc., 139-52.

[12] Dekel, E., D., Fudenberg, and S. Morris (2006): "Topologies on Types," *Theoretical Economics*, 1, 275-309.

[13] Ely, J. and M. Peski (2006): "Hierarchies of Belief and Interim Rationalizability," *Theoretical Economics*, 1, 19-65.

[14] Ely, J. and M. Peski (2007): "Critical Types," mimeo.

[15] Feinberg, Y. and A. Skrzypacz (2005): "Uncertainty about Uncetainty and Delay in Bargaining," *Econometrica*, 73, 69-91.

[16] Fudenberg, D., D. Levine, and J. Tirole (1985): "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in A. Roth (ed), *Game-Theoretic Models of Bargaining*, Cambridge: Cambridge University Press, pp. 73-98.

[17] Gul, F., H. Sonnenschiein, and R. Wilson (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, 155-190.

[18] Gul, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66, 923-927.

[19] Harsanyi, J. (1967-1968): "Games with Incomplete information Played by Bayesian Players," *Management Sciences*, 14, 159-182, 320-334, 486-502.

[20] Heifetz, A and D. Samet (1998): "Topology-Free Typology of Beliefs," *Journal of Economic Theory*, 82, (1998), 324-341.

[21] Heifetz, A. and Z. Neeman (2006): "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design," *Econometrica*, 74, 213-233.

[22] Kajii, A. and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283-1310.

[23] Liu, Q. (2005): "Representation of Belief Hierarchies in Games with Incomplete Information," mimeo, University of Pennsylvania.

[24] Liu, Q. (2009): "On Redundant Types and Bayesian Formulation of Incomplete Information," *Journal of Economic Theory*, 144, 115-145.

[25] Mertens, J. F. and S. Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information." *International Journal of Game Theory*, 10, 619-632.

[26] Monderer, D., and D. Samet (1989): "Approximating Common Knowledge with Common Beliefs, *Games and Economic Behavior*, 1, 170-190.

[27] Morris, S. (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327-1347.

[28] Morris, S. (1995): "The Common Prior Assumption in Economic Theory." *Economics and Philosophy*, 11, 227-253.

[29] Morris, S. and H. S. Shin (2000):"Rethinking Multiple Equilibria in Macroeconomics," *NBER Macroeconomics Annual*, 139161. MIT Press.

[30] Morris, S. and H. S. Shin (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics* (Proceedings of the Eighth World Congress of the Econometric Society), edited by M. Dewatripont, L. Hansen and S. Turnovsky. Cambridge, England: Cambridge University Press.

[31] Myerson, R. (2004): "Comments on "Games with Incomplete Information Played by 'Bayesian, Players", I-III": Harsanyi's Games with Incomplete Information," *Management Sciences*, 50, 1818-1824.

[32] Neeman, Z. (2004): "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory*, 117, 55-77.

[33] Rubinstein, A., (1989): "The electronic mail game: Strategic Behavior under Almost Common Knowledge," *American Economic Review*. 79, 385-391.

[34] Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989): *Recursive Methods in Economic Dynamics*, Canbridge: Harvard University Press.

[35] Weinstein, J. and M. Yildiz (2007a): "A Structure Theorem for Rationality with Application to Robust Predictions of Refinements," *Econometrica*, 75, 365-400.

[36] Weinstein, J. and M. Yildiz (2007b): "Impact of higher-order uncertainty," *Games and Economic Behavior*, 60: 200-212.

Appendix: Necessary and Sufficient Conditions Embedding Heterogeneous Priors

Proposition 5 *Let $\mathcal{P} = (P_1, P_2, \cdots, P_n)$ be a profile of naive priors with a common support. Then, $\mathcal{P}$ can be embedded in an order-2 prior if and only if there exist a subset $\mathcal{F} \subseteq \Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})$ and a function $h_i : \Theta_i \times \Delta(\Theta_{-i}) \longrightarrow \Re_+$ for each i such that*

(a) $(\theta, P(\cdot|\theta)) \in \mathcal{F}$ for all $\theta \in \Theta$ such that $\theta \in \mathrm{supp} P_i$;

(b) $h_i(\theta_i, \phi_i) > 0$ if and only if $(\theta_i, \phi_i) \in \mathrm{proj}_{\Theta_i \times \Delta(\Theta_{-i})} \mathcal{F}$;

(c) For all i and for all $c_i : \Theta \times \Delta(\Theta_{-i}) \longrightarrow \Re$ with $c_i(\theta, \phi_i) = 0$ for $(\theta, \phi_i) \notin \mathrm{proj}_{\Theta \times \Delta(\Theta_{-i})} \mathcal{F}$,

$$\max_{(\tilde{\theta}, \tilde{\phi}) \in \mathcal{F}} \sum_{i=1}^n c_i(\tilde{\theta}, \tilde{\phi}_i) \geq \sum_{i=1}^n \sum_{(\theta, \phi_i) \in \mathrm{proj}_{\Theta_i \times \Delta(\Theta_{-i})} \mathcal{F}} c_i(\theta, \phi_i) \phi_i(\theta_{-i}) h_i(\theta_i, \phi_i).$$

Proof. Suppose $\mathcal{P} = (P_1, \cdots, P_n)$ can be embedded in an order-2 prior F. Set

$$\mathcal{F} = \mathrm{supp} F \text{ and } \mathcal{F}_i = \mathrm{proj}_{\Theta \times \Delta(\Theta_i)} \mathcal{F}, \ \forall i.$$

For all i and $(\theta, \phi_i) \in \mathcal{F}_i$, define $I_{\theta, \phi_i} : \mathcal{F} \longrightarrow \Re$ by

$$I_{\theta, \phi_i}(\tilde{\theta}, \tilde{\phi}) = \delta_\theta(\tilde{\theta}) \delta_{\phi_i}(\tilde{\phi}_i), \ (\tilde{\theta}, \tilde{\phi}) \in \mathcal{F}. \tag{37}$$

Now, define $I : \mathcal{F} \longrightarrow \Pi_{i=1}^n \Re^{\mathcal{F}_i}$ by

$$I(\tilde{\theta}, \tilde{\phi}) = ((I_{\theta, \phi_i}(\tilde{\theta}, \tilde{\phi}))_{(\theta, \phi_i) \in \mathcal{F}_i})_{i=1}^n, \ \forall (\tilde{\theta}, \tilde{\phi}) \in \mathcal{F}. \tag{38}$$

Notice

$$F(\theta, \phi_i) = \sum_{(\tilde{\theta}, \tilde{\phi}): \tilde{\theta} = \theta, \tilde{\phi}_i = \phi_i} F(\tilde{\theta}, \tilde{\phi}) = \sum_{\tilde{\phi}: \tilde{\phi}_i = \phi_i} F(\theta, \tilde{\phi})$$

for all $(\theta, \phi_i) \in \mathcal{F}_i$. Thus, by (37) and (38),

$$((F(\theta, \phi_i))_{(\theta, \phi_i) \in \mathcal{F}_i})_{i=1}^n = \sum_{(\tilde{\theta}, \tilde{\phi})} F(\tilde{\theta}, \tilde{\phi}) I(\tilde{\theta}, \tilde{\phi}).$$

Since $\sum_{(\tilde{\theta},\tilde{\phi})} F(\tilde{\theta},\tilde{\phi}) = 1$, the preceding equation implies that $((F(\theta,\phi_i))_{(\theta,\phi_i)\in\mathcal{F}_i})_{i=1}^n$ is in the convex hull of the vectors $I(\tilde{\theta},\tilde{\phi})$, $(\tilde{\theta},\tilde{\phi}) \in \mathcal{F}$. Since $I(\tilde{\theta},\tilde{\phi}) \in \Pi_{i=1}^n \Re^{\mathcal{F}_i}$ for $(\tilde{\theta},\tilde{\phi}) \in \mathcal{F}$, it follows from the separating hyperplane theorem that for all i and for all $c_i \in \Re^{\mathcal{F}_i}$,

$$\max_{(\tilde{\theta},\tilde{\phi})\in\mathcal{F}} \sum_i \sum_{(\theta,\phi_i)\in\mathcal{F}_i} c_i(\theta,\phi_i) I_{\theta,\phi_i}(\tilde{\theta},\tilde{\phi}) \geq \sum_i \sum_{(\theta,\phi_i)\in\mathcal{F}_i} c_i(\theta,\phi_i) F(\theta,\phi_i).$$

Since $I_{\theta,\phi_i}(\tilde{\theta},\tilde{\phi}) \neq 0$ if and only if $\tilde{\theta} = \theta$ and $\tilde{\phi}_i = \phi_i$, the above inequality is equivalent to

$$\max_{(\tilde{\theta},\tilde{\phi})\in\mathcal{F}} \sum_i c_i(\tilde{\theta},\tilde{\phi}_i) \geq \sum_i \sum_{(\theta,\phi_i)\in\mathcal{F}_i} c_i(\theta,\phi_i) F(\theta,\phi_i).$$

By Definition 2 and the consistency of F, the necessity of conditions (a)-(c) is thus proved by letting $h_i : \Theta_i \times \Delta(\Theta_{-i})$ be defined by $h_i(\theta_i,\phi_i) = \text{marg}_{\Theta_i\times\Delta(\Theta_{-i})} F(\theta_i,\phi_i)$ for all $(\theta_i,\phi_i) \in \Theta_i \times \Delta(\Theta_{-i})$ and for all i.

To prove the sufficiency, let $\mathcal{F}$ be a subset of $\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i})$ and let h_i be a function from $\Theta_i \times \Delta(\Theta_{-i})$ to $\Re$ for all i such that conditions (a)-(c) are satisfied. Then, by condition (c), the vector $((\phi_i(\theta_{-i})h_i(\theta_i,\phi_i))_{(\theta,\phi_i)\in\mathcal{F}_i})_{i=1}^n$ is in the convex hull of the vectors $I(\tilde{\theta},\tilde{\phi})$, $(\tilde{\theta},\tilde{\phi}) \in \mathcal{F}$. Thus, there exists an element $F \in \Delta(\Theta \times \Pi_{i=1}^n \Delta(\Theta_{-i}))$ such that

$$((\phi_i(\theta_{-i})h_i(\theta_i,\phi_i))_{(\theta,\phi_i)\in\mathcal{F}_i})_{i=1}^n = \sum_{(\tilde{\theta},\tilde{\phi})\in\mathcal{F}} F(\tilde{\theta},\tilde{\phi}) I(\tilde{\theta},\tilde{\phi}).$$

By (37) and (38), the preceding equation equivalent to

$$\phi_i(\theta_{-i})h_i(\theta_i,\phi_i) = F(\theta,\phi_i),\ \forall(\theta,\phi_i) \in \mathcal{F}_i,\ \forall i. \tag{39}$$

Since $\phi \in \Delta(\Theta_{-i})$ for all i and for all $(\theta,\phi_i) \in \mathcal{F}_i$, it follows from (39) that

$$h(\theta_i,\phi_i) = F(\theta_i,\phi_i). \tag{40}$$

This together with condition (b) implies

$$F(\theta_i,\phi_i) > 0, \forall(\theta,\phi_i) \in \mathcal{F}_i,\ \forall i. \tag{41}$$

Putting (39)-(41) together, F is a consistent order-2 nature with supp$F = \mathcal{F}$. Hence, by condition (a), $\mathcal{P} = (P_1, \cdots, P_n)$ is embedded in F. This concludes the proof of the sufficiency. ■