



Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Beliefs, payoffs, information: On the robustness of the BDP property in models with endogenous beliefs[☆]

Alia Gizatulina^{*}, Martin Hellwig

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, D-53113 Bonn, Germany

ARTICLE INFO

Article history:

Received 25 February 2013

Received in revised form

23 August 2013

Accepted 27 October 2013

Available online xxxx

Keywords:

Abstract type spaces

Mechanism design

Correlated information

Belief functions

BDP property

ABSTRACT

Neeman (2004) and Heifetz and Neeman (2006) have shown that, in auctions with incomplete information about payoffs, full surplus extraction is only possible if agents' beliefs about other agents are fully informative about their own payoff parameters. They argue that the set of incomplete-information models with common priors that satisfy this so-called BDP property ("beliefs determine preferences") is negligible. In contrast, we show that, in models with finite-dimensional abstract type spaces, the set of belief functions with this property is topologically generic in the set of all belief functions. Our result implies genericity of (non-common or common) priors with the BDP property.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Neeman (2004) and Heifetz and Neeman (2006) have drawn attention to the so-called BDP property and the role this property plays in mechanism design models with correlated values. The label "BDP" – beliefs determine preferences – refers to the fact that, if this property holds and if one knows an agent's beliefs, then one also knows his preferences. They show that this property underlies the findings of Crémer and McLean (1988) or McAfee and Reny (1992) that correlations of agents types can be used to eliminate information rents. In Crémer and McLean (1988) or McAfee and Reny (1992), differences in an agent's beliefs about the other agents' characteristics induce differences in attitudes towards state-contingent payment schemes; such schemes are then used to screen agents in order to extract rents. Such rent extraction is necessarily incomplete if there are different states of the world in which an agent has different payoff parameters and the same beliefs.¹

Heifetz and Neeman (2006) suggest that the set of incomplete-information models having the BDP property is a negligible subset of the set of all incomplete-information models that are consistent with common priors.² Their suggestion is based on the view, that, in a model with private values, the preferences of an agent can be specified independently of his beliefs about other agents (Heifetz and Neeman, 2006, p. 215).

We want to take issue with this view. Beliefs are the result of agents' forming expectations on the basis of whatever information they have.³ This information includes their own preferences. The view that preferences and beliefs can be specified independently presumes that the information that an agent has about his preferences is not relevant for forming expectations about other agents' characteristics. This would be the case, for example, if the agent believed that his own payoff parameters and the other agents' characteristics were stochastically independent. If the agent believed that his own payoff parameters and the other agent's characteristics were correlated, his information about his

[☆] For helpful comments and discussions, we would like to thank an anonymous referee, as well as Beth Allen, Sophie Bade, Felix Bierbrauer, Yi-Chun Chen, Jacques Crémer, Christoph Engel, Amanda Friedenberg, Bruno Jullien, Aniol Llorente-Saguer, Benny Moldovanu, Stephen Morris, Zvika Neeman, Wojciech Olszewski, and Siyang Xiong.

^{*} Corresponding author. Tel.: +49 228 91416 820.

E-mail addresses: gizatulina@coll.mpg.de (A. Gizatulina), hellwig@coll.mpg.de (M. Hellwig).

¹ Neeman (2004) applies this reasoning to a public-good provision problem with participation constraints. He shows that feasible provision levels are close to zero

when there are many agents and certain violations of BDP are uniform over agents and states of the world, regardless of the number of participants. Gizatulina and Hellwig (2010) show that such uniform violations of BDP are incompatible with the notion that, in a large economy, individuals may be informationally small.

² Formally, they study the status of BDP models within a fixed family of incomplete-information models. Assuming that this family is what they call "closed under finite unions", they show that, if the family contains at least one BDP model, then within this family, failure of the BDP property is generic in a geometric and in a measure theoretic sense. We discuss their results in Section 6.

³ For a forceful statement of this view, see Aumann (1987, 1998).

own preferences would be relevant for his forming expectations about the other agents. This might be the case, for example, if the agent believes that other agents have observed an informative signal about his own payoff parameters, and he expects this signal to affect their beliefs. The value of the signal is part of the other agents' characteristics, and the agent's expectations about the other agents' signal observations will be based on his own observation of the payoff parameter itself.⁴

Once beliefs are treated as a result of agents' forming expectations on the basis of the information they have, the notion that the preferences of an agent can be specified independently of his beliefs about other agents is problematic. From this perspective, independence is the exception. Indeed, stochastic independence of the agent's payoff parameters with the other agents' characteristics would be incompatible with correlated values, the very specification that the literature on surplus extraction is concerned with.⁵

However, even when payoff parameters provide information about other agents' characteristics and agents condition their beliefs on their payoff parameters as well as any other information they may have, the BDP property, i.e., the ability to infer payoff parameters from beliefs, cannot be taken for granted. If an agent has many sources of information, it may be impossible to tell which combination of information variables has given rise to a given belief. Suppose, for example, that in an auction with two participants, an agent observes not only the value he assigns to the object at auction but also a signal of the value that other agent assigns to the object. With correlated private values, a given belief that the other agent assigns a high value to the object may then be due to the fact that the agent's own value is high or to the fact that the agent has received an optimistic signal about the other agent.

In summary, if beliefs are obtained by conditioning on information and if this information includes the agent's preferences, the question is whether the map from information to beliefs is invertible. If it is, then the agent's information, including his preferences, can be inferred from his beliefs. If the map from information to beliefs is not invertible, this inference is not possible for every specification of beliefs that might occur. The confounding of influences of different information variables on beliefs makes it impossible to recover the value of any information variable from the beliefs.

Even if the agent observes only his payoff characteristics, e.g. the value he assigns to the object that is auctioned off, correlations might be such that the map from information variables to beliefs is not monotonic. In this case, too, a given belief about the other agent's valuation might be generated by different realizations of the agent's own valuation.

In summary, the BDP property is a property of the mapping from information variables to beliefs. If this mapping is invertible, the BDP property holds. The question then is whether a lack of invertibility of the map from information variables to beliefs should be considered to be the rule or the exception.

Heifetz and Neeman (2006) do not actually study this question. They impose a common-prior assumption and discuss the genericity of the BDP property as a property of the common prior without investigating the implications of the requirement that beliefs must be given by conditional distributions.

This is where our paper steps in. We explicitly consider the mappings from information variables to beliefs, modelled as conditional distributions. Our main result shows that a confounding of influences of different information variables is

unlikely to occur if the set of objects about which the agent forms his beliefs is sufficiently rich. In this case, the BDP property is the rule, rather than the exception.

Throughout the paper, we consider an agent's "type" as reflecting not just his preference parameters but also additional information variables that he may observe. We do not limit ourselves to the "naive" type spaces considered by Crémer and McLean (1988), where agents' types are defined by their preference parameters only and differences in preference parameters are the only source of heterogeneity in beliefs. If type sets are finite, the logic of Crémer and McLean (1988) still implies that, if all agents' type sets have the same cardinality, the BDP property is generic.⁶

Our main results, however, concern models with a continuum of types. We treat the type t_i of agent i as a vector in \mathbb{R}^{n_i} for some natural number n_i . The vector of all agents' types is a vector in \mathbb{R}^N where $N = \sum_i n_i$. Priors are probability distributions on \mathbb{R}^N . Belief functions (regular conditional distributions) are mappings from agents' own type spaces into the spaces of probability distributions over the other agents' type spaces. We restrict the analysis to belief functions that are continuous and to priors that admit continuous belief functions as regular conditional distributions and that, moreover, have marginal distributions of agents' types which have full supports.

For any agent i , a belief function is given by a continuous function from the space \mathbb{R}^{n_i} of agent i 's types into the space $M(\mathbb{R}^{N-n_i})$ of probability distributions over the space of the other agents' types. We endow the space of such functions with the topology of uniform convergence. We also endow the space of priors with the coarsest topology under which the maps from priors to marginal distributions of agents' types and to continuous regular conditional probability distributions over the other agents' types are continuous.

With this specification of topologies, we show that, for any agent i , the BDP property holds for a residual set of belief functions, i.e. for a countable intersection of open and dense subsets of the space of continuous functions from \mathbb{R}^{n_i} to $M(\mathbb{R}^{N-n_i})$. We also show that, for any agent i , the BDP property holds for a residual set of priors, i.e., on a countable intersection of open and dense sets of probability distributions on \mathbb{R}^N .

These results are based on an extension of the classical Embedding Theorem for continuous functions.⁷ An embedding of a metric space X in a metric space Y is a one-to-one bicontinuous function from X into (a subset of) Y . The standard version of the Embedding Theorem asserts that, for any natural numbers n and m , if $m \geq 2n + 1$, then the set of embeddings of a set $X \subset \mathbb{R}^n$ in $[0, 1]^m$ contains a residual subset of the set of continuous functions from X into $[0, 1]^m$. We use this theorem to show that, for any natural number n , the set of embeddings of a compact set $X \subset \mathbb{R}^n$ into the space $Y = \mathcal{M}(Z)$ of probability measures on a compact subset $Z \subset \mathbb{R}^{N-n_i}$ that has infinitely many elements contains a residual subset of the set of continuous functions from X into $\mathcal{M}(Z)$. The proof of this result is based on the observation that, no matter what n_i and N may be, the dimension of $\mathcal{M}(Z)$ exceeds $2n_i + 1$.

Because embeddings are injective, the Embedding Theorem immediately implies that the BDP property holds for a residual set of belief functions, i.e. continuous regular conditional distributions for any agent i . Under the given topology on the set of priors, it

⁶ This confirms a conjecture in Compte and Jehiel (2009, p. 188).

⁷ See Chapter V in Hurewicz and Wallman (1941). In economics, the literature on generic existence of completely revealing rational expectations equilibria has made extensive use of Embedding Theorems; see, in particular, Allen (1981). That literature, however, relied on Whitney's Embedding Theorem for C^1 functions (Hirsch, 1994, p. 35). Lacking the requisite differentiability properties, we use the Embedding Theorem for C^0 functions.

⁴ This effect is central to the analysis of Gizatulina and Hellwig (2010).

⁵ In a common-prior framework, Börgers and Oh (2012) show that stochastic independence of payoffs and beliefs actually implies that the payoff characteristics of any one agent are stochastically independent of all other agents' payoff characteristics and beliefs.

follows that the set of priors giving rise to belief functions with the BDP property must contain a residual subset of the set of all priors that give rise to continuous regular conditional distributions. For any agent i , the BDP property thus holds on a residual set of priors.

Because the set of priors that satisfy the BDP property for all agents is given by the (finite) intersection of the sets of priors that satisfy the BDP property for agent i , $i = 1, \dots, I$, the BDP property also holds for all agents on a residual set of common priors. On the set of common priors that have continuous density functions, the topology that yields residualness of the BDP property turns out to be the uniform topology for density functions.

The relation of our analysis to Heifetz and Neeman (2006) and the subsequent literature is discussed in detail in Section 6 below. There we show that the difference between our analysis and the analysis of Heifetz and Neeman (2006) is due to their considering genericity of the BDP property *within* a fixed family of models. If the family in question is the family of all common-prior models in the universal type space, this approach involves no loss of generality. In finite-dimensional abstract type spaces, however, the analysis of Heifetz and Neeman begs the question whether the families of models within which their genericity results hold are themselves robust. In a companion paper, Gizatulina and Hellwig (2012), we show that this is not the case. Using the results of the present paper, we show there that the set of families within which the BDP property is robust is itself a residual set in the set of all families of incomplete-information models.

In contrast to this paper, Barelli (2009) and Chen and Xiong (2011) work with the universal type space. Barelli (2009) asserts that non-BDP models are topologically generic. Chen and Xiong (2011) point to a flaw in his analysis and show that the BDP property holds on a residual set of models in the universal type space if this space is endowed with the weak* topology. They use the fact that, in the weak* topology on the priors of the universal type space, priors over finite type sets are dense in the set of all models and, by the logic of Crémer and McLean (1988), BDP priors are dense in the set of finite priors. In subsequent work, Chen and Xiong (2013) use the same logic to establish a genericity result for full surplus extraction.

The universal type space approach diverts attention away from the fact that the BDP property is a property of belief functions. In a universal type space approach, belief functions are fixed and given by the formalism. Each “type” is defined by a vector consisting of the agent’s payoff parameters and the agent’s beliefs about the other agents’ payoff parameters, the agent’s beliefs about the other agents’ beliefs about other agents’ payoff parameters, etc. In this setting, that function that assigns beliefs to “types” is simply the projection from the type space to the space of belief hierarchies or, if we think of the universal type space itself as an example of an abstract type space, the function that assigns to each “universal type” of an agent the unique measure on the space of the other agents’ universal types that is compatible with the belief hierarchy encoded in the universal type. In this setting, the notion that an agent’s beliefs reflect the information contained in his “type” seems trivial because the beliefs themselves are a part of what makes up his type. There is however a question as to whether the information contained in the other components of an agent’s type, in particular his payoff parameters, is properly reflected in his beliefs.

Because belief functions are fixed and given by the formalism, the universal type space approach is ill-suited to studying the BDP property as a property of belief functions. It can however be introduced as a property of a subset of the universal type space. A subset of the universal type space has the BDP property for a given agent if any two distinct types of the agent that belong to the set involve different belief hierarchies; the set must not contain be two types with the same belief hierarchies but different payoff parameters.

For arbitrarily chosen subsets of the universal type space, there is no reason why such a property should hold. For belief-closed subsets, however, such a property would be implied by suitable restrictions on the agents’ interactive belief systems. For any belief-closed subset of the universal type space, the BDP property holds for agent i if this agent believes that the beliefs of all agents $j \neq i$ about agent i ’s type are concentrated on a set for which the BDP property holds. Genericity of the BDP property can then be studied in terms of the genericity of belief-closed subsets with this property in the space of all belief-closed subsets of the universal type space. Barelli (2009) and Chen and Xiong (2011) actually impose an even stronger consistency condition on interactive belief systems and treat the BDP property as a property of common priors.

As a property of an agent’s belief function, however, the BDP property for can be studied without any additional presumptions about the interactive belief systems, i.e. without any reference to the beliefs of other agents, let alone the consistency conditions inherent in working with a belief-closed set or with a common prior. When different agent’s beliefs are derived from different priors, it still makes sense to ask whether an agent’s belief function exhibits the BDP property. For this question, however, one needs a framework in which the function mapping an agent’s information into his beliefs is not trivially given by having the beliefs themselves encoded as part of his information. An abstract type space approach offers such a framework.

In the following, Section 2 introduces the basic framework of our analysis, including a definition of the BDP property for belief functions. Section 3 discusses the BDP property for priors and relates our approach to that of Heifetz and Neeman (2006). Section 4 illustrates the role of information for belief formation by means of several examples. Section 5 presents our genericity results. In particular, Section 5.1 gives the result for finite type spaces. Sections 5.2–5.4 give the results for continuous type spaces: Section 5.2 deals with the BDP property for a single agent, Section 5.3 with the BDP property for all agents in models with common priors. Section 5.4 deals with the BDP property in models with common priors that have continuous density functions. Section 6 discusses the relation of our analysis to the literature.

2. The basic framework

An abstract (Harsanyi) type space formulation of an incomplete-information model with $I \geq 2$ agents involves a collection

$$\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I \quad (1)$$

such that, for any i , T_i is a set of abstract “types” for agent i , $\hat{\theta}_i$ is a mapping from T_i into a set Θ_i of payoff parameter vectors for agent i , and $\hat{\pi}_i$ is a mapping from T_i into the set $\mathcal{M}(T_{-i})$ of probability distributions on the space $T_{-i} := \prod_{j \neq i} T_j$ of the other agents’ abstract types.⁸ For any $t_i \in T_i$, $\hat{\theta}_i(t_i)$, the payoff type of agent i , indicates the agent’s payoff parameters when his abstract type is t_i ; $\hat{\pi}_i(t_i)$, the belief type of agent i , represents the agent’s probabilistic beliefs about the other agents’ types. Following Heifetz and Neeman (2006), we assume that the spaces T_i and Θ_i are complete separable metric spaces, and that the functions $\hat{\theta}_i : T_i \rightarrow \Theta_i$ and $\hat{\pi}_i : T_i \rightarrow \mathcal{M}(T_{-i})$ are continuous, where $\mathcal{M}(T_{-i})$ has the topology of weak convergence of probability measures, i.e. the weak* topology.⁹

⁸ See, e.g., Bergemann and Morris (2005).

⁹ We use the word “model” in a wide sense, like Bergemann and Morris (2005) or Heifetz and Neeman (2006). In contrast, Barelli (2009), as well as Chen and Xiong (2011), Chen and Xiong (2013) use the word more specifically for a common prior model with a minimal consistent belief subspace of the universal type space.

Definition 2.1. An incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ exhibits the BDP property for agent i if for any t_i and t'_i in T_i ,

$$\hat{\pi}_i(t_i) = \hat{\pi}_i(t'_i) \text{ implies } \hat{\theta}_i(t_i) = \hat{\theta}_i(t'_i). \tag{2}$$

If there are no restrictions on the mappings $\hat{\theta}_i$ and $\hat{\pi}_i$, there is no reason why an incomplete-information model should exhibit this property. One expects the set of models exhibiting this property to be in some sense negligible in the set of all incomplete-information models.

This conclusion cannot be taken for granted, however, if the belief mapping $\hat{\pi}_i$ must satisfy

$$\hat{\pi}_i(t_i) = b_i(t_i, v_i), \tag{3}$$

for all $t_i \in T_i$, where $v_i \in \mathcal{M}(T_i \times T_{-i})$ is a prior for agent i and $b_i(t, v_i)$ is the value of a regular conditional distribution for t_{-i} given t_i that is induced by v_i , i.e. a function from T_i to $\mathcal{M}(T_{-i})$ such that for any bounded continuous function $f : T_{-i} \rightarrow \mathbb{R}$, $\int_{T_{-i}} f(t_{-i}) b_i(dt_{-i}|t_i, v_i)$ is the conditional expectation of $f(t_{-i})$ given t_i . Heifetz and Neeman impose this restriction with the additional requirement that v_i be common to all agents.¹⁰ In this paper, we shall consider both, the case of agent-specific priors v_i , $i = 1, \dots, I$, and the case of a common prior ν such that $v_i = \nu$ for all i .

The information on which agent i conditions his beliefs includes his payoff type $\hat{\theta}_i(t_i)$. To make this dependence explicit, we find it convenient to write (3) in the form

$$\hat{\pi}_i(\cdot) = \hat{b}_i(\hat{\theta}_i(\cdot), \hat{s}_i(\cdot), v_i), \tag{4}$$

where, for any $t_i \in T_i$, $\hat{s}_i(t_i)$ is a vector of payoff-irrelevant information variables that agent i observes in addition to his payoff type $\hat{\theta}_i(t_i)$, and the mapping $\hat{s}_i(\cdot)$ takes values in some set S_i .¹¹ In this formulation, the underlying type space T_i matters only to the extent that t_i affects the payoff type $\hat{\theta}_i(t_i)$ and the information vector $\hat{s}_i(t_i)$. Therefore there is no loss of generality in identifying abstract types with pairs of payoff and signal vectors and writing

$$T_i = \Theta_i \times S_i, \tag{5}$$

with the understanding that, for any $t_i \in T_i$,

$$t_i = (\theta_i, s_i) \text{ implies } \hat{\theta}_i(t_i) = \theta_i \text{ and } \hat{s}_i(t_i) = s_i, \tag{6}$$

i.e., the maps $\hat{\theta}_i(\cdot)$ and $\hat{s}_i(\cdot)$ are the projections from T_i to Θ_i and S_i . The representation (1) of the incomplete-information model then takes the form

$$\hat{\mathcal{T}} = \{T_i, \hat{\pi}_i\}_{i=1}^I = \{\Theta_i \times S_i, \hat{\pi}_i\}_{i=1}^I, \tag{7}$$

with the understanding that payoff and information mappings are the projections and that belief mappings satisfy (4) for some prior v_i . In this formulation, the beliefs

$$\hat{\pi}_i(t_i) = \hat{b}_i(\hat{\theta}_i(t_i), \hat{s}_i(t_i), v_i) = b_i(t_i, v_i) \tag{8}$$

of agent i concern the pairs $t_j = (\hat{\theta}_j(t_j), \hat{s}_j(t_j))$, $j \neq i$, of the other agents' payoff and information vectors.¹²

As is standard for abstract type space formulations and is proved in the Appendix to Heifetz and Neeman (2006), the mapping $\hat{\pi}_i(\cdot) = b_i(\cdot, v_i)$ can be used to build an infinite hierarchy of beliefs of i about the distribution of θ_{-i} , the joint distribution of θ_{-i} and the other agents' beliefs about preferences, etc. The abstract type space can thus be mapped into the Θ -based universal type space of Mertens and Zamir (1985). If there is a common prior, i.e., if the priors v_i are all the same, the universal type space image of the abstract type space model must have the prior assign all probability mass to a belief-closed subset of the universal type space with the additional property that the belief of any type must correspond to the value of a regular conditional distribution given the information that this type has.

3. The BDP property

In the preceding section, we have introduced the BDP property as a property of the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$, without any reference to the priors that give rise to the belief functions $\hat{\pi}_i$, $i = 1, \dots, I$. In contrast, Heifetz and Neeman (2006) study the BDP property as a property of agents' priors. We shall also consider the BDP property as a property of priors but base the definition of this property on Definition 2.1.

In the following, we refer to a measure v_i on the set $T := \prod_{j=1}^I T_j$ as a prior for agent i in the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ if, under the measure v_i , $\hat{\pi}_i(\cdot)$ is a regular conditional distribution for t_{-i} given t_i , i.e., if $\hat{\pi}_i(t_i) = b_i(t_i, v_i)$ for all $t_i \in T_i$. We refer to a measure ν as a common prior for the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ if it is a prior for every agent in the incomplete-information model $T = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$.

Definition 3.1. For any j , let v_j be a measure on T that is a prior for agent j in the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$. The prior v_j has the BDP property for agent j if the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ has the BDP property for agent j , i.e., if, for any $t_j = (\theta_j, s_j)$ and $t'_j = (\theta'_j, s'_j)$ in $T_j = \Theta_j \times S_j$, $\hat{\pi}_j(t_j) = \hat{\pi}_j(t'_j)$ implies $\hat{\theta}_j(t_j) = \hat{\theta}_j(t'_j)$.

Definition 3.2. A common prior ν for the incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ has the BDP property if it has the BDP property for each agent $k = 1, \dots, n$.

For Definitions 3.1 and 3.2 to be unambiguous, we need to restrict the space of priors under consideration. The reason is that any given prior can be compatible with multiple belief functions. If $\hat{\pi}_j(\cdot) = b_j(\cdot|v_j)$ is a regular conditional distribution for t_{-j} given t_j under the prior v_j , then any function $\bar{\pi}_j(\cdot)$ such that $\bar{\pi}_j(t_j) = \hat{\pi}_j(t_j)$ for v_j -almost all $t_j \in T_j$ is also a regular conditional distribution for t_{-j} given t_j under the prior v_j . If the incomplete-information model $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ has the BDP property for agent j and the incomplete-information model $\{T_i, \hat{\theta}_i, \bar{\pi}_i\}_{i=1}^I$ that is obtained from $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ by replacing $\hat{\pi}_j(\cdot)$ with $\bar{\pi}_j(\cdot)$ does not have the BDP property, Definition 3.1 would suggest that v_j both has and fails to have the BDP property for agent j .

To see this point, consider the following example: Let $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^2$ be such that, for $i = 1, 2$, $T_i = \Theta_i \times S_i$, where $\Theta_i = S_i = [0, 1]$, and $\hat{\pi}_i : [0, 1]^2 \rightarrow \mathcal{M}([0, 1]^2)$ is a continuous injection. Let $\bar{v}_1 \in \mathcal{M}([0, 1]^2)$ be such that all probability mass is assigned to the set $[0, 1] \times \{1\}$, i.e. $v_1([0, 1] \times \{1\}) = 1$ and $v_1([0, 1] \times \{0\}) = 0$, and let $v_1 \in \mathcal{M}([0, 1]^4)$ be the distribution of the vector $(\theta_1, s_1, \theta_2, s_2)$ that is induced by $\bar{v}_1(\cdot)$ and the regular conditional distribution $\hat{\pi}_1(\cdot|\cdot)$ for (θ_2, s_2) given (θ_1, s_1) . Because $\hat{\pi}_1 : [0, 1]^2 \rightarrow \mathcal{M}([0, 1]^2)$ is a continuous injection, the model $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^2$ has the BDP property for agent 1, so, Definition 3.1 implies that v_1 has the BDP property for agent 1.

¹⁰ Dekel et al. (2006) suggest that, in a non-common-prior environment, the approach of Heifetz and Neeman would not even be well defined.

¹¹ This is without loss of generality. Any model $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ can trivially be rewritten in the form $\{\hat{T}_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ where $\hat{T}_i = \Theta_i \times T_i$, $\hat{\theta}_i$ is the projection from $\Theta_i \times T_i$ to Θ_i , and, with \hat{s}_i given as the projection from $\Theta_i \times T_i$ to T_i , $\hat{\pi}_i$ satisfies (4). In this formulation, the original type space T_i itself is interpreted as a space of signals on which expectations are conditioned. For a discussion in the context of the universal type space, see Section 6 below.

¹² A formulation with conditioning on other information variables, in addition to payoff parameters, was previously proposed by Compte and Jehiel (2009, p. 188).

Next, consider the function $\bar{\pi}_1 : [0, 1]^2 \rightarrow \mathcal{M}([0, 1]^2)$ such that

$$\bar{\pi}_1(\theta_1, s_1) \equiv \hat{\pi}_1 \left(\frac{\theta_1}{2 - s_1}, 1 \right) \tag{9}$$

for all $(\theta_1, s_1) \in [0, 1]^2$. Because $\hat{\pi}_1$ is continuous, the function $\bar{\pi}_1$ is continuous as well, so $\{T_i, \hat{\theta}_i, \bar{\pi}_i\}_{i=1}^I$ is also an incomplete-information model. Moreover, (9) implies $\bar{\pi}_1(\theta_1, s_1) \equiv \hat{\pi}_1(\theta_1, s_1)$ whenever $s_1 = 1$. Since $\nu_1([0, 1] \times \{1\}) = 1$, it follows that, like $\hat{\pi}_1$, under the prior ν_1 , $\bar{\pi}_1$ is a regular conditional distribution for (θ_2, s_2) given (θ_1, s_1) , i.e., ν_1 is a prior for agent 1 in the incomplete-information model $\{T_i, \hat{\theta}_i, \bar{\pi}_i\}_{i=1}^I$, as well as $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$. However, the incomplete-information model $\{T_i, \hat{\theta}_i, \bar{\pi}_i\}_{i=1}^I$ does not have the BDP property for agent 1 because, under (9), we have $\bar{\pi}_1(\theta_1, s_1) = \bar{\pi}_1(\theta'_1, s'_1)$ whenever $\frac{\theta'_1}{2-s'_1} = \frac{\theta_1}{2-s_1}$.

To avoid this kind of paradox, we restrict our analysis to priors on the set $T := \prod_{i=1}^I T_i$ for which the marginal distributions on the individual type spaces have full support, i.e., for $i = 1, \dots, I$, the support of the distribution of the type t_i of agent i is T_i . The following lemma shows that, for such priors, the requirement that belief functions must be continuous eliminates the multiplicity of belief functions that are compatible with a given prior.¹³ For such priors, there is no ambiguity about the meaning of Definitions 3.1 and 3.2.

Lemma 3.3. For $j = 1, \dots, I$, let T_j be a complete separable metric space, and let $T := \prod_{i=1}^I T_i$. For any i , let $\hat{\pi}_i : T_i \rightarrow \mathcal{M}(T_{-i})$ be a continuous function and let $\nu_i \in \mathcal{M}(T)$ be such that (3) holds, i.e., ν_i is a prior for $\hat{\pi}_i$. If the support of the marginal distribution on agent i 's types that is induced by ν_i is equal to T_i , then $\hat{\pi}_i$ is the unique continuous regular conditional distribution for t_{-i} given t_i that is induced by ν_i .

Proof. If the lemma is false, there exists a prior $\nu_i \in \mathcal{M}(T)$ such that the support of the marginal distribution on agent i 's types is equal to T_i and there exist two belief functions $b_i^1(\cdot, \nu_i)$, $b_i^2(\cdot, \nu_i)$ be such that $b_i^1(\cdot, \nu_i)$ and $b_i^2(\cdot, \nu_i)$ are both continuous and both regular conditional distributions for t_{-i} given t_i under ν_i . If $b_i^1(\cdot, \nu_i) \neq b_i^2(\cdot, \nu_i)$, there exists $t_i \in T_i$ such that $b_i^1(t_i, \nu_i) \neq b_i^2(t_i, \nu_i)$. Because $b_i^1(\cdot, \nu_i)$ and $b_i^2(\cdot, \nu_i)$ are both regular conditional distributions for t_{-i} given t_i under ν_i , it must be the case that $\nu_i(\{t_i\} \times T_{-i}) = 0$. Because the support of the marginal distribution on agent i 's types is equal to T_i , it must be the case that, for any $\varepsilon > 0$, $\nu_i(B_\varepsilon(t_i) \times T_{-i}) > 0$, where $B_\varepsilon(t_i)$ is an ε -neighbourhood of t_i . However, because $b_i^1(\cdot, \nu_i)$ and $b_i^2(\cdot, \nu_i)$ are both continuous, $b_i^1(t_i, \nu_i) \neq b_i^2(t_i, \nu_i)$ implies $b_i^1(t'_i, \nu_i) \neq b_i^2(t'_i, \nu_i)$ for all $t'_i \in B_\varepsilon(t_i)$ if $\varepsilon > 0$ is sufficiently small. But then, $b_i^1(t'_i, \nu_i) \neq b_i^2(t'_i, \nu_i)$ for

¹³ Without continuity, the multiplicity of regular conditional distributions for a given prior is even more vexing. For an example, suppose again that $\theta_i = [0, 1]$, $S_i = [0, 1]$, hence $T_i = [0, 1]^2$. For a given prior ν_i , let $b_i(\cdot|\cdot, \nu_i)$ be a regular conditional distribution for t_{-i} given $t_i = (\theta_i, s_i)$ such that the BDP property in the sense of Definition 3.1 is satisfied. Construct another regular conditional distribution for t_{-i} given t_i by setting

$$\hat{b}_i(\cdot|\theta_i, s_i, \nu_i) = b_i(\cdot|\theta_i, s_i, \nu_i) \quad \text{if } \theta_i \in (0, 1], s_i \in [0, 1]$$

and

$$\hat{b}_i(\cdot|\theta_i, s_i, \nu_i) = b_i(\cdot|\psi(s_i), \nu_i) \quad \text{if } \theta_i = 0, s_i \in [0, 1],$$

where $\psi : [0, 1] \rightarrow [0, 1]^2$ is Peano's space-filling function. If ν_i assigns measure zero to the event $\theta_i = 0$, we have $\hat{b}_i(t_i, \nu) = b_i(t_i, \nu)$ for ν_i -almost all (θ_i, s_i) , implying that $\hat{b}_i(\cdot, \nu_i)$ is indeed another regular conditional probability distribution for t_{-i} given t_i under ν_i . However, for every type vector $(\theta_i, s_i) \in (0, 1] \times [0, 1]$ for agent i , there is $s'_i \in [0, 1]$ such that $\hat{b}_i(\theta_i, s_i, \nu_i) = \hat{b}_i(0, s'_i, \nu_i)$.

all t'_i in a set that has positive measure under $\nu_i(\cdot \times T_{-i})$. This is incompatible with the assumption that $b_i^1(\cdot, \nu_i)$ and $b_i^2(\cdot, \nu_i)$ are both regular conditional distributions for t_{-i} given t_i under ν_i . ■

In contrast to our approach, Heifetz and Neeman (2006) do not require marginal type distributions to have full support. Instead they use a definition of the BDP property which neglects null sets. In their analysis, a prior on the underlying type space satisfies the BDP property for agent i if there exists a subset \hat{T}_i of the type space for agent i such that the prior assigns probability one to \hat{T}_i and, moreover, the BDP condition (2) holds for all t_i and t'_i in \hat{T}_i .¹⁴ This definition of the BDP property for priors implies that, if the BDP property in their definition holds for one regular conditional distribution that is induced by a given prior, then it must hold for all regular conditional distributions that are induced by this prior. The ambiguity in the definition of the BDP property that is created by the multiplicity of regular conditional distributions is eliminated.

Because null sets are defined in terms of priors, the BDP property here is also defined as a property of priors. This contrasts with the specification of the incomplete-information model $T = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$, which does not refer to any prior but takes the belief functions $\hat{\pi}_i$ as primitives. Taking belief functions as primitives, we find it is more natural to start with a definition of the BDP property for belief functions.

The treatment of the BDP property in Heifetz and Neeman (2006) stands in the context of the discussion about auction design and surplus extraction with correlated values. Proposition 2, p. 218, in their paper asserts that any prior under which an auction mechanism can be designed so that all surplus is extracted must be a BDP prior.¹⁵ In the following, we prove an analogous result, showing that, if the incomplete-information model $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ admits full surplus extraction, then it must exhibit the BDP property in the sense of Definition 2.1.

Let $\{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$ be given and suppose that payoff parameters are scalars. Let A be a set of allocation choices and let $v_i(a, \theta_i)$ be the payoff that agent i associates with the choice $a \in A$ when the value of his payoff parameter is θ_i . An *direct mechanism* is a vector $(\hat{a}(\cdot), m_1(\cdot), \dots, m_I(\cdot))$ of functions on T such that, for any vector $(t_1, \dots, t_I) \in T$, $a(t_1, \dots, t_I)$ is the allocation choice and $m_i(t_1, \dots, t_I), \dots, m_I(t_1, \dots, t_I)$ are the different agents' payments when agents' types are t_1, \dots, t_I . The mechanism is *incentive-compatible* if

$$\int_{T_{-i}} [v_i(\hat{a}(t_i, t_{-i}), \hat{\theta}_i(t_i)) - m_i(t_i, t_{-i})] d\hat{\pi}_i(t_{-i}|t_i) \geq \int_{T_{-i}} [v_i(\hat{a}(t'_i, t_{-i}), \hat{\theta}_i(t_i)) - m_i(t'_i, t_{-i})] d\hat{\pi}_i(t_{-i}|t_i) \tag{10}$$

for all i and all t_i and t'_i in T_i . The mechanism is *individually rational* if

$$\int_{T_{-i}} [v_i(\hat{a}(t_i, t_{-i}), \hat{\theta}_i(t_i)) - m_i(t_i, t_{-i})] d\hat{\pi}_i(t_{-i}|t_i) \geq 0 \tag{11}$$

for all i and all t_i in T_i . The mechanism induces *up-to- ε surplus extraction from agent i* if

$$\int_{T_{-i}} [v_i(\hat{a}(t_i, t_{-i}), \hat{\theta}_i(t_i)) - m_i(t_i, t_{-i})] d\hat{\pi}_i(t_{-i}|t_i) \leq \varepsilon \tag{12}$$

for all t_i in T_i .

¹⁴ The concept of a BDP* model in Chen and Xiong (2011) also requires the BDP property to hold only with probability one.

¹⁵ As stated, the proposition is not actually correct. The strict inequality in the proof of the proposition on p. 219 presumes that the type θ'_i with a lower willingness to pay for the object actually has a positive probability of getting the object. If this type has zero probability of getting the object, the strict inequality is incorrect. A correct version of the proposition would have to state the additional assumption.

where

$$(\alpha_\theta \quad \alpha_s) = (\text{cov}(s_2, \theta_1) \quad \text{cov}(s_2, s_1)) \times \begin{pmatrix} \text{var}\theta_1 & \text{cov}(\theta_1, s_1) \\ \text{cov}(\theta_1, s_1) & \text{vars}_1 \end{pmatrix}^{-1}. \tag{19}$$

In this case, as in Example 4.1, agent 1’s belief about agent 2’s signal is affected by agent 1’s payoff type unless $\text{cov}(s_2, \theta_1) = 0$. However, if agent 1’s own signal is also correlated with s_2 , it is not possible to infer θ_1 from the belief variable $E_1[s_2|\theta_1, s_1]$ and the parameters of the prior ν_1 . For such an inference, one would also have to know the realization of agent 1’s own signal. In this setting, for any prior on T that is multivariate normal, the BDP property fails to hold for agent 1 except in the negligible case where $\text{cov}(s_2, s_1) = 0$.

Example 4.3. Let $I = 2$, $\Theta_1 = S_1 = \mathbb{R}$, $\Theta_2 = \{0\}$, $S_2 = \mathbb{R}^2$, and suppose that ν_1 is a multivariate normal distribution. Then agent 1’s conditional expectations satisfy

$$\begin{pmatrix} E_1[s_2^1|\theta_1, s_1] \\ E_1[s_2^2|\theta_1, s_1] \end{pmatrix} = A \begin{pmatrix} \theta_1 - E_1\theta_1 \\ s_1 - E_1s_1 \end{pmatrix} + \begin{pmatrix} E_1s_2^1 \\ E_1s_2^2 \end{pmatrix}, \tag{20}$$

where $A = \Sigma_{21}\Sigma_{11}^{-1}$, with Σ_{21} , Σ_{11} as submatrices of the variance–covariance matrix Σ of $s_2^1, s_2^2, \theta_1, s_1$, partitioned so as to reflect the distinction between the variables that agent 1 observes and the variables that he does not observe. In this specification, both θ_1 and s_1 can be inferred from the conditional expectations $E_1[s_2^1|\theta_1, s_1]$ and $E_1[s_2^2|\theta_1, s_1]$ whenever the matrix A is invertible, i.e., whenever the matrix Σ_{21} is nonsingular. Because the set of nonsingular two-by-two matrices is open and dense in the set of all two-by-two matrices, the BDP property is generic for agent 1 within the set of models covered by this example.

Whereas Example 4.2 involves a failure of BDP due to a confounding of influences of different information variables, the following Example 4.4 shows that BDP will also fail if, for given parameters of the prior ν_1 , the map from payoff types to conditional expectations is not monotonic (not one-to-one). Subsequently, Example 4.5 will show that this problem is likely to disappear if there are more variables about which to form expectations so that the vector of conditional expectations has a sufficiently high dimension.

Example 4.4. Let $I = 2$, $\Theta_1 = S_2 = \mathbb{R}$, $\Theta_2 = S_1 = \{0\}$, and suppose that ν_1 is the distribution that is generated when

$$s_2 = \theta_1^2 + \varepsilon, \tag{21}$$

where θ_1 and ε are independent normal random variables. In this case, agent 1’s conditional distribution for s_2 given θ_1 is normal with mean

$$E_1[s_2|\theta_1, s_1] = (\theta_1^2 - E_1\theta_1^2) + E_1s_2 \tag{22}$$

and variance $\text{Var}\varepsilon$. From the belief $E[s_2|\theta_1, s_1]$, one can infer θ_1^2 , but one cannot tell whether it is the positive or the negative solution of the equation

$$\theta_1^2 = E_1[s_2|\theta_1, s_1] - E_1s_2 + E_1\theta_1^2. \tag{23}$$

The BDP property fails to hold.

Example 4.5. Let $I = 2$, $\Theta_1 = \mathbb{R}$, $\Theta_2 = S_1 = \{0\}$, $S_2 = \mathbb{R}^2$, and suppose that ν_1 is the distribution that is generated when

$$s_2^1 = \theta_1^2 + \varepsilon, \tag{24}$$

$$s_2^2 = A\theta_1 + B\theta_1^2 + \eta \tag{25}$$

where θ_1, ε , and η are independent normal random variables and A and B are constants. In this case, agent 1’s conditional distribution for s_2^1 and s_2^2 given θ_1 is normal with means

$$E_1[s_2^1|\theta_1, s_1] = (\theta_1^2 - E_1\theta_1^2) + E_1s_2^1, \tag{26}$$

$$E_1[s_2^2|\theta_1, s_1] = A(\theta_1 - E_1\theta_1) + B(\theta_1^2 - E_1\theta_1^2) + E_1s_2^2 \tag{27}$$

and variance–covariance matrix $\begin{pmatrix} \text{Var}\varepsilon & 0 \\ 0 & \text{Var}\eta \end{pmatrix}$. If $A \neq 0$, then from (26) and (27), one obtains

$$\theta_1 = E_1\theta_1 + \frac{1}{A} [E_1[s_2^2|\theta_1, s_1] - E_1s_2^2 - BE_1[s_2^1|\theta_1, s_1] - E_1s_2^1], \tag{28}$$

which shows that θ_1 can be inferred by looking at $E_1[s_2^1|\theta_1, s_1]$ and $E_1[s_2^2|\theta_1, s_1]$ jointly. By looking at the two belief variables together, one overcomes the difficulty that neither belief variable alone is injective in θ_1 . The BDP property holds for agent 1 except in the negligible case $A = 0$.

5. Genericity results

5.1. The BDP property with finite type sets

Turning from these examples to a more general analysis, we first consider the case where the type sets T_i are finite. If each type set T_i is finite, with n_i elements, the space $T := \prod_{i=1}^I T_i$ is also finite, with $N = \prod_{i=1}^I n_i$ elements, and the prior F is represented by a vector $\Pi \in \mathbb{R}^N$, such that $\sum_{k=1}^N \Pi_k = 1$. The set of such vectors is endowed with the usual (Euclidean) topology.

Proposition 5.1. Assume that, for each i , T_i is a finite set with n_i distinct elements. For any i , let

$$N_{-i} := \prod_{\substack{j=1 \\ j \neq i}}^I n_j \tag{29}$$

be the cardinality of the set T_{-i} . If $n_i \leq N_{-i}$ for all i , then, for each i , the set P_i of priors that exhibit the BDP property for agent i is open and dense in the set of all priors on T .

Proof. Fixing i , we note that any N vector Π^N of probabilities on T can be written in matrix form as $\Pi(i) = (\pi_{t_i t_{-i}})$ where the different rows refer to different types t_i of agent i and different columns refer to the different elements t_{-i} of T_{-i} . The belief $b_i(\cdot|t_i)$ of agent i , conditional on t_i , about the other agents’ types is represented by a vector of conditional probabilities on T_{-i} . Bayes’ Law implies that, under the prior Π , this vector is proportional to the vector $(\pi_{t_i t_{-i}})$; one can write:

$$b_i(t_{-i}|t_i) = \lambda(t_i) \times \pi_{t_i t_{-i}} \tag{30}$$

for all $t_{-i} \in T_{-i}$, where

$$\lambda(t_i) = \frac{1}{\sum_{t_{-i} \in T_{-i}} \pi_{t_i t_{-i}}} \tag{31}$$

is chosen to ensure that the entries in (30) sum to one. Eqs. (30) and (31) imply that, if the rows of the matrix $\Pi(i)$ are linearly independent, as well as strictly positive, then so are the belief vectors $b_i(\cdot|t_i)$, $t_i \in T_i$. This implies, in particular, that the belief vectors $b_i(\cdot|t_i)$, $t_i \in T_i$, are all distinct and the function $t_i \rightarrow b_i(\cdot|t_i)$ is invertible, i.e. one can infer the type t_i of agent i from his belief vector. Given that $t_i = (\theta_i, s_i)$, this means, in particular, that one can infer θ_i from $b_i(\cdot|t_i)$. The proposition follows because, by standard arguments, for $n_i \leq N_{-i}$, the set of $n_i \times N_{-i}$ matrices with linearly independent rows is an open and dense subset P_i of $\mathbb{R}^{n_i} \times \mathbb{R}^{N_{-i}} = \mathbb{R}^N$. ■

for $t_{-i} \in T_{-i}$. With this density function, we associate the probability distribution $b_i(\cdot|t_i, \nu)$ such that

$$b_i(B_{-i}|t_i, \nu) := \int_{B_{-i} \cap \hat{T}_{-i}} \beta_i(t_{-i}|t_i, \nu) dt_{-i} \tag{37}$$

for any measurable set $B_{-i} \subset \mathbb{R}^{N-i}$. The function $b_i(\cdot|t_i, \nu)$ is obviously a regular conditional probability distribution for t_{-i} given t_i .

By inspection of (36), the density $\beta_i(t_{-i}|t_i, \nu)$ depends continuously on $t_i \in T_i$ and $t_{-i} \in T_{-i}$. By standard arguments, again using Lebesgue's bounded-convergence theorem, it follows that the function $t_i \rightarrow b_i(\cdot|t_i, \nu)$ that is given by (37) maps the domain T_i of the marginal density f_i^ν continuously into $\mathcal{M}(T_{-i})$. By the same argument as in Lemma 3.3, this is the *only* regular conditional distribution for t_{-i} given t_i that maps T_i continuously into $\mathcal{M}(T_{-i})$.

By construction, we have $\mathcal{M}_+^d(T) \subset \mathcal{N}^c(T)$. The continuity property of the regular conditional distribution here is actually stronger than in the preceding analysis. For $\nu \in \mathcal{M}_+^d(T)$, the belief function $b_i(\cdot|t_i, \nu)$ actually takes values in the subspace $\mathcal{M}^d(T_{-i})$ of $\mathcal{M}(T_{-i})$ consisting of those measures that have densities that are continuous on T_{-i} . If we endow $\mathcal{M}^d(T_{-i})$ with the topology that is induced by the uniform topology for density functions, we find that the function $b_i(\cdot|t_i, \nu)$ maps the set T_i continuously into $\mathcal{M}^d(T_{-i})$, i.e., that $b_i(\cdot|t_i, \nu)$ is an element of the space $C(T_i, \mathcal{M}^d(T_{-i}))$ of continuous functions from T_i into $\mathcal{M}^d(T_{-i})$. If we endow $C(T_i, \mathcal{M}^d(T_{-i}))$ with the uniform topology, we obtain the following analogue of Proposition 5.3.

Proposition 5.6. *For $i = 1, \dots, I$, let T_i be a compact subset of \mathbb{R}^{n_i} and let $T_{-i} := \prod_{j \neq i} T_j$ and $T = \prod_{j=1}^I T_j$. Assume that the sets T_i have nonempty interiors \hat{T}_i . Then for any i , the set $\mathcal{E}(T_i, \mathcal{M}^d(T_{-i}))$ of embeddings of T_i in $\mathcal{M}^d(T_{-i})$ is a residual subset, i.e., it contains a countable intersection of open and dense subsets of $C(T_i, \mathcal{M}^d(T_{-i}))$.*

For a proof of Proposition 5.6, the reader is referred to Appendix A.2.

To translate this result into a proposition about priors, we endow the space $\mathcal{M}_+^d(T)$ with the coarsest topology under which the maps

$$\nu \rightarrow \psi_i(\nu) = (\bar{v}_i(\nu), b_i(\cdot, \nu)), \tag{38}$$

from distributions on T into marginal distributions and belief functions, are continuous. In this context, the range $\mathcal{M}_+^d(T_i)$ of the function $\nu \rightarrow \bar{v}_i(\nu)$ is taken to have the topology that is induced by the uniform topology for density functions, the range $C(T_i, \mathcal{M}^d(T_{-i}))$ of the function $\nu \rightarrow b_i(\cdot, \nu)$ the topology of uniform convergence.

Proposition 5.7. *For $i = 1, \dots, I$, let T_i be a compact subset of \mathbb{R}^{n_i} and let $T_{-i} := \prod_{j \neq i} T_j$ and $T = \prod_{j=1}^I T_j$. Assume that the sets T_i have nonempty interiors \hat{T}_i . Let $\mathcal{M}_+^d(T)$ is endowed with the coarsest topology under which the mapping*

$$\nu \rightarrow \psi(\nu) := (\psi_1(\nu), \dots, \psi_I(\nu)), \tag{39}$$

with $\psi_i(\nu) = (\bar{v}_i(\nu), b_i(\cdot, \nu))$, from $\mathcal{M}_+^d(T)$ to $\prod_{i=1}^I [C(T_i, \mathcal{M}_+^d(T_i)) \times C(T_i, \mathcal{M}^d(T_{-i}))]$ is continuous. Then the set $\mathcal{N}^{**}(T)$ of priors in $\mathcal{M}_+^d(T)$ that have the BDP property for all agents is a residual subset of $\mathcal{M}_+^d(T)$.

Proposition 5.7 follows from Proposition 5.6 by the same argument by which Proposition 5.5 was derived from Proposition 5.3. The details are left to the reader.

Corollary 5.8. *For $i = 1, \dots, I$, let T_i be a compact subset of \mathbb{R}^{n_i} and let $T = \prod_{j=1}^I T_j$. Assume that the sets T_i have nonempty interiors \hat{T}_i . If $\mathcal{M}_+^d(T)$ is endowed with the topology that is induced by the uniform topology for density functions, then the set $\mathcal{N}^{**}(T)$ of priors on T that have the BDP property for all agents is a residual subset of $\mathcal{M}_+^d(T)$.*

Corollary 5.8 follows immediately from Proposition 5.7 and the following lemma, the proof of which is given in Appendix A.3.

Lemma 5.9. *The topology that is induced by the topology of uniform convergence of density functions is the coarsest topology under which the mapping $\nu \rightarrow \psi(\nu)$ in Proposition 5.7 is continuous.*

6. Relation to the literature

The thrust of our results runs counter to that of Heifetz and Neeman (2006) and Barelli (2009), and parallels (Chen and Xiong, 2011). It is therefore appropriate to discuss the relation of our analysis to theirs. Heifetz and Neeman (2006) consider families $\{\mathcal{T}^k\}_{k \in K}$ of incomplete-information models of the form

$$\mathcal{T}^k = \{T_i^k, \hat{\theta}_i^k, \hat{\pi}_i^k\}_{i=1}^I \tag{40}$$

where, for any k in some index set K , for any i , T_i^k is a set of abstract "types" for agent i , $\hat{\theta}_i^k$ is a mapping from T_i^k into a set Θ_i^k of payoff parameter vectors for agent i , and $\hat{\pi}_i^k$ is a continuous mapping from T_i^k into the set $\mathcal{M}(T_{-i}^k)$ of probability distributions on the space T_{-i}^k of the other agents' abstract types. They restrict their attention to incomplete-information models that are consistent with common priors and study the genericity of the BDP property in the set

$$\mathcal{P} \subset \mathcal{M} \left(\prod_{i=1}^I \bigcup_{k \in K} T_i^k \right) \tag{41}$$

such that $F \in \mathcal{P}$ if and only if F is a common prior for one model T^k in the family that is being considered. Under the assumption that the family $\{\mathcal{T}^k\}_{k \in K}$ is "closed under finite unions", they show that the set \mathcal{P} is convex: If F^{k^1} and F^{k^2} are common priors for the incomplete-information models $\mathcal{T}^{k^1}, \mathcal{T}^{k^2}$, then, for any $\alpha \in [0, 1]$, the measure

$$\alpha F^{k^1} + (1 - \alpha) F^{k^2} \in \mathcal{M} \left(\prod_{i=1}^I (T_i^{k^1} \cup T_i^{k^2}) \right) \tag{42}$$

is a common prior for the model $\mathcal{T}^{\hat{k}} = \{(T_i^{k^1} \cup T_i^{k^2}), \hat{\theta}_i^{\hat{k}}, \hat{\pi}_i^{\hat{k}}\}_{i=1}^I$, where $\mathcal{T}^{\hat{k}}$ is the model that corresponds to the "union" of \mathcal{T}^{k^1} and \mathcal{T}^{k^2} , with $\hat{\theta}_i^{\hat{k}}, \hat{\pi}_i^{\hat{k}}$ specified so that

$$(\hat{\theta}_i^{\hat{k}}(t_i), \hat{\pi}_i^{\hat{k}}(t_i)) = (\hat{\theta}_i^{k^1}(t_i), \hat{\pi}_i^{k^1}(t_i)) \quad \text{if } t_i \in T_i^{k^1} \tag{43}$$

and

$$(\hat{\theta}_i^{\hat{k}}(t_i), \hat{\pi}_i^{\hat{k}}(t_i)) = (\hat{\theta}_i^{k^2}(t_i), \hat{\pi}_i^{k^2}(t_i)) \quad \text{if } t_i \in T_i^{k^2}. \tag{44}$$

Given this finding, they go on to show that any prior $F^{\hat{k}} \in \mathcal{P}$ that can be represented in the form

$$F^{\hat{k}} = \sum_{j=1}^J \alpha_j F^{k^j}, \tag{45}$$

with $\alpha_j > 0$ for all j , has the BDP property if and only if every one of the distributions F^{k^j} has the BDP property. This leads them to conclude that, unless the incomplete-information models $\mathcal{T}^k, k \in K$, admit *only* BDP priors, the set of non-BDP priors will be

geometrically and measure-theoretically generic in \mathcal{P} . Specifically, if the incomplete-information models \mathcal{T}^k , $k \in K$, admit one or more non-BDP priors, the set of BDP priors will be a proper face of the convex set \mathcal{P} . Moreover, under certain additional regularity conditions, the set of BDP priors will be finitely shy in \mathcal{P} .

In contrast to Heifetz and Neeman (2006) we follow a topological approach. However, with finite-dimensional abstract type spaces, the difference between our results and those of Heifetz and Neeman (2006) is not only a matter of topological versus geometric or measure-theoretic genericity but also of genericity in the full space versus genericity in a specially chosen subspace. Heifetz and Neeman (2006) are concerned with the genericity of non-BDP priors relative to the set of all priors that are associated with a fixed family $\{T^k\}_{k \in K}$ of incomplete-information models that is closed under finite unions. If the family $\{\mathcal{T}^k\}_{k \in K}$ itself is the set of common-prior models in the universal type space, this approach involves no loss of generality other than the presumption that a common prior exists. With finite-dimensional abstract type spaces, however, the requirement that the family $\{\mathcal{T}^k\}_{k \in K}$ be closed under finite unions is quite restrictive. In a companion paper, Gizatulina and Hellwig (2012), we show that, if the type spaces have finitely many dimensions and non-empty interiors, then a family that satisfies this requirement is at most countable; more precisely, such a family consists of unions of members of a countable family of models with type sets that do not intersect each other. As an implication of Proposition 5.7, we then find that the set of such families for which all models satisfy the BDP property contains a residual set. The set of families of models for which (Heifetz and Neeman, 2006) obtain geometric and measure-theoretic genericity on non-BDP priors is itself a sparse set.

Topological genericity of the BDP property is also discussed by Barelli (2009) and Chen and Xiong (2011). Barelli suggests that the measure-theoretic approach of Heifetz and Neeman (2006) is problematic and states a topological genericity result for non-BDP priors. Chen and Xiong (2011) point to an error in Barelli's argument and prove a topological genericity result for BDP priors. Chen and Xiong (2011) specify the type space $T = \prod_{i=1}^I T_i$ as the Θ -based universal type space, i.e., the space of payoff parameters and belief hierarchies that is generated by the payoff type space $\Theta = \prod_{i=1}^I \Theta_i$.²⁴ More recently, Chen and Xiong (2013) have extended the analysis of Chen and Xiong (2011) to show that priors allowing for full surplus extraction also form a residual set in the set of all common priors on the universal type space when this space is endowed with the weak* topology.

The results of Chen and Xiong are based on a result of Mertens et al. (1994), which shows that, if the universal type space is endowed with the product topology, then any incomplete information situation can be approximated by situations that involve only finitely many potential types for each player. Thus if the space of common priors on this space is endowed with the topology of weak convergence of probability measures, i.e., the weak* topology, the set of common priors with finite supports is dense. Because the set of BDP priors is generic in the set of finite priors and a dense subset of a dense set is itself dense in the ambient space, Chen and Xiong (2011) conclude that the set of BDP priors is dense in the space of all priors endowed with the weak* topology. Chen and Xiong (2013) use a similar argument to establish the genericity of priors that allow full surplus extraction.

The analysis of Chen and Xiong is more general than ours in one sense and more special in another. Whereas they consider the genericity of the BDP property in the context of the universal type space that admit common priors; we consider it only for finite-dimensional (though uncountable) abstract type spaces. However, their analysis is restricted to models with common priors; our analysis allows for models that do not admit common priors, as well as those that do.

Allowing for models that do not admit common priors is important if one thinks about incomplete-information models as having belief functions as a primitive, i.e., if one does not start from the hypothesis that, at some prior stage, agents had the same beliefs, as represented by a common prior, and that the differences in beliefs that we model are all due to differences in information.²⁵ If one thinks about belief functions as a primitive of an incomplete-information model, then, as explained in Section 3, it is appropriate to think about the BDP property as a property of belief functions. For priors, which in this approach are introduced as a derived concept, the BDP property is then introduced as a derived concept, as in our Definition 3.1.

In the universal type space, a meaningful assessment of the BDP property as a property of belief functions is hardly possible. We may think about the universal type space itself in terms of an abstract incomplete-information model $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$, but then all the relevant features of the model are encoded in the spaces T_i , i.e. the vectors specifying the agents' payoff parameters, probabilistic beliefs about the other agents' payoff parameters, probabilistic beliefs about the other agents' beliefs about the other agents' payoff parameters, etc. The mappings $\hat{\theta}_i$ and $\hat{\pi}_i$, by contrast, are fixed and given in that, for any element t_i of T_i , the payoff type $\theta_i(t_i)$ is given by the projection from T_i to the space of payoff parameters of agent i and the belief type π_i is given by the unique measure on T_{-i} that is compatible with the belief hierarchy of agent i that is given by the projection from T_i to the space of belief hierarchies of agent i . While, formally, the universal type space can be written in the form $\mathcal{T} = \{T_i, \hat{\theta}_i, \hat{\pi}_i\}_{i=1}^I$, the interpretation is very different from the usual abstract type space where it makes sense to ask how strategic behaviour would change if the belief functions $\hat{\pi}_i$ were to change.

In the universal type space, the BDP property can only be studied as a property of subsets of the space. A subset of the universal type space exhibits the BDP property for agent i if any two distinct elements of the set involve different beliefs of agent i about the other agents. Genericity of the BDP property must then be assessed as a property of subsets of the universal type space. In the absence of any additional restrictions on the set of subsets under consideration, this leads to the conclusion that, since T_i is a product of a space of payoff parameters and a space of belief hierarchies, the BDP property is not to be expected; failure of the BDP property is likely to be generic in any suitable topology on the set of subsets of the universal type space.

This conclusion, however, cannot be taken for granted if the set of subsets under consideration is taken to be the set of belief-closed subsets of the universal type space or even the set of belief-closed subsets of the universal type space that support common priors. Although the set of universal types of agent i has a product structure, the degrees of freedom provided by this structure are irrelevant if we consider a model in which agent i believes that agents $j \neq i$ hold beliefs under which the set of agent i 's types has the BDP property. In this approach, the question of genericity of the BDP property has bite because the consistency property of interactive belief systems that is inherent in the notion of subsets

²⁴ Barelli (2009) presents his analysis as if he was considering arbitrary type spaces. However, like Heifetz and Neeman (2006), he works with fixed belief functions. By the arguments in Gizatulina and Hellwig (2012), this implies that distinct indecomposable type sets cannot intersect. As he goes on to talk about convergence of sequences of models in terms of Hausdorff convergence of the corresponding sequences of type sets, he cannot be working with arbitrary abstract type spaces but must implicitly be working with the universal type space.

²⁵ On this point see Aumann (1987, 1998) and also Gul (1998).

of the universal type space being belief-closed imposes strong restrictions on the subsets of the universal type space that are being considered. If the existence of a common prior is imposed, these restrictions become even stronger.

However, as a property of belief functions, the BDP property is a feature of individual belief formation that does not depend on how the other agents form their beliefs. As such, it has little to do with the consistency conditions on interactive belief systems that are implied by considering belief-closed subsets of the universal type space, or by the existence of a common prior. The abstract type space approach has the advantage that it leaves room for studying the BDP property precisely as a feature of individual belief formation, without any need to impose a consistency condition on the implied interactive belief system.

These considerations point to a deeper issue concerning the universal type space. In the universal type space approach, belief hierarchies are taken as given, and there is no account of where beliefs come from. To be sure, the notion that beliefs are conditional distributions can be articulated in this approach as well as the abstract type space approach, simply by postulating a prior for which a regular conditional distribution is given by the map that associates to each type the unique measure over other agents' types that is compatible with the belief hierarchy encoded in the type. However, this is not a useful way to think about the dependence of beliefs on information.

To understand our concern, go back to the examples in Section 4. In each of those examples, we asked how the beliefs of agent 1 would vary with the payoff parameter θ_1 , and in particular, whether the value of the payoff parameter might be recoverable from a knowledge of the agent's beliefs. The underlying idea was that the type $t_i = (\theta_i, s_i)$ represents agent i 's information, and that the BDP property pertains to the way in which the agent uses this information to form his beliefs.

Suppose now that, along the lines of the universal type space approach, we treat the beliefs themselves as additional signals, i.e. as a part of agents' types. In Example 4.1 for instance, this would mean that we consider the "signal"

$$s_1^* := \frac{\text{cov}(s_2, \theta_1)}{\text{var}\theta_1}(\theta_1 - E_1\theta_1) + E_1s_2, \tag{46}$$

along with the payoff characteristic θ_1 and the signal s_1 that we had considered before. According to Eq. (17), s_1^* is just the agent's conditional expectation of the signal s_2 of the other agent given θ_1 and s_1 . Trivially, the agent's conditional expectation of s_2 given θ_1 and s_1^* satisfies the equation

$$E[s_2|\theta_1, s_1, s_1^*] \equiv s_1^*. \tag{47}$$

One might therefore infer that the payoff parameter θ_1 is irrelevant for the agent's expectations formation. It would indeed be irrelevant if s_1^* were a real signal. However if s_1^* is merely a notational device, introduced as a shorthand in order to treat the agent's beliefs as a part of type, the payoff parameter θ_1 is not at all irrelevant for the agent's expectations formation because in fact his expectations are based on his observation of θ_1 . The dependence of his beliefs on θ_i is however lost out of sight if one thinks about his beliefs in terms of Eq. (47) only.

Whether we think of s_1^* as a real signal or as a notational device, in both case, Eq. (46) remains valid, and with it the BDP property, but by looking at (47), one would not know it. To assess the BDP property one needs to know how the agents' conditional expectations are related to his payoff parameters and more generally "real" signals that he observes. In the universal type space, encoding beliefs in types gives rise to more sophisticated versions of Eq. (47), which raise the same kinds of problems.

Our interest in the BDP property as a property of belief functions is one reason why, in contrast to Chen and Xiong, we do not

rely on finite approximations. Finite approximations of continuous type spaces may give rise to information discontinuities. To see why, consider any function that is defined on the real numbers and consider also the restrictions of this function to finite subsets of the real numbers. Regardless of whether the function itself is injective (has the BDP property) or not, restrictions of the function to finite subsets of its domain will generically be injective. For an illustration, go back to Example 4.4, where agent 2 has a noisy signal $s_2 = (\theta_1)^2 + \varepsilon$ and agent 1's belief about the signal s_2 reveals $(\theta_1)^2$ but not θ_1 . If the support of agent 1's type distribution was the finite set

$$\left\{ -\frac{2n}{\sqrt{n}}, -\frac{2(n-1)}{\sqrt{n}}, \dots, -\frac{2}{\sqrt{n}}, 0, \frac{1}{\sqrt{n}}, \frac{1+2}{\sqrt{n}}, \dots, \frac{1+2(n-1)}{\sqrt{n}}, \frac{1+2n}{\sqrt{n}} \right\} \tag{48}$$

and the support of the noise term ε contains only even multiples of $\frac{1}{\sqrt{n}}$, the signal s_2 would actually reveal θ_1 and so would agent 1's belief about s_2 . In this case, if s_2 is an even multiple of $\frac{1}{\sqrt{n}}$, agent 2 will know that θ_1 is negative; if s_2 is an odd multiple of $\frac{1}{\sqrt{n}}$, agent 2 will know that θ_1 is positive. If we consider a sequence of type distributions with supports (48) approximating the continuous type distribution, there is an informational discontinuity because, in the continuous limit, agent 2 cannot distinguish whether θ_1 is negative or positive.²⁶

The informational discontinuity in the transition from finite models to the continuous models is likely to introduce a discontinuity in strategic behaviour in any game of incomplete information that involves these type spaces. Specifically, if $A_i, i = 1, 2$, are the action sets and $u_i : A_1 \times A_2 \times T_1 \times T_2 \rightarrow \mathbb{R}$ are the payoff functions of the two agents, then in the transition from the models with finite type sets for agent 1, as given by (48), to the continuous model in Example 4.4, the interim Nash equilibrium correspondence violates upper hemi-continuity as well as approximate lower hemi-continuity in all but most trivial specifications.²⁷

The use by Chen and Xiong of finite approximations and the weak* topology on the space of priors associated with the product topology on the universal type space is also problematic for other reasons. The product topology on the universal type space has been criticized on the grounds that, in this topology, the correspondence of interim strictly ε -rationalizable strategies is not generally lower hemi-continuous in types. The product topology is strictly coarser than the strategic topology proposed by Dekel et al. (2006), i.e., the coarsest topology on the universal type space in which, for every $\varepsilon > 0$, the correspondence of interim strictly ε -rationalizable strategies is lower hemi-continuous in types and the correspondence of rationalizable strategies is upper hemi-continuous. In this topology, however, models with finitely many types are not dense, so the argument of Chen and Xiong (2011) is not available.²⁸

²⁶ Interestingly, this informational discontinuity does not appear in a universal type space representation of the example. With beliefs encoded in types, weak convergence of type distributions implies that the equation $E_2[\theta_1|t_2] = \theta_1$ should hold in the continuous limit as well as the finite approximations. The fact that the signal s_2 only reveals $(\theta_1)^2$ is irrelevant if the agent is just assumed to "observe" his own belief $s_2^* := E_2[\theta_1|t_2]$.

²⁷ For a discussion of the need for uniform convergence of conditional distributions to obtain approximate lower as well as upper hemi-continuity of the interim Nash equilibrium correspondence, see Kajii and Morris (1994).

²⁸ An important question is how our topology on belief functions relates to the different topologies on the universal type space. We are currently doing research on that. Within the context of abstract type spaces, we note that, by the arguments

7. Concluding remarks

We conclude this paper with several additional remarks. First, the genericity properties established in Propositions 5.6 and 5.7 are still obtained if we replace the space $\mathcal{M}_+^d(T)$ by the space $\mathcal{M}_+^d(\dot{T})$ of measures with densities that are defined on the interior of T and if we endow the spaces $\mathcal{M}_+^d(\dot{T}_i)$ and $\mathcal{M}^d(\dot{T}_{-i})$ of marginal distributions on the interiors of the type spaces of agent i and agent other than i with the topologies that is induced by the compact open topologies, rather than the uniform topologies for the associated density functions. In this case, however, an analogue of Corollary 5.8 is only obtained if an additional condition of uniform boundedness is imposed on the density functions of the priors under consideration.

Second, in Gizatulina and Hellwig (2012), we show that the genericity properties established in Propositions 5.6 and 5.7 hold also for an infinite-dimensional type space of the form $T = \cup_{\ell=1}^\infty T^\ell$, where, for each ℓ , $T^\ell = \prod_{i=1}^I T_i^\ell$ is a finite-dimensional set. This corresponds to a specification with *ex ante* uncertainty about the incomplete-information model that is going to be relevant. If the set of such models is countable and each model has a finite-dimensional type space, the BDP property is still generic.

Third, the assumption that the type sets T_i in Proposition 5.7 are taken as fixed and given is not essential. In principle, the set T in Proposition 5.7 could be any compact subset of \mathbb{R}^N that has a nonempty interior \dot{T} . Given a measure ν with a density that is continuous on \dot{T} , the belief function $b_i(\cdot|\cdot, \nu)$ is defined on the set D_i^ν of types of agent i for which the marginal density $\bar{f}_i^\nu(t_i)$ is strictly positive. D_i^ν is an element of the set \mathcal{D}_i^ν of open subsets of \mathbb{R}^{n_i} that have compact closures. Thus $b_i(\cdot|\cdot, \nu)$ can be treated as an element of the space $\cup_{D \in \mathcal{D}_i^\nu} C(D, \mathcal{M}^d(T_{-i}))$.

If the topology on $\cup_{D \in \mathcal{D}_i^\nu} C(D, \mathcal{M}^d(T_{-i}))$ is specified so that a sequence $\{b_i^k\}$ of functions in $\cup_{D \in \mathcal{D}_i^\nu} C(D, \mathcal{M}^d(T_{-i}))$ converges to a limit $b_i \in \cup_{D \in \mathcal{D}_i^\nu} C(D, \mathcal{M}^d(T_{-i}))$ if and only if, every compact subset K of the domain D_{b_i} of b_i is also a subset of the domains $D_{b_i^k}$ of the functions b_i^k for any sufficiently large k , and, moreover, $\lim_{k \rightarrow \infty} b_i^k(t_i) = b_i(t_i)$, uniformly on K , an extension of the argument given to prove Proposition 5.6 can be used to show that the set $\cup_{D \in \mathcal{D}_i^\nu} \mathcal{E}(D, \mathcal{M}^d(T_{-i}))$ of embeddings is a residual subset of $\cup_{D \in \mathcal{D}_i^\nu} C(D, \mathcal{M}^d(T_{-i}))$. A similar generalization of Proposition 5.3 is also available.

Fourth, in a previous version of this paper, we had applied the Embedding Theorem to the conditional-expectations functions $\bar{t}_{-i}^\nu(\cdot)$, which are given by the formula

$$\bar{t}_{-i}^\nu(t_i) := \int_{T_{-i}^\nu} t_{-i} \beta_i(t_{-i}|t_i, \nu) dt_{-i} = \frac{\int_{T_{-i}^\nu} t_{-i} f_i^\nu(t_i, t_{-i}) dt_{-i}}{\bar{f}_i^\nu(t_i)}$$

Under the dimensionality assumption that $2n_i + 1 \leq N_{-i}$, the Embedding Theorem for continuous functions implies that, for any $\varepsilon > 0$ there exists an embedding $t_i \rightarrow \hat{t}_{-i}^\varepsilon(t_i)$ such that $\|\hat{t}_{-i}^\varepsilon(t_i) - \bar{t}_{-i}^\nu(t_i)\| < \varepsilon$ for all $t_i \in D_i^\nu$. Given this embedding, approximating densities can be defined by writing

$$f^{v^\varepsilon}(t_i, t_{-i}) = \beta(t_{-i} + \hat{t}_{-i}^\varepsilon(t_i) - \bar{t}_{-i}^\nu(t_i)|t_i) \bar{f}_i^\nu(t_i) \quad \text{if } t_i \in D_i^\nu, \quad (49)$$

and

$$f^{v^\varepsilon}(t_i, t_{-i}) = 0 \quad \text{if } t_i \notin D_i^\nu. \quad (50)$$

of Kajii and Morris (1994), for any specification of action sets and payoff functions of the different agents, the interim Nash equilibrium correspondence of the induced game will be upper hemi-continuous and approximately lower hemi-continuous as the belief functions of the different agents vary in $C(T_i, \mathcal{M}^d(T_{-i}))$.

The BDP property of approximating priors is then directly embodied in the conditional expectations, and there is no need to appeal to the fact that probability distributions themselves are infinite-dimensional objects.

Finally, if we strengthen the assumptions on the set of priors under consideration so that we can use a differentiable approach, then, under the dimensionality assumption that $2n_i + 1 \leq N_{-i}$, we can strengthen the claim that BDP priors are generic from residual to open and dense. Specifically, if \bar{N}^{cd} is a subset of \bar{N}^c so that, for any $\nu \in \bar{N}^{cd}$, the support T^ν of ν is a compact manifold and the conditional-expectations functions \bar{t}_{-i}^ν are continuously differentiable on T_i^ν , and if the topology on \bar{N}^{cd} is such that, for all i , the map from priors into agent i 's conditional-expectations functions is continuous when the range of this map is given the strong C^1 topology, then, under the dimensionality assumption in Proposition 5.7, the set of common priors exhibiting the BDP property is open and dense in \bar{N}^{cd} . The proof is basically the same as the proof of Proposition 5.7, except that the Embedding Theorem for continuous functions with compact domains must be replaced by Whitney's Embedding Theorem for C^1 functions, see, e.g., Hirsch (1994, p. 35).

Appendix

In this appendix, we state and prove the embedding theorems that are used in Section 5. For any separable metric space X , a subset of X is said to be residual in X if it contains a countable intersection of open and dense subsets of X . For any two separable metric spaces X and Y , $C(X, Y)$ will denote the space of continuous functions from X to Y . The space $C(X, Y)$ will be endowed with the uniform topology. The set of embeddings of X into Y will be denoted as $\mathcal{E}(X, Y)$.

A.1. Proof of Proposition 5.3

Proposition 5.3 in the text is an instance of the following result, with $X = T_i$ and $Z = T_{-i}$.

Proposition A.1. *Let $X \subset \mathbb{R}^n$ be compact, and let Z be a compact metric space with infinitely many elements. Let $Y = \mathcal{M}(Z)$ be the space of probability measures on Z , endowed with the topology of weak convergence of probability measures. Then the set $\mathcal{E}(X, Y)$ of embeddings of X in Y is a residual subset of $C(X, Y)$.*

In proving Proposition A.1, we will repeatedly use the following lemma from topology.

Lemma A.2. *Let X, Y be any two topological spaces and let g be a continuous and open function from X to Y . Then, for any open and dense set $U \subset Y$, the inverse image $g^{-1}(U)$ is an open and dense subset of X . For any residual set $Y^r \subset Y$, the inverse image $g^{-1}(Y^r)$ is a residual subset of X .*

Proof. Continuity of $g : X \rightarrow Y$ implies that $g^{-1}(U)$ is open in X whenever U is open in Y . Further, openness of g implies that $g(V)$ is open in Y whenever V is open in X . Thus, if V is open in X and U is dense in Y , we must have $g(V) \cap U \neq \emptyset$. For any $y \in g(V) \cap U$, there exists $v \in V$ such that $y = g(v) \in U$ and, $v \in V \cap g^{-1}(U)$. Thus, $g(V) \cap U \neq \emptyset$ implies $V \cap g^{-1}(U) \neq \emptyset$. It follows that $g^{-1}(U)$ is dense in X whenever U is dense in Y . The first statement of the lemma is proved.

To prove the second statement, it suffices to observe that, if $Y^r \subset Y$ satisfies $Y^r \supset \cap_{k=1}^\infty U_k$ for some sequence $\{U_k\}$ of open and dense subsets of Y , then

$$g^{-1}(Y^r) \supset g^{-1}(\cap_{k=1}^\infty U_k) = \cap_{k=1}^\infty g^{-1}(U_k),$$

and, by the first statement of the lemma, the sets $g^{-1}(U_k)$, $k = 1, 2, \dots$, are open and dense in X . ■

Lemma A.3. Let X and Y be as specified in Proposition A.1. If there exists a separable metric space Q that is homeomorphic to $[0, 1]^{2n+1}$ and if there exists a mapping Φ from $C(X, Y)$ to $C(X, Q)$ that is continuous and open, then the set $\mathcal{E}(X, Y)$ of embeddings of X in Y is a residual subset of $C(X, Y)$.

Proof. Let h be a homeomorphism from Q to $[0, 1]^{2n+1}$. We claim that the formula

$$H(f) := h \circ f$$

defines a homeomorphism H from $C(X, Q)$ to $C(X, [0, 1]^{2n+1})$. To establish this claim, we note that, because h is continuous, H takes values in $C(X, [0, 1]^{2n+1})$. Continuity of h also implies that H is continuous. Because h is invertible, the inverse of H is well defined, with

$$H^{-1}(\hat{f}) = h^{-1} \circ \hat{f}$$

for any $\hat{f} \in [0, 1]^{2n+1}$. Continuity of h^{-1} implies that H^{-1} takes values in $C(X, Q)$ and that H^{-1} is continuous.

By the classical Embedding Theorem,²⁹ the set $\mathcal{E}(X, [0, 1]^{2n+1})$ of embeddings of X in $[0, 1]^{2n+1}$ is a residual subset of $C(X, [0, 1]^{2n+1})$. By Lemma A.2 therefore, $H^{-1}(\mathcal{E}(X, [0, 1]^{2n+1}))$ is a residual subset of $C(X, Q)$. If the mapping Φ from $C(X, Y)$ to $C(X, Q)$ is continuous and open, then, again by Lemma A.2, it follows that the set $\Phi^{-1}(H^{-1}(\mathcal{E}(X, [0, 1]^{2n+1})))$ is a residual subset of $C(X, Y)$.

To prove the lemma, it therefore suffices to show that

$$\Phi^{-1}(H^{-1}(\mathcal{E}(X, [0, 1]^{2n+1}))) \subset \mathcal{E}(X, Y). \tag{51}$$

For this purpose, suppose that $f \in \Phi^{-1}(H^{-1}(\mathcal{E}(X, [0, 1]^{2n+1}))) \setminus \mathcal{E}(X, Y)$. Then there exist $x_1, x_2 \in X$ such that $f(x_1) = f(x_2)$ and $h(\Phi(x_1|f)) \neq h(\Phi(x_2|f))$. Because h is a homeomorphism, $h(\Phi(x_1|f)) \neq h(\Phi(x_2|f))$ implies $\Phi(x_1|f) \neq \Phi(x_2|f)$, which in turn implies $f(x_1) \neq f(x_2)$. The assumption that $\Phi^{-1}(H^{-1}(\mathcal{E}(X, [0, 1]^{2n+1}))) \setminus \mathcal{E}(X, Y) \neq \emptyset$ thus leads to a contradiction, which proves (51). The lemma follows immediately. ■

To prove Proposition A.1, it thus suffices to show that there exist Q and Φ as specified in the lemma. We define Q as the unit simplex in \mathbb{R}_+^{2n+1} , i.e.,

$$Q := \left\{ \mathbf{q} \in \mathbb{R}_+^{2n+1} \mid \sum_{i=1}^{2n+1} q_i \leq 1 \right\}. \tag{52}$$

Lemma A.4. The set Q that is defined by (52) is homeomorphic to $[0, 1]^{2n+1}$.

Proof. Because both Q and $[0, 1]^{2n+1}$ are compact convex subset of \mathbb{R}_+^{2n+1} and have nonempty interiors, they are both homeomorphic to the closed unit ball in \mathbb{R}_+^{2n+1} and therefore to each other.³⁰ ■

To construct the mapping Φ , we proceed as follows: Exploiting the fact that Z has infinitely many elements, we fix $2n + 1$ distinct elements z_1, \dots, z_{2n+1} of Z . For this purpose, we note that, as compact metric space, Z is separable. Therefore, there exists some $\varepsilon > 0$ so that the open ε -balls $B_\varepsilon(z_1), \dots, B_\varepsilon(z_{2n+1})$ around z_1, \dots, z_{2n+1} do not intersect each other, i.e., $B_\varepsilon(z_i) \cap B_\varepsilon(z_j) = \emptyset$ for all i and all $j \neq i$. Moreover, by Urysohn's lemma, there exist continuous functions g_1, \dots, g_{2n+1} from Z into $[0, 1]$ such that, for

any i , $g_i(z_i) = 1$ and $g_i(z) = 0$ for $z \in Z \setminus B_\varepsilon(z_i)$. For any $z \in Z$, therefore, $\sum_{i=1}^{2n+1} g_i(z) \in [0, 1]$, and, for any $\mu \in \mathcal{M}(Z)$, the vector

$$\varphi(\mu) = \left(\int_Z g_1(z) d\mu(z), \dots, \int_Z g_{2n+1}(z) d\mu(z) \right) \tag{53}$$

is an element of the set Q that is defined by (52). Given the mapping $\varphi : \mathcal{M}(Z) \rightarrow Q$, the formula

$$\Phi(b) = \varphi \circ b \tag{54}$$

defines a mapping from $C(X, \mathcal{M}(Z))$ into the space of functions from X to Q . We need to show that $\Phi(b) \in C(X, Q)$ for all $b \in C(X, \mathcal{M}(Z))$ and that the mapping Φ is continuous and open. We begin by showing that the mapping φ has these properties.

Lemma A.5. The mapping $\varphi : \mathcal{M}(Z) \rightarrow Q$ that is defined by formula (53) is continuous and open.

Proof. Continuity of φ is immediate from (53) and the definition of the topology of weak convergence.

Note that, if Z is a compact metric space and if the space $Y = \mathcal{M}(Z)$ of probability measures on Z is endowed with the topology of weak convergence of probability measures, i.e., the weak* topology, then Y is a compact metric space.³¹ If $\{h_1, h_2, \dots\}$ is a countable dense set of continuous functions from Z into $[0, 1]$, then, by an argument (Parthasarathy, 1967, p. 43), the mapping

$$\mu \rightarrow h(\mu) := \left(\int_Z h_1(z) d\mu(z), \int_Z h_2(z) d\mu(z), \dots \right)$$

defines a homeomorphism between $\mathcal{M}(Z)$ and a subspace S of the infinite product $[0, 1]^\infty$. Without loss of generality, we may assume that $h_i = g_i$ for $i = 1, \dots, 2n + 1$. Let $\pi_{2n+1} : [0, 1]^\infty \rightarrow [0, 1]^{2n+1}$ be the natural projection from $[0, 1]^\infty$ to the first $2n + 1$ factors of the infinite product. Then

$$\varphi = \pi_{2n+1} \circ h.$$

Because the homeomorphism $h : \mathcal{M}(Z) \rightarrow S$ is open and the composition of two open mappings is also open, openness of the mapping $\varphi : \mathcal{M}(Z) \rightarrow Q$ will follow if the restriction π_{2n+1}^S of π_{2n+1} to S is shown to be an open mapping from S to Q .

For this purpose, we note that φ maps $\mathcal{M}(Z)$ onto Q , i.e., that $\varphi(\mathcal{M}(Z)) = Q$: For any $\mathbf{q} = (q_1, \dots, q_{2n+1}) \in Q$, any $\mu \in \mathcal{M}(Z)$ such that $\mu(\{z_i\}) = q_i$ for $i = 1, 2, \dots, 2n + 1$, and $\mu(Z \setminus \cup_i B_\varepsilon(z_i)) = 1 - \sum_{i=1}^{2n+1} q_i$ satisfies $\varphi(\mu) = \mathbf{q}$. Thus, we obtain

$$\pi_{2n+1}(S) = \pi_{2n+1} \circ h(\mathcal{M}(Z)) = \varphi(\mathcal{M}(Z)) = Q.$$

Now let V be any open subset of S . By definition of the subspace topology for S , there exists an open set $U \subset [0, 1]^\infty$ such that $V = S \cap U$. Thus,

$$\begin{aligned} \pi_{2n+1}^S(V) &= \pi_{2n+1}^S(S \cap U) = \pi_{2n+1}(S \cap U) \\ &= \pi_{2n+1}(S) \cap \pi_{2n+1}(U) = Q \cap \pi_{2n+1}(U). \end{aligned}$$

Because the projection π_{2n+1} is an open mapping from $[0, 1]^\infty$ to $[0, 1]^{2n+1}$, the set $\pi_{2n+1}(U)$ is an open subset of $[0, 1]^{2n+1}$. Therefore, the set $\pi_{2n+1}^S(V) = Q \cap \pi_{2n+1}(U)$ is open in the subspace topology for Q as a subset of $[0, 1]^{2n+1}$. Thus, π_{2n+1}^S maps any open subset of $S = h(\mathcal{M}(Z))$ into an open subset of Q . This proves that the mapping $\pi_{2n+1}^S : S \rightarrow Q$ is open. The mapping $\varphi = \pi_{2n+1}^S \circ h : \mathcal{M}(Z) \rightarrow Q$ is therefore also open. ■

Lemma A.6. The mapping Φ that is defined by formula (54) maps $C(X, \mathcal{M}(Z))$ continuously into $C(X, Q)$.

²⁹ See Theorem V.2, p. 56, in Hurewicz and Wallman (1941).

³⁰ See Proposition 4.26 in Lee (2000).

³¹ See Theorem 6.4, p. 45, in Parthasarathy (1967).

Proof. Because φ is continuous, obviously, $\Phi(b) \in C(X, Q)$ if $b \in C(X, \mathcal{M}(Z))$. To prove that Φ is continuous, we note that, as mentioned above, if Z is a compact metric space, then $\mathcal{M}(Z)$ is a compact metric space. A metric ρ for $\mathcal{M}(Z)$ is given by the formula

$$\rho(\mu, \nu) = \sup_k \left| \int h_k(z\nu) d\mu(z) - \int h_k(z) d\nu(z) \right|, \quad (55)$$

where $\{h_1, h_2, \dots\}$ is a countable dense subset of the set of continuous functions from Z into $[0, 1]$. Again there is no of generality in assuming that $h_k = g_k$ for $k = 1, \dots, 2n + 1$.

For any $\gamma > 0$, define $\eta(\gamma) := \frac{\gamma}{2n+1}$. Then, for any μ and ν in $\mathcal{M}(Z)$, $\rho(\mu, \nu) < \eta(\gamma)$ implies

$$\left| \int h_k(z) d\mu(z) - \int h_k(z) d\nu(z) \right| < \frac{\gamma}{2n+1}$$

for all k . Thus, $\rho(\mu, \nu) < \eta(\gamma)$ implies $\|\varphi(\mu) - \varphi(\nu)\| \leq \gamma$, where $\|\cdot\|$ is the Euclidean norm on \mathbb{R}^{2n+1} . If b and \hat{b} in $C(X, \mathcal{M}(Z))$ are such that

$$\sup_{x \in X} \rho(b(x), \hat{b}(x)) < \eta(\gamma) = \frac{\gamma}{2n+1},$$

it follows that

$$\sup_{x \in X} \|\varphi(b(x)) - \varphi(\hat{b}(x))\| \leq \gamma,$$

i.e., if b and \hat{b} are $\eta(\gamma)$ -close, then $\Phi(b)$ and $\Phi(\hat{b})$ are γ -close. This proves that Φ is continuous. ■

Lemma A.7. The mapping $\Phi : C(X, \mathcal{M}(Z)) \rightarrow C(X, Q)$ that is defined by (54) is open.

Proof. By Proposition 2.2.b, p. 269, in Clausing (1978), the lemma follows from the fact that $Y = \mathcal{M}(Z)$ and Q are compact convex sets and that the mapping $\varphi : Y \rightarrow Q$ is affine, as well as continuous and open. ■

Proposition A.1 now follows from Lemmas A.3, A.6 and A.7.

A.2. Proof of Proposition 5.6

Proposition 5.6 in the text is an instance of the following result, with $X = T_i$ and $Z = T_{-i}$.

Proposition A.8. Let $X \subset \mathbb{R}^n$, and let $Z \subset \mathbb{R}^{N-n}$ be compact sets with nonempty interiors. Let $Y = \mathcal{M}(Z)$ be the space of probability measures on Z that have continuous density functions and let the topology on Y be induced by the uniform topology for density functions. Then the set $\mathcal{E}(X, Y)$ of embeddings of X in Y is a residual subset of $C(X, Y)$.

The proof proceeds along the same lines as the proof of Proposition A.1. We note that $\mathcal{M}^d(Z)$ is a separable metric space. A metric ρ for $\mathcal{M}^d(Z)$ is given by the formula

$$\rho(\mu, \nu) = \sup_{z \in Z} |f_\mu(z) - f_\nu(z)|, \quad (56)$$

where f_μ and f_ν are the continuous densities associated with μ and ν .

The conclusion of Lemma A.3 remains valid with $X \subset \mathbb{R}^n$ and $Y = \mathcal{M}^d(Z)$ having the topology that is induced by uniform convergence of density functions. To prove Proposition A.8, it therefore suffices to specify a mapping $\hat{\Phi} : C(X, \mathcal{M}^d(Z)) \rightarrow C(X, [0, 1]^{2n+1})$ that is continuous and open. We do so by setting

$$\hat{\Phi}(b) = \hat{\varphi} \circ b \quad (57)$$

for $b \in C(X, \mathcal{M}^d(Z))$, where $\hat{\varphi} : \mathcal{M}^d(Z) \rightarrow [0, 1]^{2n+1}$ is given by the formula:

$$\hat{\varphi}(\mu) = \left(\frac{f_\mu(z_1)}{1 + f_\mu(z_1)}, \dots, \frac{f_\mu(z_{2n+1})}{1 + f_\mu(z_{2n+1})} \right), \quad (58)$$

where z_1, \dots, z_{2n+1} is an arbitrary but fixed collection of distinct elements of the interior of Z .

Lemma A.9. The mapping $\hat{\varphi} : \mathcal{M}^d(Z) \rightarrow [0, 1]^{2n+1}$ is continuous and open.

Proof. Continuity is immediate from (58). Openness follows from observing that $\hat{\varphi}$ is the composition of the homeomorphism $\mu \rightarrow f_\mu$ between $\mathcal{M}^*(Z)$ and $C(Z, \mathbb{R}_+)$, the projection $f \rightarrow (f(z_1), \dots, f(z_{2n+1}))$ from $C(Z, \mathbb{R}_+)$ to \mathbb{R}_+^{2n+1} , and the homeomorphism

$$(f_1, \dots, f_{2n+1}) \rightarrow \left(\frac{f_1}{(1 + f_1)}, \dots, \frac{f_{2n+1}}{(1 + f_{2n+1})} \right)$$

from \mathbb{R}_+^{2n+1} to the product $[0, 1]^{2n+1}$. ■

Lemma A.10. The function $\hat{\Phi}$ that is defined by (57) and (58) maps $C(X, \mathcal{M}^d(Z))$ continuously into $C(X, [0, 1]^{2n+1})$.

Proof. By routine calculations, (58) implies

$$\|\hat{\varphi}(\mu) - \hat{\varphi}(\bar{\mu})\| \leq (2n + 1) \max_i |f_\mu(z_i) - f_{\bar{\mu}}(z_i)|$$

for any μ and $\bar{\mu}$ in $\mathcal{M}^d(Z)$, where $\|\cdot\|$ is again the Euclidean norm on \mathbb{R}^{2n+1} . For any μ and $\bar{\mu}$ in $\mathcal{M}^d(Z)$, therefore,

$$\|\hat{\varphi}(\mu) - \hat{\varphi}(\bar{\mu})\| \leq (2n + 1)\rho(\mu, \bar{\mu}).$$

The function $\hat{\varphi}$ thus is uniformly continuous. By the same argument as in the proof of Lemma A.6, it follows that $\hat{\Phi}$ takes values in $C(X, [0, 1]^{2n+1})$ and that $\hat{\Phi}$ is continuous. ■

The proof that $\hat{\Phi}$ is also open is more involved. The range $\mathcal{M}^d(Z)$ of belief functions now consists of measures with continuous density functions and is therefore not compact.³² Therefore we cannot rely on the result of Clausing (1978) to infer that $\hat{\Phi}$ is open if $\hat{\varphi}$ is open. Instead we need a new argument.

Lemma A.11. For any $(\mathbf{q}, \mu) \in [0, 1]^{2n+1} \times \mathcal{M}^d(Z)$, the infimum $\rho^*(\mathbf{q}, \mu)$ of the distance $\rho(\mu, \nu)$ over the set of measures $\nu \in \mathcal{M}^d(Z)$ that satisfy

$$\hat{\varphi}(\nu) = \mathbf{q} \quad (59)$$

is well defined and satisfies

$$\rho^*(\mathbf{q}, \mu) = \max_i \left| \frac{q_i}{1 - q_i} - f_\mu(z_i) \right|. \quad (60)$$

Proof. Fix $\mathbf{q} \in [0, 1]^{2n+1}$ and $\mu \in \mathcal{M}^d(Z)$ and let f_μ be the density of μ . For any $\varepsilon > 0$ and any i , let $B_\varepsilon(z_i)$ be the open ε -ball around z_i and let $g_i^\varepsilon \in C(Z, [0, 1])$ be such that $g_i^\varepsilon(z_i) = 1$ and $g_i^\varepsilon(z) = 0$ for all $z \notin B_\varepsilon(z_i)$. For any $\varepsilon > 0$ and any $\alpha > 0$, consider the function $\hat{f}_{\alpha\varepsilon}$ that is given by the formula

$$\begin{aligned} \hat{f}_{\alpha\varepsilon}(z) &= \max \left[0, \alpha f_\mu(z) + \sum_{i=1}^{2n+1} \left(\frac{q_i}{1 - q_i} - \alpha f_\mu(z_i) \right) g_i^\varepsilon(z) \right]. \end{aligned} \quad (61)$$

³² Compactness would require that the density functions belong to an equicontinuous set of functions.

By construction, the function $\hat{f}_{\alpha\varepsilon}(\cdot)$ from Z into \mathbb{R}_+ is continuous. Thus, if ε and α are such that

$$\int_Z \hat{f}_{\alpha\varepsilon}(z) dz = 1, \tag{62}$$

then $\hat{f}_{\alpha\varepsilon}(\cdot)$ is the density of a measure $\nu_{\alpha\varepsilon} \in \mathcal{M}^d(Z)$.

We claim that, if $\varepsilon > 0$ is sufficiently small, there exists $\alpha(\varepsilon)$ such that Eq. (62) holds for $\alpha = \alpha(\varepsilon)$ and ε . Moreover, if ε is close to zero, $\alpha(\varepsilon)$ is close to one. To prove this claim, we note that, because $\int_Z f_\mu(z) dz = 1$, (61) implies

$$\begin{aligned} \int_Z \hat{f}_{\alpha\varepsilon}(z) dz &\leq \alpha + \sum_{i=1}^{2n+1} \frac{q_i}{1-q_i} \int_Z g_i^\varepsilon(z) dz \\ &\leq \alpha + H(\varepsilon) \end{aligned} \tag{63}$$

and

$$\begin{aligned} \int_Z \hat{f}_{\alpha\varepsilon}(z) dz &\geq \alpha + \sum_{i=1}^{2n+1} \left[\frac{q_i}{1-q_i} - f_\mu(z_i) \right] \int_Z g_i^\varepsilon(z) dz \\ &\geq \alpha - H(\varepsilon), \end{aligned} \tag{64}$$

where

$$\begin{aligned} H(\varepsilon) &:= \sum_{i=1}^{2n+1} \max \left[\frac{q_i}{1-q_i}, f_\mu(z_i) \right] \int_Z g_i^\varepsilon(z) dz \\ &\leq \sum_{i=1}^{2n+1} \max \left[\frac{q_i}{1-q_i}, f_\mu(z_i) \right] \int_{B_\varepsilon(z_i)} dz. \end{aligned} \tag{65}$$

Then, for $\alpha \leq 1 - H(\varepsilon)$, (63) and (65) imply $\int_Z \hat{f}_{\alpha\varepsilon}(z) dz \leq 1$ and, for $\alpha \leq 1 + H(\varepsilon)$, (64) and (65) imply $\int_Z \hat{f}_{\alpha\varepsilon}(z) dz \geq 1$. By the intermediate value theorem, there exists $\alpha(\varepsilon) \in [1 - \eta, 1 + \eta]$ such that $\int_Z \hat{f}_{\alpha(\varepsilon)}(z) dz = 1$. Moreover, by (65), $\lim_{\varepsilon \rightarrow 0} H(\varepsilon) = 0$ and, therefore, $\lim_{\varepsilon \rightarrow 0} \alpha(\varepsilon) = 1$.

Consider the measures $\nu_{\alpha(\varepsilon)\varepsilon} \in \mathcal{M}^d(Z)$ that are induced by the density functions $\hat{f}_{\alpha(\varepsilon)\varepsilon}$ for $\varepsilon > 0$ sufficiently small and $\alpha(\varepsilon)$ chosen so that (62) holds. Because the points z_1, \dots, z_{2n+1} in Z are distinct, there exists $\bar{\varepsilon} > 0$ such that, for $\varepsilon < \bar{\varepsilon}$, no two of the open balls $B_\varepsilon(z_i)$ intersect each other. In this case, (61) implies

$$\hat{f}_{\alpha\varepsilon}(z_i) = \frac{q_i}{1-q_i} \tag{66}$$

for all i . For $\alpha = \alpha(\varepsilon)$, it follows that the measure $\nu_{\alpha(\varepsilon)\varepsilon} \in \mathcal{M}^d(Z)$ satisfies $\hat{\varphi}(\nu_{\alpha\varepsilon}) = \mathbf{q}$. This implies, in particular, that the set of measures satisfying (59) is nonempty so that the infimum $\rho^*(\mathbf{q}, \mu)$ is well defined.

For any one of the measures $\nu_{\alpha(\varepsilon)\varepsilon}$ with $\varepsilon < \bar{\varepsilon}$, we compute

$$\begin{aligned} \rho(\mu, \nu_{\alpha(\varepsilon)\varepsilon}) &= \sup_{z \in Z} \left| f_\mu(z) - \hat{f}_{\alpha(\varepsilon)\varepsilon}(z) \right| \\ &\leq \sup_{z \in Z} \left| (1 - \alpha(\varepsilon)) f_\mu(z) - \sum_{i=1}^{2n+1} \left(\frac{q_i}{1-q_i} - \alpha(\varepsilon) f_\mu(z_i) \right) g_i^\varepsilon(z) \right| \\ &\leq (1 - \alpha(\varepsilon)) \sup_{z \in Z} f_\mu(z) \\ &\quad + \max_i \left| \frac{q_i}{1-q_i} - f_\mu(z_i) \right| + (1 - \alpha(\varepsilon)) f_\mu(z_i). \end{aligned}$$

If ε and $\alpha(\varepsilon)$ go to zero, the first term and the third term on the right-hand side vanish and only the middle term remains. Therefore,

$$\rho^*(\mathbf{q}, \mu) \leq \max_i \left| \frac{q_i}{1-q_i} - f_\mu(z_i) \right|.$$

Eq. (60) follows because, for ν satisfying (59), we also have

$$\begin{aligned} \rho(\mu, \nu) &= \sup_{z \in Z} |f_\mu(z) - f_\nu(z)| \\ &\geq \max_i |f_\mu(z_i) - f_\nu(z_i)| \\ &= \max_i \left| f_\mu(z_i) - \frac{q_i}{1-q_i} \right|, \end{aligned}$$

and hence

$$\rho^*(\mathbf{q}, \mu) \geq \max_i \left| \frac{q_i}{1-q_i} - f_\mu(z_i) \right|. \quad \blacksquare$$

Lemma A.12. For any $\varepsilon > 0$, there exists a continuous function v_ε from $[0, 1]^{2n+1} \times \mathcal{M}^d(Z)$ to $\mathcal{M}^d(Z)$ such that, for any $(\mathbf{q}, \mu) \in [0, 1]^{2n+1} \times \mathcal{M}^d(Z)$, $\hat{\varphi}(v_\varepsilon(\mathbf{q}, \mu)) = \mathbf{q}$ and

$$\rho(\mu, v_\varepsilon(\mathbf{q}, \mu)) \leq \rho^*(\mathbf{q}, \mu) + \varepsilon, \tag{67}$$

where $\rho^*(\mathbf{q}, \mu)$ is again the infimum of the distance $\rho(\mu, \nu)$ over the set of measures $\nu \in \mathcal{M}^d(Z)$ that satisfy (59).

Proof. Fix $\varepsilon > 0$. For any $\mathbf{q} \in Q$ and any $\mu \in \mathcal{M}^d(Z)$, let

$$\begin{aligned} \psi_\varepsilon(\mathbf{q}, \mu) &:= \{ \nu \in \mathcal{M}^d(Z) \mid \varphi(\nu) \\ &= \mathbf{q} \text{ and } \rho(\mu, \nu) < \rho^*(\mathbf{q}, \mu) + \varepsilon \}, \end{aligned} \tag{68}$$

and let $\bar{\psi}_\varepsilon(\mathbf{q}, \mu)$ be the closure of $\psi_\varepsilon(\mathbf{q}, \mu)$. To prove the lemma, it suffices to show that the correspondence $\bar{\psi}_\varepsilon$ from $Q \times \mathcal{M}^d(Z)$ into $\mathcal{M}^d(Z)$ has a continuous selection.

By Theorem 1.2 of Michael (1964), existence of a continuous selection of the closed-valued correspondence $\bar{\psi}_\varepsilon$ from $Q \times \mathcal{M}^d(Z)$ into $\mathcal{M}^d(Z)$ is guaranteed if $\bar{\psi}_\varepsilon$ is convex-valued and lower hemi-continuous. To show that $\bar{\psi}_\varepsilon$ has these properties, it suffices that to show that ψ_ε has them: Convex-valuedness of $\bar{\psi}_\varepsilon$ then follows from the convex-valuedness of ψ_ε and the observation that the closure of a convex set is convex; similarly, lower hemi-continuity of $\bar{\psi}_\varepsilon$ follows from the lower hemi-continuity of ψ_ε .³³

To establish convex-valuedness of ψ_ε , fix $(\mathbf{q}, \mu) \in Q \times \mathcal{M}^d(Z)$. Let ν_1, ν_2 be any two elements of $\psi_\varepsilon(\mathbf{q}, \mu)$, and, for some $\lambda \in (0, 1)$, consider the convex combination $\lambda\nu_1 + (1-\lambda)\nu_2$ of ν_1 and ν_2 . Because ν_1 and ν_2 belong to $\psi_\varepsilon(\mathbf{q}, \mu)$, the densities $f_{\nu_1} f_{\nu_2}$ of ν_1, ν_2 satisfy

$$f_{\nu_1}(z_i) = f_{\nu_2}(z_i) = \frac{q_i}{1-q_i} \tag{69}$$

for all i and

$$\sup_{z \in Z} \rho(\mu, \nu_j) < \rho^*(\mathbf{q}, \mu) + \varepsilon \tag{70}$$

for $j = 1, 2$. From (69), we immediately obtain

$$\begin{aligned} f_{\lambda\nu_1 + (1-\lambda)\nu_2}(z_i) &= \lambda f_{\nu_1}(z_i) + (1-\lambda) f_{\nu_2}(z_i) \\ &= \frac{q_i}{1-q_i} \end{aligned}$$

for all i , hence, $\hat{\varphi}(\lambda\nu_1 + (1-\lambda)\nu_2) = \mathbf{q}$. Moreover, (70) implies

$$\begin{aligned} \rho(\mu, \lambda\nu_1 + (1-\lambda)\nu_2) &= \sup_{z \in Z} |f_\mu(z) - \lambda f_{\nu_1}(z) - (1-\lambda) f_{\nu_2}(z)| \\ &\leq \lambda \sup_{z \in Z} |f_\mu(z) - f_{\nu_1}(z)| + (1-\lambda) \sup_{z \in Z} |f_\mu(z) - \lambda f_{\nu_2}(z)| \\ &< \rho^*(\mathbf{q}, \mu) + \varepsilon. \end{aligned}$$

Thus, $\lambda\nu_1 + (1-\lambda)\nu_2 \in \psi_\varepsilon(\mathbf{q}, \mu)$.

³³ Hildenbrand (1974, pp. 36, 26).

To prove lower hemi-continuity of ψ_ε , consider any pair $(\mathbf{q}, \mu) \in Q \times \mathcal{M}^d(Z)$, any measure $\nu \in \psi_\varepsilon(\mathbf{q}, \mu)$, and any sequence $\{(\mathbf{q}^r, \mu^r)\}$ that converges to (\mathbf{q}, μ) . Because $\hat{\phi}$ is open, there exists a sequence $\{\nu^r\}$ converging to ν such that $\varphi(\nu^r) = \mathbf{q}^r$ for all r . For any r , the triangle inequality implies

$$\rho(\mu^r, \nu^r) \leq \rho(\mu^r, \mu) + \rho(\mu, \nu) + \rho(\nu, \nu^r).$$

Since $\nu \in \psi_\varepsilon(\mathbf{q}, \mu)$ satisfies (70), it follows that, for some $\eta > 0$, we have

$$\rho(\mu^r, \nu^r) \leq \rho(\mu^r, \mu) + \rho^*(\mathbf{q}, \mu) + \varepsilon + \rho(\nu, \nu^r) - \eta$$

for all r . If r is large enough so that $\rho(\mu^r, \mu)$, $\rho(\nu, \nu^r)$, and $|\rho^*(\mathbf{q}, \mu) - \rho^*(\mathbf{q}^r, \mu^r)|$ are all less than $\frac{\eta}{3}$, it follows that

$$\rho(\mu^r, \nu^r) < \rho^*(\mathbf{q}^r, \mu^r) + \varepsilon$$

and, hence, that $\nu^r \in \psi_\varepsilon(\mathbf{q}^r, \mu^r)$. This proves that ψ_ε is lower hemi-continuous. ■

Lemma A.13. *The mapping $\hat{\phi} : C(X, \mathcal{M}^d(Z)) \rightarrow C(X, [0, 1]^{2n+1})$ is open.*

Proof. The lemma is equivalent to the statement that $\hat{\phi}^{-1}$ is a lower hemi-continuous correspondence. To verify lower hemi-continuity, consider any functions $b \in C(X, \mathcal{M}^d(Z))$ and $\beta \in C(X, Q)$ such that $\beta = \hat{\phi}(b)$. Consider any sequence $\{\beta^r\}$ in $C(X, Q)$ that converges to β . For any r , define a function b^r by setting

$$b^r(x) = \nu_{1/r}(\beta^r(x), b(x))$$

for any $x \in X$, where $\nu_{1/r}$ is the function given by Lemma A.12 with $\varepsilon = \frac{1}{r}$. Because $\nu_{1/r}$, β^r , and b , are continuous, $b^r \in C(X, \mathcal{M}^d(Z))$. Lemma A.12 implies that, for any $\gamma > 0$, there exists $\eta(\gamma) > 0$ such that, if $\|\hat{\phi}(b(x)) - \beta^r(x)\| < \eta(\gamma)$ for all x , then

$$\rho(b(x), b^r(x)) \leq \rho^*(\beta^r(x), b(x)) + \frac{1}{r}$$

for all x . By Lemma A.11, it follows that

$$\begin{aligned} \rho(b(x), b^r(x)) &\leq \max_i \left| \frac{\beta_i^r(x)}{1 - \beta_i^r(x)} - f_{b(x)}(z_i) \right| + \frac{1}{r} \\ &= \max_i \left| \frac{\beta_i^r(x)}{1 - \beta_i^r(x)} - \frac{\beta_i(x)}{1 - \beta_i(x)} \right| + \frac{1}{r} \end{aligned}$$

for all x , where $f_{b(x)}$ is the density of the measure $b(x)$. Now the uniform convergence of the sequence $\{\beta^r\}$ to β implies that, for any $\eta > 0$, there exists $R(\eta)$ such that for all $r \geq R(\eta)$, $\rho(b(x), b^r(x)) < \eta$ for all x . This proves that the sequence $\{b^r\}$ converges uniformly to b . Lower hemi-continuity of $\hat{\phi}$ follows immediately. ■

A.3. Proof of Lemma 5.9

Proof. We first show that, if $\mathcal{M}_+^d(T)$ is endowed with the topology that is induced by the uniform topology for density functions, then the maps $\nu \rightarrow \psi_i(\nu) = (\bar{\nu}_i(\nu), b_i(\cdot, \nu))$ are continuous and open. We begin by proving continuity. If $\{\nu^k\}$ is a sequence in $\mathcal{M}_+^d(T)$ that converges to a limit $\nu \in \mathcal{M}_+^d(T)$, the associated densities satisfy

$$\lim_{k \rightarrow \infty} \int_{T-i} f^{\nu^k}(t_i, t_{-i}) = \int_{T-i} f^\nu(t_i, t_{-i}), \tag{71}$$

uniformly on T . Because the continuous function f^ν is bounded on the compact set T , it follows that the densities f^{ν^k} are uniformly bounded. For any i , therefore, Lebesgue's bounded convergence theorem implies that

$$\begin{aligned} \lim_{k \rightarrow \infty} \bar{f}_i^{\nu^k}(t_i) &= \lim_{k \rightarrow \infty} \int_{T-i} f^{\nu^k}(t_i, t_{-i}) dt_{-i} \\ &= \int_{T-i} f^\nu(t_i, t_{-i}) dt_{-i} = \bar{f}_i^\nu(t_i), \end{aligned} \tag{72}$$

uniformly on T_i , which proves that the sequence $\{\bar{\nu}_i(\nu^k)\}$ converges to $\bar{\nu}_i(\nu)$.

Because $\nu \in \mathcal{M}_+^d(T)$, we have $\bar{f}_i^\nu(t_i) > 0$ for all $t_i \in T_i$; indeed, because $\bar{f}_i^\nu(\cdot)$ is continuous, $\bar{f}_i^\nu(t_i)$ is bounded away from zero on T_i . Because the marginal densities $f_i^{\nu^k}$ converge uniformly to \bar{f}_i^ν , it follows that they are uniformly bounded away from zero. If we combine (71) and (72) with (36) in the text, we obtain

$$\lim_{k \rightarrow \infty} \beta_i(t_{-i}|t_i, \nu^k) = \beta_i(t_{-i}|t_i, \nu), \tag{73}$$

uniformly on T . For any $t_i \in T_i$, therefore, the sequence $\{b_i(\cdot|t_i, \nu^k)\}$ converges to $b_i(\cdot|t_i, \nu)$, uniformly on T_i . Continuity of the map $\psi_i(\nu) = (\bar{\nu}_i(\nu), b_i(\cdot, \nu))$ is thus proved.

To show that ψ_i is open, we note that, for any ψ_i has an inverse. For any $\bar{\nu}_i \in \mathcal{M}_+^d(T_i)$ and any $b_i \in C(T_i, \mathcal{M}^d(T_{-i}))$, let $\bar{f}_i^{\bar{\nu}_i}$ and $\beta_i(\cdot|\cdot)$ be the associated marginal and conditional density functions and consider the measure $\psi_i^{-1}(\bar{\nu}_i, b_i)$ that is defined by the formula

$$\psi_i^{-1}(B_i \times B_{-i}|\bar{\nu}_i, b_i) = \int_{B_i} \int_{B_{-i}} \beta_i(t_{-i}|t_i) \bar{f}_i^{\bar{\nu}_i}(t_i) dt_{-i} dt_i. \tag{74}$$

By another application of Lebesgue's bounded convergence theorem, one sees that $\psi_i^{-1}(\bar{\nu}_i, b_i) \in \mathcal{M}_+^d(T)$ and that ψ_i^{-1} is continuous. Therefore ψ_i is open.

Next, we show that the topology that is induced by uniform convergence of densities is in fact the coarsest topology under which the maps $\nu \rightarrow \psi_i(\nu) = (\bar{\nu}_i(\nu), b_i(\cdot, \nu))$. Let B be any open subset of $\mathcal{M}_+^d(T)$ in the topology that is induced by the uniform topology for density functions and let \mathcal{T} be any other topology on $\mathcal{M}_+^d(T)$ such that the maps ψ_i are continuous. Because, in the topology that is induced by the uniform topology for density functions, ψ_i is an open mapping, $\psi_i(B)$ is an open subset of $\mathcal{M}_+^d(T_i) \times C(T_i, \mathcal{M}^d(T_{-i}))$. Because ψ_i is continuous when $\mathcal{M}_+^d(T)$ has the topology \mathcal{T} , it follows that $\psi_i^{-1}(\psi_i(B)) = B \in \mathcal{T}$. Thus, the topology on $\mathcal{M}_+^d(T)$ that is induced by the uniform topology for density functions is at least as coarse as \mathcal{T} . ■

References

Allen, B., 1981. Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information. *Econometrica* 49, 1173–1199.
 Aumann, R., 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
 Aumann, R., 1998. Common priors: a reply to Gul. *Econometrica* 66, 929–938.
 Barelli, P., 2009. On the genericity of full surplus extraction in mechanism design. *J. Econom. Theory* 144, 1320–1333.
 Bergemann, D., Morris, S., 2005. Robust mechanism design. *Econometrica* 73, 1521–1534.
 Börgers, T., Oh, T., 2012. Common Prior Type Spaces in Which Payoff Types and Belief Types are Independent. Working paper, University of Michigan.
 Chen, Y., Xiong, S., 2011. The genericity of beliefs-determine-preferences models revisited. *J. Econom. Theory* 146, 751–761.
 Chen, Y., Xiong, S., 2013. Genericity and robustness of full surplus extraction. *Econometrica* 81, 825–847.
 Clausing, A., 1978. Retractions and open mappings between convex sets. *Math. Z.* 160, 263–274.
 Compte, O., Jehiel, P., 2009. Veto constraint in mechanism design: inefficiency with correlated types. *Amer. Econom. J. Microeconom.* 1, 182–206.
 Crémer, J., McLean, R., 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56, 1247–1257.
 Dekel, E., Fudenberg, D., Morris, S., 2006. Topologies on types. *Theoret. Econom.* 1, 275–309.
 Gizatulina, A., Hellwig, M., 2010. Informational smallness and the scope for limiting informational rents. *J. Econom. Theory* 145, 2260–2281.

- Gizatulina, A., Hellwig, M., 2012. On the Robustness of the BDP Property for Families of Incomplete-Information Models, Working paper, Max Planck Institute for Research on Collective Goods.
- Gul, F., 1998. A comment on Aumann's Bayesian view. *Econometrica* 66, 923–927.
- Heifetz, A., Neeman, Z., 2006. On the generic (im)possibility of full surplus extraction in mechanism design. *Econometrica* 117, 213–233.
- Hellwig, M., 1996. Sequential decisions under uncertainty and the maximum theorem. *J. Math. Econom.* 25, 443–464.
- Hildenbrand, W., 1974. *Core and Equilibria of a Large Economy*. Princeton University Press, Princeton, NJ.
- Hirsch, M., 1994. *Differential Topology*. Springer-Verlag, New York.
- Hurewicz, W., Wallman, H., 1941. *Dimension Theory*. Princeton Mathematical Series, vol. 4., Princeton, NJ.
- Jordan, J., 1977. The continuity of optimal dynamic decision rules. *Econometrica* 45, 1365–1376.
- Kajii, A., Morris, S., 1994. Payoff Continuity in Incomplete Information Games and Almost Uniform Convergence of Beliefs, Caress working paper no. 94-17, University of Pennsylvania.
- Lee, J., 2000. *Introduction to Topological Manifolds*. Springer, NY.
- McAfee, P., Reny, P., 1992. Correlated information and mechanism design. *Econometrica* 60, 395–421.
- Mertens, J.-F., Sorin, S., Zamir, S., 1994. Repeated Games: Part A: Background Material, Working paper, CORE Université de Louvain.
- Mertens, J., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* 10, 619–632.
- Michael, E., 1964. Three mapping theorems. *Proc. Amer. Math. Soc.* 15, 410–415.
- Neeman, Z., 2004. The relevance of private information in mechanism design. *J. Econom. Theory* 117, 55–77.
- Parthasarathy, K., 1967. *Probability Measures on Metric Spaces*. Academic Press, New York and London.