

Online Appendix: The Context of the Game

Amanda Friedenberg*

Martin Meier†

August 2008

Appendix C Proofs for Section 2

Lemma C1 *Fix strategies s^a, s^b and also a type $t^a \in T^a$. Define $f^a : \Theta \times T^b \rightarrow \mathbb{R}$ with $f^a(\theta, t^b) = \pi^a(\theta, s^a(t^a), s^b(t^b))$ for each $(\theta, t^b) \in \Theta \times T^b$. Then f^a is measurable.*

Proof. Define the map $\vec{s}^b : \Theta \times \{s^a(t^a)\} \times T^b \rightarrow \Theta \times \Delta(C^a) \times \Delta(C^a)$ so that each $\vec{s}^b(\theta, s^a(t^a), t^b) = (\theta, s^a(t^a), s^b(t^b))$. Using the fact that s^b is measurable, the map \vec{s}^b is also measurable. Now, define a map $g^a : \Theta \times T^b \rightarrow \Theta \times \{s^a(t^a)\} \times T^b$ so that $g^a(\theta, t^b) = (\theta, s^a(t^a), t^b)$. Of course, g^a is measurable. Finally, note that $f^a = \pi^a \circ \vec{s}^b \circ g^a$ and so is measurable. ■

Appendix D Proofs for Section 3

Throughout this appendix, we continue to assume that the type structures are non-redundant. Again, a formal definition will not be necessary. Instead, we point to a consequence of this, which can be found in Friedenberg-Meier [5, 2008]. For this result, recall: Say a function $f : \Omega \rightarrow \Phi$ is bimeasurable if it is measurable and, for each Borel set E in Ω , $f(E)$ is Borel in Φ .

Lemma D1 (Friedenberg-Meier [5, 2008]) *If (h^a, h^b) is a type morphism from Λ to Λ_* , then h^a and h^b are injective, bimeasurable and uniquely defined.*

Properties 3.1-3.2 are immediate from Lemma D1. For Property 3.3, we use another consequence of non-redundant structures.

Lemma D2 (Friedenberg-Meier [5, 2008]) *Suppose (h^a, h^b) is a type morphism from Λ to Λ_* and (h_*^a, h_*^b) is a type morphism from Λ_* to Λ . Then, $h_*^a = (h^a)^{-1}$ and $h_*^b = (h^b)^{-1}$.*

*Department of Economics, W.P. Carey School of Business, Arizona State University, P.O. Box 873806, Tempe, AZ, 85287-3806, amanda.friedenberg@asu.edu.

†Instituto de Análisis Económico - CSIC, Barcelona, martin.meier@iae.csic.es

Proof of Property 3.3. Suppose (h^a, h^b) is a type morphism from Λ to Λ_* . Apply Lemma D1 to get that h^a and h^b are measurable embeddings. Suppose there is also a type morphism, viz. (h_*^a, h_*^b) , from Λ_* to Λ . By Lemma D2, $h_*^a = (h^a)^{-1}$ and $h_*^b = (h^b)^{-1}$. It follows that h^a and h^b are measurable isomorphisms of type spaces. ■

Proof of Lemma 3.1. Suppose, contra hypothesis, $h^a(T^a) = T_*^a$ and $h^b(T^b) = T_*^b$. Define $h_*^a : T_*^a \rightarrow T^a$ (resp. $h_*^b : T_*^b \rightarrow T^b$) so that $h_*^a = (h^a)^{-1}$ (resp. $h_*^b = (h^b)^{-1}$). Then, h_*^a (resp. h_*^b) is bijective. Fix some $t_*^a \in T_*^a$ and note that $h^a(h_*^a(t_*^a)) = t_*^a$. Also note that, for any event $E \subseteq \Theta \times T^b$, $((\text{id} \times h^b)(E)) = ((\text{id} \times h_*^b)^{-1}(E))$. These facts will be used below.

We will show that Λ_* can be embedded into Λ via (h_*^a, h_*^b) . For this, fix an event $E \subseteq \Theta \times T^b$. Note that

$$\begin{aligned} \lambda^a(h_*^a(t_*^a))(E) &= \lambda_*^a(h^a(h_*^a(t_*^a)))((\text{id} \times h^b)(E)) \\ &= \lambda_*^a(t_*^a)((\text{id} \times h^b)(E)) \\ &= \lambda_*^a(t_*^a)((\text{id} \times h_*^b)^{-1}(E)), \end{aligned}$$

where the first line follows from Lemma A2 and the fact that (h^a, h^b) is a type morphism. An analogous argument, replacing a and b , gives that (h_*^a, h_*^b) is a type morphism. This contradicts that Λ can be properly embedded into Λ_* . ■

Proof of Lemma 3.2. Immediate from Lemma A2. ■

Appendix E Proofs for Section 4

Lemma E1 Take (Γ_0, Λ_0) and (Γ_1, Λ_2) as in either Section 4.1 or 4.2.

(i) Fix a strategy for the Bayesian Game (Γ_1, Λ_2) , viz. s_2^a , where $s_2^a(t^a)(C_0^a) = 1$ for all $t^a \in T_0^a$. Define a map $s_0^a : T_0^a \rightarrow \Delta(C_0^a)$ so that, for all $t^a \in T_0^a$ and all events E in C_0^a , $s_0^a(t^a)(E) = s_2^a(t^a)(E)$. Then, s_0^a is measurable.

(ii) Fix a strategy for the Bayesian Game (Γ_0, Λ_0) , viz. $s_0^a : T_0^a \rightarrow \Delta(C_0^a)$. Consider any map $s_2^a : T_2^a \rightarrow \Delta(C_1^a)$ so that, for all $t^a \in T_0^a$ and all events E in C_1^a , $s_2^a(t^a)(E) = s_0^a(t^a)(E \cap C_0^a)$. Then, s_2^a is measurable.

Proof. Both parts use the following fact: There exists a homeomorphism

$$f^a : \Delta(C_0^a) \rightarrow \{\mu \in \Delta(C_1^a) : \mu(C_0^a) = 1\}.$$

(See Kechris [6, 1995; Exercise 17.28].)

Begin with Part (i). Fix an event F_0 in $\Delta(C_0^a)$. Note, $f^a(F_0)$ is an event in $\Delta(C_1^a)$. So, $(s_2^a)^{-1}(f^a(F_0))$ is an event in T_2^a . Note:

$$(s_0^a)^{-1}(F_0) = (s_2^a)^{-1}(f^a(F_0)) \cap T_0^a.$$

Since T_0^a is endowed with the induced topology, $(s_0^a)^{-1}(F_0)$ is an event in T_0^a .

Turn to Part (ii). Fix an event F_1 in $\Delta(C_1^a)$. Let $G_1 = F_1 \cap \{\mu \in \Delta(C_1^a) : \mu(C_0^a) = 1\}$ and notice that G_1 is measurable in $\Delta(C_1^a)$. Since

$$(s_2^a)^{-1}(F_1) = \begin{cases} (s_0^a)^{-1}((f^a)^{-1}(G_1)) \cup \{t_1^a\} & \text{if } s_2^a(t_1^a) \in F_1 \\ (s_0^a)^{-1}((f^a)^{-1}(G_1)) & \text{otherwise,} \end{cases}$$

$(s_2^a)^{-1}(F_1)$ is an event in T_2^a . ■

Lemma E2 *Let Γ_1 and Λ_2 be as described in Section 4.1. There is no equilibrium of (Γ_1, Λ_2) .*

Proof. Suppose, contra hypothesis, there is a Bayesian Equilibrium, viz. (s_2^a, s_2^b) , of the game (Γ_1, Λ_2) . Notice that, for any $\theta \in \Theta_0$, c_1^a is strictly dominated on the θ -section of the game Γ_1 , by any choice c_0^a in C_0^a . As such, for each $t^a \in T_0^a$, $s_2^a(t^a)$ assigns probability 1 to C_0^a . With this, we can define a map $s_0^a : T_0^a \rightarrow \Delta(C_0^a)$ as follows: Fix some $t^a \in T_0^a$. For every event E in C_0^a , let $s_0^a(t^a)(E) = s_2^a(t^a)(E)$. Note, s_0^a is a strategy for the game (Γ_0, Λ_0) . (This uses Lemma E1 and the fact that, for each $t^a \in T_0^a$, $s_2^a(t^a)(C_0^a) = 1$.) Define s_0^b analogously. Then,

$$\Pi_0^a(t^a, s_0^a, s_0^b) = \Pi_2^a(t^a, s_2^a, s_2^b) \quad \text{for all } t^a \in T_0^a.$$

And similarly with a and b interchanged.

Now, fix a strategy for Ann in the game (Γ_0, Λ_0) , viz. $r_0^a : T_0^a \rightarrow \Delta(C_0^a)$. This strategy can be extended to a map $r_2^a : T_2^a \rightarrow \Delta(C_2^a)$. Specifically, for each type $t^a \in T_0^a$ and each event E in C_2^a , $r_2^a(t^a)(E) = r_0^a(t^a)(E \cap C_0^a)$ and let $r_2^a(t_1^a)$ be an arbitrary element of $\Delta(C_2^a)$. Note, by Lemma E1, r_2^a is a strategy for the game (Γ_1, Λ_2) . Under this extension,

$$\Pi_0^a(t^a, r_0^a, s_0^b) = \Pi_2^a(t^a, r_2^a, s_2^b) \quad \text{for all } t^a \in T_0^a.$$

And similarly with a and b interchanged.

Using the above if, for each $t^a \in T_0^a$,

$$\Pi_2^a(t^a, s_2^a, s_2^b) \geq \Pi_2^a(t^a, r_2^a, s_2^b) \quad \text{for all } r_2^a \in S_2^a$$

then, for each $t^a \in T_0^a$,

$$\Pi_0^a(t^a, s_0^a, s_0^b) \geq \Pi_0^a(t^a, r_0^a, s_0^b) \quad \text{for all } r_0^a \in S_0^a.$$

And likewise with a and b interchanged. As such, given that (s_2^a, s_2^b) is a Bayesian equilibrium of (Γ_1, Λ_2) , we also have that (s_0^a, s_0^b) is a Bayesian Equilibrium of the game (Γ_0, Λ_0) . But no such equilibrium exists, a contradiction. ■

Appendix F Rationalizability

In this appendix, we show that rationalizability satisfies Extension and Pull-Back properties. This is the case, despite the fact that the set of rationalizable choice-type pairs may be empty.

We focus on correlated rationalizability. We extend the definitions in Dekel-Fudenberg-Morris [3, 2007] and Battigalli-Di Tillio-Grillo-Penta [1, 2008] to arbitrary (i.e., non-finite) games. Let us begin with the definition.

Fix a Θ -based Bayesian game (Γ, Λ) . The m -rationalizable and rationalizable sets will be defined relative to this game.

Definition F1 Set $R_0^a = C^a \times T^a$ and $R_0^b = C^b \times T^b$. For each $m \geq 0$, put $(c^a, t^a) \in R_{m+1}^a$ if there exists a measure $\mu \in \Delta(\Theta \times C^b \times T^b)$ so that

$$(i) \text{ for any } d^a \in C^a, \int_{\Theta \times C^b} \pi^a(\theta, c^a, c^b) d \text{marg}_{\Theta \times C^b} \mu \geq \int_{\Theta \times C^b} \pi^a(\theta, d^a, c^b) d \text{marg}_{\Theta \times C^b} \mu;$$

$$(ii) \mu(\Theta \times R_m^b) = 1; \text{ and}$$

$$(iii) \lambda^a(t^a) = \text{marg}_{\Theta \times T^b} \mu.$$

And likewise with a and b interchanged. Call $R_m^a \times R_m^b$ the set of **m-rationalizable** choice-type pairs.

Remark F1 Suppose, for each m , R_m^a, R_m^b are non-empty and measurable. Then, for each m , $R_{m+1}^a \times R_{m+1}^b \subseteq R_m^a \times R_m^b$.

Definition F2 Call $R^a \times R^b = \bigcap_m [R_m^a \times R_m^b]$ the set of **rationalizable** choice-type pairs.

Battigalli-Di Tillio-Grillo-Penta [1, 2008] provide epistemic conditions for rationalizability when the parameter and choice sets are finite. Definitions F1-F2 appear to be natural extensions to the infinite case. (Of course, this is a far cry from a formal epistemic treatment.) We also note that rationalizability may not be the appropriate solution concept when the question is about “small type structures.” (The appropriate concept may correspond to best-response sets, i.e., as in Pearce [7, 1984].¹)

We begin by pointing out that there are games for which the rationalizable set is empty.

¹The rationale for this statement is analogous to the rationale for self-admissible sets in Brandenburger-Friedenberg-Keisler [2, 2008].

Example F1 (Dufwenberg-Stegeman [4, 2002; Example 2]) Let $\Theta = \{\theta\}$ and consider a Θ -based type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ with $T^a = \{t^a\}$ and $T^b = \{t^b\}$. Define a Θ -based game as follows: Let $C^a = C^b = [0, 1]$. Set

$$\pi^a(\theta, c^a, c^b) = \begin{cases} 1 - c^a & \text{if } 2c^a \geq c^b > 0 \\ c^a & \text{otherwise,} \end{cases}$$

and likewise with a and b interchanged. Then, for each $m \geq 1$, $R_m^a = (0, \frac{1}{2^{m-1}}] \times \{t^a\}$ and $R_m^b = (0, \frac{1}{2^{m-1}}] \times \{t^b\}$. So, $\bigcap_m [R_m^a \times R_m^b] = \emptyset$.

There could also be a second existence problem. To see this, note:

Lemma F1 If $R^a \times R^b \neq \emptyset$, then, for each m , R_m^a and R_m^b are measurable.

Proof. Suppose R_m^b is not measurable. Then, there is no $\mu \in \Delta(\Theta \times C^b \times T^b)$ with $\mu(\Theta \times R_m^b) = 1$. As such, $R_{m+1}^a = \emptyset$ and so $R^a = \emptyset$. ■

Lemma F1 says that, if there is some m so that R_m^a is not measurable, then $R^a \times R^b = \emptyset$. We don't know if, for a general game, the sets R_m^a, R_m^b must be measurable. As such, we don't know if this non-existence problem is possible.

Fix a Θ -based game Γ and Θ -based interactive structures Λ and Λ_* . Write $R_m^a \times R_m^b$ (resp. $R^a \times R^b$) for the set of m -rationalizable (resp. rationalizable) choice-type pairs for the Bayesian game (Γ, Λ) , and write $R_{m,*}^a \times R_{m,*}^b$ (resp. $R_*^a \times R_*^b$) for the set of m -rationalizable (resp. rationalizable) choice-type pairs for the Bayesian game (Γ, Λ_*) .

Definition F3 Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then, the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Rationalizable Extension Property for the Θ -based game Γ** if the following holds: If $(c^a, t^a, c^b, t^b) \in R^a \times R^b$ then $(c^a, h^a(t^a), c^b, h^b(t^b)) \in R_*^a \times R_*^b$. Say the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Rationalizable Extension Property** if it satisfies the Rationalizable Extension Property for each Θ -based game Γ .

Definition F4 Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then, the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Rationalizable Pull-Back Property for the Θ -based game Γ** if the following holds: If $(c^a, h^a(t^a), c^b, h^b(t^b)) \in R_*^a \times R_*^b$, then $(c^a, t^a, c^b, t^b) \in R^a \times R^b$. Say the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Rationalizable Pull-Back Property** if it satisfies the Rationalizable Pull-Back Property for each Θ -based game Γ .

Proposition F1 Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* . Fix, also, a Θ -based game Γ . If, for each m , $R_m^a, R_m^b, R_{m,*}^a, R_{m,*}^b$ are measurable, then $\langle \Lambda, \Lambda_* \rangle$ satisfies the Rationalizable Extension and Pull-Back Properties for Γ .

The following Corollary is an immediate consequence of Proposition F1.

Corollary F1 Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* . If, for each Θ -based game Γ and each m , $R_m^a, R_m^b, R_{m,*}^a, R_{m,*}^b$ are measurable, then $\langle \Lambda, \Lambda_* \rangle$ satisfies the Rationalizable Extension and Pull-Back Properties.

The proof of Proposition F1 will follow from the following Lemma.

Lemma F2 Fix a Θ -based game Γ and Θ -based structures Λ and Λ_* , where Λ can be properly embedded into Λ_* via (h^a, h^b) . Suppose, for each m , $R_m^a, R_m^b, R_{m,*}^a, R_{m,*}^b$ are measurable. Then:

- (i) $(c^a, t^a) \in R_m^a$ implies $(c^a, h^a(t^a)) \in R_{m,*}^a$;
- (ii) $(c^a, t^a) \in [C^a \times T^a] \setminus R_m^a$ implies $(c^a, h^a(t^a)) \in [C^a \times T_*^a] \setminus R_{m,*}^a$.

And, likewise, with a and b reversed.

Proof. The proof is by induction on m . For $m = 0$, the result is immediate. Assume the result holds for m . We will show that it also holds for $m + 1$.

Begin with part (i): Fix $(c^a, t^a) \in R_{m+1}^a$. Then, we can find a measure $\mu \in \Delta(\Theta \times C^b \times T^b)$ satisfying (i)-(iii) of Definition F1. Extend h^b to $\vec{h}^b: \Theta \times C^b \times T^b \rightarrow \Theta \times C^b \times T_*^b$, with $\vec{h}^b(\theta, c^b, t^b) = (\theta, c^b, h^b(t^b))$ for each $(\theta, c^b, t^b) \in \Theta \times C^b \times T^b$. Let μ_* be the image measure of μ under \vec{h}^b . We will show that μ_* satisfies analogs of (i)-(iii), relative to the structure Λ_* , i.e.,

- (i_{*}) for each $d^a \in C^a$, $\int_{\Theta \times C^b} \pi^a(\theta, c^a, c^b) d \text{marg}_{\Theta \times C^b} \mu_* \geq \int_{\Theta \times C^b} \pi^a(\theta, d^a, c^b) d \text{marg}_{\Theta \times C^b} \mu_*$;
- (ii_{*}) $\mu_*(\Theta \times R_{m,*}^b) = 1$; and
- (iii_{*}) $\lambda_*^a(h^a(t^a)) = \text{marg}_{\Theta \times T_*^a} \mu_*$.

This will establish that $(c^a, h^a(t^a)) \in R_{m+1,*}^a$.

Condition (i_{*}) follows from (i) and the fact that $\text{marg}_{\Theta \times C^b} \mu = \text{marg}_{\Theta \times C^b} \mu_*$. For part (ii_{*}), note that

$$\begin{aligned} \mu_*(\Theta \times R_{m,*}^b) &= \mu((\vec{h}^b)^{-1}(\Theta \times R_{m,*}^b)) \\ &= \mu(\Theta \times R_m^b) \\ &= 1, \end{aligned}$$

where the first line uses the fact that $R_{m,*}^b$ is measurable, the second line follows from parts (i)-(ii) of the induction hypothesis, and the third line follows from (ii) of Definition F1. For part (iii_{*}), fix some event $E_* \subseteq \Theta \times T_*^b$ and note that

$$\begin{aligned}
\lambda_*^a(h^a(t^a))(E_*) &= \lambda^a(t^a)((\text{id} \times h^b)^{-1}(E_*)) \\
&= \text{marg}_{\Theta \times T^b} \mu((\text{id} \times h^b)^{-1}(E_*)) \\
&= \mu(C^b \times (\text{id} \times h^b)^{-1}(E_*)) \\
&= \mu((\vec{h}^b)^{-1}(C^b \times E_*)) \\
&= \mu_*(C^b \times E_*) \\
&= \text{marg}_{\Theta \times T_*^b} \mu_*(E_*),
\end{aligned}$$

where the first line follows from the definition of a type morphism, the second line follows from condition (iii), and the fourth and fifth lines follow from construction.

Now we turn to part (ii). Suppose $(c^a, h^a(t^a)) \in R_{m+1,*}^a$. Then, we can find a measure μ_* satisfying conditions (i_*) - (iii_*) above. Let $\vec{g}^b : \Theta \times C^b \times h^b(T^b) \rightarrow \Theta \times C^b \times T^b$ be a map, with $\vec{g}^b(\theta, c^b, h^b(t^b)) = (\theta, c^b, t^b)$. Note that h^b is injective (Lemma D1); as such, \vec{g}^b is well-defined. Also note that h^b is bimeasurable (Lemma D1); as such, \vec{g}^b is the product of measurable maps and so measurable. Using (ii_*) , $\mu_*(\Theta \times C^b \times h^b(T^b)) = 1$. So, the image measure of μ_* under \vec{g}^b is well-defined. Write μ for this measure. We need to show that μ satisfies conditions (i)-(iii) of Definition F1.

Conditions (i)-(ii) are shown by repeating the arguments for (i_*) and (ii_*) above. We focus on condition (iii). Fix some event E in $\Theta \times T^b$ and note that

$$\begin{aligned}
\mu(C^b \times E) &= \mu_*((\vec{g}^b)^{-1}(C^b \times E)) \\
&= \mu_*(C^b \times (\text{id} \times h^b)(E)) \\
&= \text{marg}_{\Theta \times T_*^b} \mu_*((\text{id} \times h^b)(E)) \\
&= \lambda_*^a(h^a(t^a))((\text{id} \times h^b)(E)) \\
&= \lambda^a(t^a)(E),
\end{aligned}$$

where the first line follows from the definition of μ , the fourth line follows from condition (iii_*) , and the last line follows from Lemma A2. This establishes (iii). ■

As a corollary of Lemma F1 and Proposition F1, we have:

Corollary F2 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* . If, for a Θ -based game Γ , $R^a \times R^b$ and $R_*^a \times R_*^b$ are both non-empty, then $\langle \Lambda, \Lambda_* \rangle$ satisfies the Rationalizable Extension and Pull-Back Properties for Γ .*

Corollary F3 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* . If, for each Θ -based game Γ , $R^a \times R^b$ and $R_*^a \times R_*^b$ are both non-empty, then $\langle \Lambda, \Lambda_* \rangle$ satisfies the Rationalizable Extension and Pull-Back Properties.*

References

- [1] Battigalli, P., A. Di Tillio, E. Grillo, A. Penta, “Interactive Epistemology and Solution Concepts for Games with Asymmetric Information,” 2008.
- [2] Brandenburger, A., A. Friedenberg, and H.J. Keisler, “Admissibility in Games,” *Econometrica*, 76, 2008, 307-352.
- [3] Dekel, E., D. Fudenberg, and S. Morris, “Interim Correlated Rationalizability,” *Theoretical Economics*, 2, 2007, 15-40.
- [4] Dufwenberg, M., and M. Stegeman, “Existence and Uniqueness of Maximal Reductions Under Iterated Strict Dominance,” *Econometrica*, 70, 2002, 2007-2023.
- [5] Friedenberg, A. and M. Meier, “A Note on the Relationship Between Type and Hierarchy Morphisms,” 2008, available at <http://www.olin.wustl.edu/faculty/Friedenberg/papers.html>.
- [6] Kechris, A., *Classical Descriptive Set Theory*, Springer-Verlag, 1995.
- [7] Pearce, D., “Rational Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52, 1984, 1029-1050.

For me context is the key—from that comes the understanding of everything.

— Kenneth Noland [?]

1 Introduction

This paper is concerned with the analysis of incomplete information games. For these games, the analyst must specify the players' choices, payoff functions, and hierarchies of beliefs (about the payoffs of the game). The importance of correctly specifying players' actual payoff functions and/or hierarchies of beliefs is well understood. (See, for instance, Kreps-Wilson [?, 1982], Milgrom-Roberts [?, 1982], Geanakoplos-Polemarchakis [?, 1982], Monderer-Samet [?, 1989], Rubinstein [?, 1989], Fudenberg-Tirole [?, 1991], Carlsson-van Damme [?, 1993], Aumann-Brandenburger [?, 1995], Kajii-Morris [?, 1997], Oyama-Tercieux [?, 2005], and Weinstein-Yildiz [?, 2007], among many others.) Here, we argue that it is also important to correctly specify the “context” within which the given hierarchies lie.

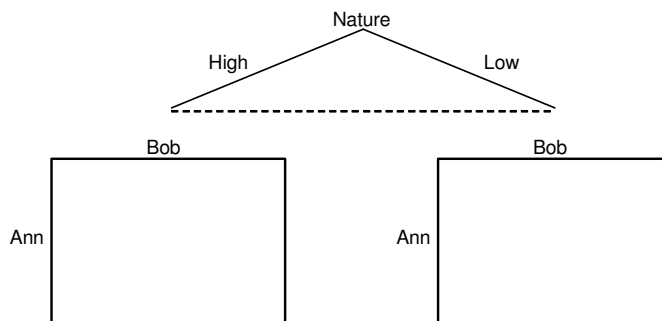


Figure 1.1

To understand this idea, let us take an example. Refer to Figure 1.1. Nature tosses a coin, whose realization is either High or Low. (This can, for instance, reflect a buyer having a High or Low valuation.) The realization of this toss results in distinct matrices (or payoff functions). Each of two players, resp. Ann (a) and Bob (b), faces uncertainty about the realization of this coin toss.

What choices should Ann and Bob make here? Presumably, Ann's choice will depend on her belief about the realization of the coin toss—after all, the realization influences which matrix is being played. But, presumably, Ann's choice will also depend on what she thinks about Bob's belief about the realization of the coin toss. After all, Bob's belief (about the realization of the coin toss) should influence his action, too. And, Ann is concerned not only with what matrix is being played, but also with what choice Bob is making within the matrix.

To analyze the situation, we must amend the description of the game to reflect these hierarchies of beliefs. In particular, we append to the game a type structure. One such type structure is given in Figure 1.2. Here, there are two possible types of Ann, viz. t^a and u^a , and one possible type of Bob, viz. t^b . Type t^a (resp. u^a) of Ann assigns probability one to Nature choosing High (resp.

Low) and Bob’s type being t^b . Type t^b of Bob assigns probability $\frac{1}{2}$ to “Nature choosing High and Ann being type t^a ” and probability $\frac{1}{2}$ to “Nature choosing Low and Ann being type u^a .” So, type t^b of Bob assigns probability $\frac{1}{2}$ to “Nature choosing High and Ann assigning probability one to High” and probability $\frac{1}{2}$ to “Nature choosing Low and Ann assigning probability one to Low.” And so on.




Figure 1.2

For a given type structure, as in Figure 1.2, we can analyze the game associated with Figure 1.1. We defer an analysis for now. Instead, we point to a particular feature of the type structure in Figure 1.2. Here, there are only two possible hierarchies of beliefs that Ann can hold and only one possible hierarchy of beliefs that Bob can hold. In particular, the type structure does not contain all hierarchies of beliefs.

What is the rationale for limiting the type structure in this way? The specified game is only one part of the picture—a small piece of a larger story. The game sits within a broader strategic situation. That is, there is a history to the game, and this history influences the players. As Brandenburger-Friedenberg-Keisler [2, 2008, p. 319] put it:

We think of a particular . . . structure as giving the “context” in which the game is played. In line with Savage’s Small-Worlds idea in decision theory [?, 1954, pp. 82-91], who the players are in the given game can be seen as a shorthand for their experiences before the game. The players’ possible characteristics—including their possible types—then reflect the prior history or context.

Under this view, the type structure, taken as a whole, reflects the context of the game. (Section 7b expands on this point, and discusses the relationship to other views of game theory.)

Here, we are concerned with the case where the players understand more than the analyst, in a particular sense. We imagine the following scenario: The analyst looks at the strategic situation and the history. Perhaps, even, the analyst deduces that certain hierarchies are inconsistent with the history. But, to the players, it is transparent that other—that is, even more—hierarchies are inconsistent with the history. Put differently, players rule out hierarchies the analyst hasn’t ruled out.

$\lambda^a(t^a)$	<i>High</i>	<i>Low</i>
t^b	1	0
u^b		

$\lambda^b(t^b)$	<i>High</i>	<i>Low</i>
t^a	$\frac{1}{2}$	0
u^a	0	$\frac{1}{2}$

$\lambda^a(u^a)$	<i>High</i>	<i>Low</i>
t^b	0	1
u^b		

$\lambda^b(u^b)$	<i>High</i>	<i>Low</i>
t^a		
u^a		

Figure 1.3

Return to the earlier example. Consider the case in which the players' type structure is as given in Figure 1.2. Suppose the analyst misspecifies the type structure, and instead studies the structure in Figure 1.3. Here, there is one extra type of Bob, viz. u^b . Type u^b is associated with some belief, distinct from type t^b 's belief. The particular belief is immaterial. What is important is that each of Ann's types assigns zero probability to this type of Bob. More to the point, each of Ann's types is associated with the exact same beliefs as in the players' type structure. So, the players' type structure can be viewed as a subset (or substructure) of the analyst's type structure.

How does this affect an analysis? Take the solution concept of Bayesian Equilibrium, applied to the game in Figure 1.1 and the type structure in Figure 1.3. For a given Bayesian Equilibrium, the analyst will have a prediction associated with the type u^b —i.e., a type that the players have ruled out. But the analyst will also have a prediction for the types t^a , u^a , and t^b . These are types in the players' structure, namely Figure 1.2.

The question is: How does the analyst's predictions for these types relate to the predictions he would have, if he had analyzed the game using the players' type structure? Presumably, the analyst's predictions shouldn't change. After all, the beliefs associated with t^a , u^a , and t^b have not changed at all. So, we can associate any equilibrium of the players' actual type structure with an equilibrium of the analyst's type structure, and vice versa.

Implicit, in the above, is that Bayesian Equilibrium satisfies Extension and Pull-Back Properties. Let us state these properties semi-formally.

Fix a type structure, viz. Λ , associated with type sets T^a and T^b . We will think of Λ as the players' type structure. Now, consider another type structure Λ_* , associated with type sets T_*^a and T_*^b . Suppose there is a map $h^a : T^a \rightarrow T_*^a$ (resp. $h^b : T^b \rightarrow T_*^b$) so that each t^a and $h^a(t^a)$ (resp. t^b

and $h^b(t^b)$) induces the same hierarchies of beliefs. We will think of Λ_* as the analyst’s structure. In our setting, we can then view the players’ type structure Λ as a subset (or a substructure) of the analyst’s structure Λ_* .² (See Lemmata 3.2 and A2.) Now, we can state the Extension and Pull-back Properties.

The Equilibrium Extension Problem. Fix an equilibrium of Λ . Does there exist an equilibrium of Λ_* so that each $h^a(t^a) \in T_*^a$ and each $h^b(t^b) \in T_*^b$ plays the same strategy as do t^a and t^b (under the original equilibrium of Λ)?

The Equilibrium Pull-Back Problem. Fix an equilibrium of Λ_* . Does there exist an equilibrium of Λ so that each $t^a \in T^a$ and each $t^b \in T^b$ plays the same strategy as do $h^a(t^a)$ and $h^b(t^b)$ (under the original equilibrium of Λ_*)?

Return to the question of whether the analyst can study the game in Figure 1.3. The answer is yes, provided the analyst won’t lose any predictions and the analyst won’t introduce any new predictions. The question of losing predictions is the Extension Problem. The question of introducing new predictions is the Pull-Back Problem.

We will see that the answer to the Extension Problem is no. This is surprising, as “types associated with the players’ structure,” viz. $h^a(T^a)$ (resp. $h^b(T^b)$), assign zero probability to “types that are in the analyst’s structure but not associated with the players’ structure,” viz. $T_*^b \setminus h^b(T^b)$ (resp. $T_*^a \setminus h^a(T^a)$). What, then, goes wrong?

The problem arises from the types that are in the analyst’s structure but not in the players’ structure. (Or, more formally, not in the structure induced by the players’ type structure, viz. $h^a(T^a)$ and $h^b(T^b)$.) There are two possible cases, each associated with a distinct problem.

- (i) These types assign zero probability to types in the players’ type structure.
- (ii) Some of these types assign strictly positive probability to types in the players’ type structure.

In the first case, we may have a problem extending any equilibrium associated with the players’ type structure. This will occur if and only if there is no equilibrium associated with the analyst’s type structure. (Note, there may be no equilibrium associated with the analyst’s type structure, despite the fact that there is an equilibrium associated with the players’ type structure.) See Sections 4.1 and 5.

In the second case, we may have a problem extending some equilibrium associated with the players’ type structure, despite the fact that there is an equilibrium associated with the analyst’s type structure. Section 4.2 will expand on this point.

These two problems shed light on what a ‘large’ type structure must look like—at least, if the goal of this ‘large’ structure is to capture all possible predictions of a Bayesian equilibrium analysis.

²Formally, we assume that no two types induce the same hierarchies of beliefs. Section 3 (specifically, Lemma 3.2) discusses what this assumption delivers formally. Section 7a discusses what this assumption delivers conceptually.

One question is how this ‘large’ structure relates to the so-called universal type structure—e.g., of Mertens-Zamir [?, 1985] and Brandenburger-Dekel [?, 1993].³ (Recall, this structure is terminal, in the sense that it contains each possible type structure as a subset or substructure.⁴) The first problem will suggest that the universal structure is too big, relative to this large structure. The second problem will suggest that the universal structure is too small, relative to this large structure. Sections 4.3 and 7d will expand on this last point. In particular, there, we will discuss the implications of the negative results for a Bayesian Equilibrium analysis of games. We will see that the necessary construction is distinct from constructions already suggested in the literature.

Many papers have asked whether type structures can be embedded in larger type structures. (See, for instance, Böge-Eisle [?, 1979], Mertens-Zamir [?, 1985], Heifetz-Samet [?, 1998], and Meier [?, 2006].) But, to the best of our knowledge, no paper has directly addressed the implication for behavior. Indeed, one contribution of this paper is to spell out the Equilibrium Extension and Pull-Back Problems.

The paper proceeds as follows. Section 2 gives notation. The Extension and Pull-Back Properties are formally defined in Section 3. There, we also state the Pull-Back result. Section 4 shows the negative results, and discusses their implications. Then we turn to positive results. By restricting both the type structure and the game, the Extension Property will obtain. The main restriction in Section 5 is on the type structure. The main restriction in Section 6 is on the game. Section 7 concludes by discussing some conceptual and formal aspects of the paper.

2 Bayesian Games

Throughout the paper, we adopt the following conventions. We will endow the product of topological spaces with the product topology, and a subset of a topological space with the induced topology. Given a Polish space Ω , endow Ω with the Borel sigma-algebra. Write $\Delta(\Omega)$ for the set of probability measures on Ω . Endow $\Delta(\Omega)$ with the topology of weak convergence, so that it is again Polish.

Let Θ be a Polish set, to be interpreted as a set of **payoff types** or the **parameter set**. A **Θ -based game** is then some $\Gamma = \langle \Theta; C^a, C^b; \pi^a, \pi^b \rangle$. Here, the players are a (or Ann) and b (or Bob). (The restriction to two-player games is irrelevant.) The sets C^a and C^b are **choice** or **action** sets; they are taken to be Polish. **Payoff functions** are measurable maps, viz. $\pi^a : \Theta \times C^a \times C^b \rightarrow \mathbb{R}$ and $\pi^b : \Theta \times C^a \times C^b \rightarrow \mathbb{R}$, whose ranges are bounded from above and below. Extend π^a, π^b to $\Theta \times \Delta(C^a) \times \Delta(C^b)$ in the usual way. (Note, the extended functions are measurable and bounded.)

To analyze the game, we will need to append to the game a Θ -based interactive type structure.

Definition 2.1 A *Θ -based interactive type structure* is some $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$, where T^a, T^b are Polish sets and λ^a, λ^b are measurable maps with $\lambda^a : T^a \rightarrow \Delta(\Theta \times T^b)$ and $\lambda^b : T^b \rightarrow$

³See, also, Ambruster-Böge [?, 1979] and Heifetz [?, 1993].

⁴Recall from Footnote 2 that we assume no two types induce the same hierarchy of beliefs. Then, this statement follows from Theorem 2.9 in Mertens-Zamir [?, 1985] and Proposition 3 in Battigalli-Siniscalchi [?, 1999].

$\Delta(\Theta \times T^a)$. We call T^a, T^b (*interactive*) *type sets*.

A Θ -based **Bayesian game** consists of a pair (Γ, Λ) , where Γ is a Θ -based game and Λ is a Θ -based interactive type structure. The Bayesian game induces strategies. A strategy for Ann, viz. s^a , is a measurable map from T^a to $\Delta(C^a)$. Let S^a be the set of strategies for Ann. And similarly for Bob.

Fix strategies s^a, s^b and also a type $t^a \in T^a$. Then, $\pi^a(\cdot, s^a(t^a), s^b(\cdot))$ can be viewed as a measurable map from $\Theta \times T^b$ to \mathbb{R} . (See Lemma C1 in the Online Appendix.) Since $\pi^a(\cdot, s^a(t^a), s^b(\cdot)) : \Theta \times T^b \rightarrow \mathbb{R}$ is measurable and bounded, we can extend π^a to a map $\Pi^a : T^a \times S^a \times S^b \rightarrow \mathbb{R}$, so that

$$\Pi^a(t^a, s^a, s^b) = \int_{\Theta \times T^b} \pi^a(\theta, s^a(t^a), s^b(t^b)) d\lambda^a(t^a).$$

The map $\Pi^b : T^b \times S^a \times S^b \rightarrow \mathbb{R}$ is defined analogously. Note, the maps Π^a, Π^b are defined relative to both Γ and Λ .

Definition 2.2 Say (s^a, s^b) is a **Bayesian equilibrium** if, for all $t^a \in T^a$,

$$\Pi^a(t^a, s^a, s^b) \geq \Pi^a(t^a, r^a, s^b) \quad \text{for all } r^a \in S^a$$

and, for all $t^b \in T^b$,

$$\Pi^b(t^b, s^a, s^b) \geq \Pi^b(t^b, s^a, r^b) \quad \text{for all } r^b \in S^b.$$

3 The Extension and Pull-Back Properties

The purpose of this section is to define the Extension and Pull-Back Properties. For the definitions—and indeed throughout the paper—we will restrict attention to particular type structures, namely type structures that are non-redundant. A type structure is **non-redundant** if any two distinct types, viz. t^a and u^a (resp. t^b and u^b), induce distinct hierarchies of beliefs. We won't need to give a formal definition. Instead, we use consequences that follow from this assumption. (To be clear: We always take the definition of a type structure to be a non-redundant structure.)

Fix two Θ -based structures $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ and $\Lambda_* = \langle \Theta; T_*^a, T_*^b; \lambda_*^a, \lambda_*^b \rangle$. We want to capture the idea that there is a **hierarchy morphism** from Λ to Λ_* , i.e., for each type t^a in T^a (resp. t^b in T^b), there is a type t_*^a in T_*^a (resp. t_*^b in T_*^b) that induces the same hierarchy of beliefs. The next definition allows us to capture this idea without explicitly describing hierarchies of beliefs.

Given a measurable map $f : \Omega \rightarrow \Phi$, write $\underline{f} : \Delta(\Omega) \rightarrow \Delta(\Phi)$ where $\underline{f}(\mu)$ is the image measure. Given maps $f_1 : \Omega_1 \rightarrow \Phi_1$ and $f_2 : \Omega_2 \rightarrow \Phi_2$, write $f_1 \times f_2$ for the map from $\Omega_1 \times \Omega_2$ to $\Phi_1 \times \Phi_2$ so that $(f_1 \times f_2)(\omega_1, \omega_2) = (f_1(\omega_1), f_2(\omega_2))$. Let $\text{id} : \Theta \rightarrow \Theta$ be the identity map.

Definition 3.1 (Mertens-Zamir [?, 1985]) Let $h^a : T^a \rightarrow T_*^a$ and $h^b : T^b \rightarrow T_*^b$ be measurable maps, so that $\underline{\text{id}} \times h^b \circ \lambda^a = \lambda_*^a \circ h^a$ and $\underline{\text{id}} \times h^a \circ \lambda^b = \lambda_*^b \circ h^b$. Then (h^a, h^b) is called a **type morphism** (from Λ to Λ_*).

Definition 3.1 can be illustrated in Figure 3.1: A type morphism, viz. (h^a, h^b) , requires that the diagram commutes.




Figure 3.1

Friedenberg-Meier [5, 2008] shows that (h^a, h^b) is a type morphism if and only if h^a and h^b are hierarchy morphisms. (This result uses the fact that the structures are non-redundant—it is not otherwise true.) As a consequence of non-redundancy and this characterization, we have the following properties.⁵

Property 3.1 *If (h^a, h^b) is a type morphism from Λ to Λ_* , then h^a and h^b are injective and uniquely defined.*

A measurable map is said to be bimeasurable if the image of each measurable set is itself measurable.

Property 3.2 *If (h^a, h^b) is a type morphism from Λ to Λ_* , then h^a and h^b are bimeasurable.*

Property 3.3 *If (h^a, h^b) is a type morphism from Λ to Λ_* , then h^a and h^b are measurable embeddings. If, in addition, (h_*^a, h_*^b) is a type morphism from Λ_* to Λ , then the maps h^a and h^b are measurable isomorphisms with $h^a = (h_*^a)^{-1}$ and $h^b = (h_*^b)^{-1}$.*

With Property 3.3 in mind, we give the following definitions.

Definition 3.2 *Say Λ can be **embedded** into Λ_* (via (h^a, h^b)) if there is a type morphism, viz. (h^a, h^b) , from Λ to Λ_* . Say Λ and Λ_* are **isomorphic** if Λ can be embedded into Λ_* and Λ_* can be embedded into Λ . Say Λ can be **properly embedded** into Λ_* , if Λ can be embedded into Λ_* but Λ_* cannot be embedded into Λ .*

We note:

Lemma 3.1 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* via (h^a, h^b) . Then, either $h^a(T^a) \subsetneq T_*^a$, $h^b(T^b) \subsetneq T_*^b$, or both.*

⁵Proofs for this section are straightforward and so relegated to the Online Appendix.

Lemma 3.2 states a consequence of embedding type structures. (See, also, Lemma A2 for a stronger result.)

Lemma 3.2 *Fix Θ -based structures Λ and Λ_* , where Λ can be embedded into Λ_* via (h^a, h^b) . Then, $h^a(T^a) \times h^b(T^b)$ forms a **belief-closed subset** of $T_*^a \times T_*^b$, i.e., for each $h^a(t^a) \in h^a(T^a)$, $\lambda_*^a(h^a(t^a))(\Theta \times h^b(T^b)) = 1$, and likewise with a and b interchanged.*

Lemma 3.2 says that, if Λ can be embedded into Λ_* via (h^a, h^b) , we can view Λ as a belief-closed subset of Λ_* . This belief-closed subset can be viewed as a “type structure” in its own right. We’ll call such a type structure **the structure induced by Λ** . This structure will consist of $\langle \Theta; h^a(T^a), h^b(T^b); \kappa^a, \kappa^b \rangle$. Note, by Property 3.2, $h^a(T^a)$ (resp. $h^b(T^b)$) is a Borel subset of the Polish space T_*^a (resp. T_*^b). The map $\kappa^a : h^a(T^a) \rightarrow \Delta(\Theta \times h^b(T^b))$ (resp. $\kappa^b : h^b(T^b) \rightarrow \Delta(\Theta \times h^a(T^a))$) is defined so that $\kappa^a(t^a)(E) = \lambda_*^a(t^a)(E)$ (resp. $\kappa^b(t^b)(E) = \lambda_*^b(t^b)(E)$) for each event E in $\Theta \times h^b(T^b)$ (resp. $\Theta \times h^a(T^a)$).⁶

Given a Θ -based game Γ , write s^a (resp. s^b) for a strategy of Ann (resp. Bob) in the Bayesian Game (Γ, Λ) , and write s_*^a (resp. s_*^b) for a strategy of Ann (resp. Bob) in the Bayesian Game (Γ, Λ_*) . Now we can state the Equilibrium Extension and Pull-Back Properties.

Definition 3.3 *Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Equilibrium Extension Property for the Θ -based game Γ** if the following holds: If (s^a, s^b) is a Bayesian Equilibrium of (Γ, Λ) , then there exists a Bayesian Equilibrium of (Γ, Λ_*) , viz. (s_*^a, s_*^b) , so that $s^a = s_*^a \circ h^a$ and $s^b = s_*^b \circ h^b$. Say the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Equilibrium Extension Property** if it satisfies the Equilibrium Extension Property for each Θ -based game Γ .*

Definition 3.4 *Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the **Equilibrium Pull-Back Property** if, for each Θ -based game Γ , the following holds: If (s_*^a, s_*^b) is a Bayesian Equilibrium of (Γ, Λ_*) , then $(s_*^a \circ h^a, s_*^b \circ h^b)$ is a Bayesian Equilibrium of (Γ, Λ) .*

Section 4 will show that the Equilibrium Extension Property may fail. Sections 5-6 will provide conditions under which the Equilibrium Extension Property is satisfied, i.e., for a particular game Γ . On the other hand, the Equilibrium Pull-Back Property is always satisfied. This is a consequence of a result in Friedenber-Meier [?, 2008].

Proposition 3.1 (Friedenberg-Meier [?, 2008]) *Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then, the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Pull-Back Property.*

⁶Note, $h^b(T^b)$ is Borel in T_*^b and endowed with the induced topology. So, if E is Borel in $\Theta \times h^b(T^b)$, then it is Borel in $\Theta \times T_*^b$. As such, the map κ^a is well-defined. That said, formally, the structure induced by Λ need not be an interactive structure in the sense of Definition 2.1. The sets $h^a(T^a)$ and $h^b(T^b)$ may not be Polish. This will be immaterial for our analysis.

Taken together, the Pull-Back Property and Property 3.3 have an immediate consequence for the Extension Property.

Corollary thesection.1 *Let Λ and Λ_* be two isomorphic Θ -based interactive type structures. Then, the pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property.*

Note, non-redundancy is an important condition here. For the redundant case, amend Definitions 3.3-3.4 to directly reflect hierarchy morphisms. Then, examples in Ely-Peski [?, 2006] and Dekel-Fudenberg-Morris [3, 2007], show that Proposition 3.1 and Corollary thesection.1 do not follow. Section 7a discusses this further.

In light of Corollary thesection.1, we will focus on the case in which Λ can be properly embedded into Λ_* .

4 The Negative Results

There are two reasons that Equilibrium Extension may fail—or, at least, two reasons that are known to us. (In Section 7f, we point out that there may be another form of an extension failure.) These reasons will be presented by way of two examples. In the example of Section 4.1, we cannot extend any equilibrium. The reason is that, in that case, there is no equilibrium associated with the analyst’s structure. In the example of Section 4.2, we will be able to extend some equilibrium, but not all equilibria. So, there, we will fail the Extension Property, despite the fact that there is an equilibrium of the analyst’s structure.

The first example is quite simple. It is essentially a corollary of non-existence of Bayesian equilibrium. The example is well-understood—it is, so to speak, “in the air.”⁷ We present it only for completeness. The second example is more involved and, arguably, novel. To show non-extension here, we will use non-existence of Bayesian Equilibrium. However, this second example is not an immediate corollary of non-existence. In particular, it makes direct use of the Bayesian equilibrium concept, i.e., it need not hold under any solution concept that fails existence. See Section 7c for a discussion.

We will present these two examples and then discuss their distinct implications. Both examples will use a Bayesian Game—satisfying certain properties—to construct a new Bayesian Game (with a different parameter set). That is, we won’t be interested in the starting game perse. Rather, it will serve as a “germ,” i.e., it will aid in constructing the game of interest. The “germ” is a Θ_0 -based Bayesian Game, viz. (Γ_0, Λ_0) , that has no equilibrium. (See Section 7c for an example. But, note, the particular game used is irrelevant.) Write $\Gamma_0 = \langle \Theta_0; C_0^a, C_0^b; \pi_0^a, \pi_0^b \rangle$ and $\Lambda_0 = \langle \Theta_0; T_0^a, T_0^b; \lambda_0^a, \lambda_0^b \rangle$, so that (Γ_0, Λ_0) has no Bayesian equilibrium.

⁷That said, to the best of our knowledge, no one has asked the question of the Extension Problem, and so no one has formally shown this result.

Recall, we take type structures to be non-redundant. As such, we are careful to choose Λ_0 so that it is non-redundant. We also normalize the game Γ_0 so that the ranges of the payoff functions lie in $[1, 2]$.

4.1 The First Extension Failure

This is an example where “types in the analyst’s structure but not in the players’ structure” assign zero probability to “types in the players’ structure.”

Start with the Θ_0 -based Bayesian Game (Γ_0, Λ_0) . Construct $\Theta_1 = \Theta_0 \cup \{\theta_1\}$ so that $\Theta_0 \subsetneq \Theta_1$. We will use the game Γ_0 to build a Θ_1 -based game $\Gamma_1 = \langle \Theta_1; C_1^a, C_1^b; \pi_1^a, \pi_1^b \rangle$. Let $C_1^a = C_0^a \cup \{c_1^a\}$ and $C_1^b = C_0^b \cup \{c_1^b\}$, where $C_0^a \subsetneq C_1^a$ and $C_0^b \subsetneq C_1^b$. Ann’s payoff function is depicted in Figure 4.1. First, consider the left-hand panel, which depicts the matrix for any given $\theta \in \Theta_0$. (We call this matrix the θ -section of a game.) Here, $\pi_1^a(\theta, c^a, c^b) = \pi_0^a(\theta, c^a, c^b)$, for all $(c^a, c^b) \in C_0^a \times C_0^b$. For all $c^b \in C_1^b$, $\pi_1^a(\theta, c^a, c^b) = 0$. For all $c^a \in C_0^a$, $\pi_1^a(\theta, c^a, c_1^b) = 1$. Next, consider the case of θ_1 , i.e., the right-hand panel of Figure 4.1. Here, $\pi_1^a(\theta_1, c^a, c^b) = 1$ if $(c^a, c^b) = (c_1^a, c_1^b)$, and $\pi_1^a(\theta_1, c^a, c^b) = 0$ otherwise. Bob’s payoffs are defined by reversing a and b .




Figure 4.1

Fix a Θ_1 -based structure, viz. Λ_1 , defined as follows. The type sets are $T_1^a = \{t_1^a\}$ and $T_1^b = \{t_1^b\}$. The measures $\lambda_1^a(t_1^a)$ and $\lambda_1^b(t_1^b)$ are concentrated on (θ_1, t_1^b) and (θ_1, t_1^a) , respectively.

Then (Γ_1, Λ_1) has some equilibrium—in fact, many.




Figure 4.2

Now, consider another Θ_1 -based structure, viz. Λ_2 , depicted in Figure 4.2. The type sets are $T_2^a = T_0^a \cup T_1^a$ and $T_2^b = T_0^b \cup T_1^b$, where we take $t_1^a \notin T_0^a$ and $t_1^b \notin T_0^b$. The left-hand panel depicts the measure $\lambda_2^a(t^a)$, for $t^a \in T_0^a$. It simply agrees with $\lambda_0^a(t^a)$, i.e., for each event E in $\Theta_1 \times T_2^b$, $\lambda_2^a(t^a)(E)$ equals $\lambda_0^a(t^a)(E \cap (\Theta_0 \times T_0^b))$. The right-hand panel depicts the measure $\lambda_2^a(t_1^a)$, which now agrees with $\lambda_1^a(t_1^a)$, i.e., for each event E in $\Theta_1 \times T_2^b$, $\lambda_2^a(t_1^a)(E)$ equals $\lambda_1^a(t_1^a)(E \cap (\Theta_1 \times T_1^b))$. The map λ_2^b is defined analogously.

Note, Λ_2 is non-redundant. To see this, recall that Λ_0 is non-redundant, so that types in T_0^a induce distinct hierarchies of beliefs in Λ_2 . We also have that $\lambda_2^a(t^a)(\Theta_0 \times T_2^b) = 1$, for each $t^a \in T_0^a$, and that $\lambda_2^a(t_1^a)(\{\theta_1\} \times T_2^b) = 1$. As such, types in T_0^a have distinct first-order beliefs from the type t_1^a .

Next, note:

Remark 4.1 *The structure Λ_1 can be embedded into Λ_2 via $(\text{id}^a, \text{id}^b)$, where $\text{id}^a : T_1^a \rightarrow T_2^a$ and $\text{id}^b : T_1^b \rightarrow T_2^b$ are the identity maps.*

That said, there is no equilibrium of (Γ_1, Λ_2) , and so no equilibrium of (Γ_1, Λ_1) can be extended to an equilibrium of (Γ_1, Λ_2) . For the idea, suppose otherwise. Recall, for each $t^a \in T_0^a$, $\lambda_2^a(t^a)$ assigns probability 1 to $\Theta_0 \times T_2^b$, and likewise for types $t^b \in T_0^b$. So, if there is a Bayesian equilibrium, viz. (s_2^a, s_2^b) , for the game (Γ_1, Λ_2) , each type $t^a \in T_0^a$ assigns probability 1 to C_0^a . And, likewise for types $t^b \in T_0^b$. But then, the restrictions of s_2^a and s_2^b to (Γ_0, Λ_0) would have been a Bayesian equilibrium of the original game. (The Online Appendix gives a proof.) As a consequence:

Proposition 4.1 *The pair $\langle \Lambda_1, \Lambda_2 \rangle$ (as defined in this subsection) fails the Equilibrium Extension Property.*

Indeed, we can go further. No equilibrium of (Γ_1, Λ_1) can be extended to an equilibrium of (Γ_1, Λ_2) . That is, for any equilibrium (s_1^a, s_1^b) of (Γ_1, Λ_1) , we cannot find an equilibrium (s_2^a, s_2^b) of (Γ_1, Λ_2) with $s_1^a = s_2^a \circ \text{id}^a$ and $s_1^b = s_2^b \circ \text{id}^b$.

4.2 The Second Extension Failure

This is an example where “types in the analyst’s structure but not in the players’ structure” assign positive probability to “types in the players’ structure.”




Figure 4.3

Again, start with the Θ_0 -based Bayesian Game (Γ_0, Λ_0) . Let Θ_1 be as in Section 4.1. Now, construct a Θ_1 -based game Γ_1 similar to above. Specifically, C_1^a and C_1^b are as in Section 4.1. But now the payoff functions are different. Ann’s payoff function is depicted in Figure 4.3. First, consider the left-hand panel, where $\theta \in \Theta_0$. Here, $\pi_1^a(\theta, c^a, c^b) = \pi_0^a(\theta, c^a, c^b)$, for all $(c^a, c^b) \in C_0^a \times C_0^b$. For all $c^a \in C_1^a$, let $\pi_1^a(\theta, c^a, c_1^b) = 1$. For all $c^b \in C_0^b$, $\pi_1^a(\theta, c_1^a, c^b) = 0$. Next, consider the case of θ_1 . This is the right-hand panel of Figure 4.3. Here, for all $(c^a, c^b) \in C_0^a \times C_0^b$, $\pi_1^a(\theta_1, c^a, c^b) = x$ where $x > 0$. Also, $\pi_1^a(\theta_1, c_1^a, c_1^b) = y$, where $y > 0$. For all other pairs of (c^a, c^b) , $\pi_1^a(\theta_1, c^a, c^b) = 0$. Bob’s payoff function is defined by reversing a and b .

For each θ -section of the game Γ_1 , the choice profile (c_1^a, c_1^b) is a pure strategy Nash equilibrium. As such, for any Θ_1 -based structure Λ , (Γ_1, Λ) has a Bayesian equilibrium, where each type of Ann chooses c_1^a with probability 1, and each type of Bob chooses c_1^b with probability 1.

Nonetheless, we construct Θ_1 -based structures Λ_1 and Λ_2 , where Λ_1 can be embedded into Λ_2 , but some equilibrium of (Γ_1, Λ_1) cannot be extended to an equilibrium of (Γ_1, Λ_2) . (Of course, there will be another equilibrium of (Γ_1, Λ_1) that can be extended to an equilibrium of (Γ_1, Λ_2) , i.e., the one just mentioned above.)

Let Λ_1 be as in Section 4.1. There is an equilibrium in which types t_1^a and t_1^b play c_1^a and c_1^b with probability 1. Yet, there is also an equilibrium, viz. (s_1^a, s_1^b) , in which $s_1^a(t_1^a)$ assigns probability 1 to C_0^a and $s_1^b(t_1^b)$ assigns probability 1 to C_0^b . (In fact, there are many such equilibria since C_0^a and C_0^b must be non-singletons. See Section 6.)




Figure 4.4

For the structure Λ_2 , refer to Figure 4.4. As in Section 4.1, let $T_2^a = T_0^a \cup T_1^a$, where $T_0^a \subsetneq T_2^a$. And likewise for Bob. Fix some $p \in (0, 1)$ and consider a type $t^a \in T_0^a$. (Note, p is chosen to be the same for each $t^a \in T_0^a$.) For this type, define $\lambda_2^a(t^a)$ as follows. Fix an event E in $\Theta_1 \times T_2^b$. If $(\theta_1, t_1^b) \in E$, let $\lambda_2^a(t^a)(E) = p\lambda_0^a(E \cap (\Theta_0 \times T_0^b)) + (1-p)$. If $(\theta_1, t_1^b) \notin E$, let $\lambda_2^a(t^a)(E) = p\lambda_0^a(E \cap (\Theta_0 \times T_0^b))$. It is readily verified that this is indeed a probability measure. Next, consider the type t_1^a and define $\lambda_2^a(t_1^a)$ so that, for each event E in $\Theta_1 \times T_2^b$, $\lambda_2^a(t_1^a)(E) = \lambda_1^a(t_1^a)(E \cap (\Theta_1 \times T_1^b))$. Define λ_2^b analogously.

Note, Λ_2 is non-redundant. To see this, recall that Λ_0 is non-redundant, so that types in T_0^a induce distinct hierarchies of beliefs in Λ_2 . We also have that $\lambda_2^a(t^a)(\Theta_0 \times T_2^b) > 0$, for each $t^a \in T_0^a$, and $\lambda_2^a(t_1^a)(\Theta_0 \times T_2^b) = 0$. As such, types in T_0^a have distinct first-order beliefs from the type t_1^a .

Just as in Section 4.1, we have:

Remark 4.2 *The structure Λ_1 can be embedded into Λ_2 via $(\text{id}^a, \text{id}^b)$, where $\text{id}^a : T_1^a \rightarrow T_2^a$ and $\text{id}^b : T_1^b \rightarrow T_2^b$ are the identity maps.*

Nonetheless, we will show:

Lemma 4.1 *If (s_2^a, s_2^b) is a Bayesian Equilibrium of (Γ_1, Λ_2) , then either $s_2^a(t_1^a)(C_0^a) < 1$ or $s_2^b(t_1^b)(C_0^b) < 1$ (or both).*

Proof. Fix a Bayesian Equilibrium, viz. (s_2^a, s_2^b) . Suppose the result is false, i.e., $s_2^a(t_1^a)(C_0^a) = 1$ and $s_2^b(t_1^b)(C_0^b) = 1$.

Fix a type $t^a \in T_0^a$. For this type, the expected payoffs from choosing some $c_0^a \in C_0^a$ are

$$\mathbb{E}(t^a, c_0^a) = p \int_{\Theta_0 \times T_0^b} \pi_1^a(\theta_0, c_0^a, s_2^b(t_0^b)) d\lambda_0^a(t^a) + (1-p)x.$$

This type's expected payoffs from choosing c_1^a are

$$\mathbb{E}(t^a, c_1^a) = p \int_{\Theta_0 \times T_0^b} \pi_1^a(\theta_0, c_1^a, s_2^b(t_0^b)) d\lambda_0^a(t^a).$$

Also, note that, for each $(\theta_0, t_0^b) \in \Theta_0 \times T_0^b$,

$$\pi_1^a(\theta_0, c_0^a, s_2^b(t_0^b)) \geq \pi_1^a(\theta_0, c_1^a, s_2^b(t_0^b)).$$

(If $s_2^b(t_0^b)(C_0^b) > 0$, then the inequality is strict.) So,

$$\int_{\Theta_0 \times T_0^b} \pi_1^a(\theta_0, c_0^a, s_2^b(t_0^b)) d\lambda_0^a(t^a) \geq \int_{\Theta_0 \times T_0^b} \pi_1^a(\theta_0, c_1^a, s_2^b(t_0^b)) d\lambda_0^a(t^a).$$

Indeed, since $1 > p$ and $x > 0$,

$$\mathbb{E}(t^a, c_0^a) > \mathbb{E}(t^a, c_1^a).$$

This says that, for each $t^a \in T_0^a$, $s_2^a(t^a)(C_0^a) = 1$. An analogous argument gives that, for each $t^b \in T_0^b$, $s_2^b(t^b)(C_0^b) = 1$.

Now, we will construct a map $s_0^a : T_0^a \rightarrow \Delta(C_0^a)$ from the map s_2^a . To do so, we will use the fact that $s_2^a(t^a)(C_0^a) = 1$ for all $t^a \in T_0^a$. Specifically, for each $t^a \in T_0^a$ and each event E in C_0^a , let $s_0^a(t^a)(E) = s_2^a(t^a)(E)$. Note that $s_0^a(t^a)$ defines a probability measure. Moreover, s_0^a is measurable (Lemma E1 in the Online Appendix), and so is a strategy of the Bayesian game (Γ_0, Λ_0) . Define s_0^b analogously. With this,

$$\Pi_2^a(t^a, s_2^a, s_2^b) = p\Pi_0^a(t^a, s_0^a, s_0^b) + (1-p)x \quad \text{for all } t^a \in T_0^a.$$

And similarly with a and b interchanged.

Fix a strategy of Ann for the game (Γ_0, Λ_0) , viz. $r_0^a : T_0^a \rightarrow \Delta(C_0^a)$. This strategy can be extended to a strategy for the game (Γ_1, Λ_2) , viz. $r_2^a : T_2^a \rightarrow \Delta(C_1^a)$. Specifically, for each type $t^a \in T_0^a$ and each event E in C_1^a , set $r_2^a(t^a)(E) = r_0^a(t^a)(E \cap C_0^a)$. Choose $r_2^a(t_1^a)$ to be any element of $\Delta(C_1^a)$ with $r_2^a(t_1^a)(C_0^a) = 1$. Then, r_2^a is measurable (Lemma E1 in the Online Appendix), and so a strategy for the Bayesian game (Γ_1, Λ_2) . Under this extension,

$$\Pi_2^a(t^a, r_2^a, s_2^b) = p\Pi_0^a(t^a, r_0^a, s_0^b) + (1-p)x \quad \text{for all } t^a \in T_0^a.$$

And similarly define strategies r_2^b .

Return to the fact that (s_2^a, s_2^b) is a Bayesian equilibrium for the game (Γ_1, Λ_2) . Then, using

the above, for each $t^a \in T_0^a$ and each $r_0^a \in S_0^a$,

$$\begin{aligned} p\Pi_0^a(t^a, s_0^a, s_0^b) + (1-p)x &= \Pi_2^a(t^a, s_2^a, s_2^b) \\ &\geq \Pi_2^a(t^a, r_2^a, s_0^b) \\ &= p\Pi_0^a(t^a, r_0^a, s_0^b) + (1-p)x, \end{aligned}$$

where r_2^a is defined as above. It follows that, for each $t^a \in T_0^a$,

$$\Pi_0^a(t^a, s_0^a, s_0^b) \geq \Pi_0^a(t^a, r_0^a, s_0^b) \quad \text{for all } r_0^a \in S_0^a.$$

And, likewise, for each $t^b \in T_0^b$,

$$\Pi_0^b(t^b, s_0^a, s_0^b) \geq \Pi_0^b(t^b, s_0^a, r_0^b) \quad \text{for all } r_0^b \in S_0^b.$$

This says that (s_0^a, s_0^b) is a Bayesian equilibrium for the game (Γ_0, Λ_0) , a contradiction. ■

The negative result is an immediate implication of Lemma 4.1.

Theorem 4.1 *The pair $\langle \Lambda_1, \Lambda_2 \rangle$ (as defined in this subsection) fails the Equilibrium Extension Property.*

Proof. Consider a strategy profile (s_1^a, s_1^b) of the game (Γ_1, Λ_1) , with $s_1^a(t_1^a)(C_0^a) = s_1^b(t_1^b)(C_0^b) = 1$. This is a Bayesian equilibrium of that game. Fix a Bayesian equilibrium (s_2^a, s_2^b) of the game (Γ_1, Λ_2) . By Lemma 4.1, either $s_1^a \neq (s_2^a \circ \text{id}^a)$, or $s_1^b \neq (s_2^b \circ \text{id}^b)$, or both. ■

In this case, some equilibrium of (Γ_1, Λ_1) can be extended to an equilibrium of (Γ_1, Λ_2) . But this does not hold for every equilibrium.

4.3 Implications

Can the analyst use a ‘larger’ type structure to analyze the game, and maintain predictions associated with the players’ actual ‘smaller’ type structure? Now we see that the answer is no.

Let’s review why. Refer back to the examples in Sections 4.1-4.2. There, the types that are in the analyst’s structure but not in the players’ structure, viz. $T_0^a \times T_0^b$, impose an equilibrium restriction on the types in the players’ structure, viz. $T_1^a \times T_1^b$. In the first example (Section 4.1), this alone imposes a difficulty for extension. (Specifically, because there is no equilibrium among the types in $T_0^a \times T_0^b$, there can be no equilibrium for the whole structure.) In the second example, this does not necessarily cause a problem for the extension. (There is some equilibrium of the structure Λ_1 that can be extended to an equilibrium of Λ_2 .) But, there, types in the players’ structure also impose an equilibrium restriction on types in the analyst’s structure. For certain equilibria, this causes a conflict.

Of course, if the analyst recognizes that he misspecified the players' type structure (in a particular way), he can simply use the players' type structure to analyze the game. Here, we study the case in which the analyst doesn't recognize this. As the "meta-analyst" or the "super-analyst," we ask: What are the implications of this result? There are two possible routes that we, the super-analyst, can take.

- (i) Change the type structure, but not the analysis.
- (ii) Change the analysis, but not the type structure.

We now discuss each of these two approaches.

(i) This route takes the idea of Bayesian Equilibrium as given—i.e., it doesn't alter the definition or application. Given this, as the super-analyst, we may want to find a large type structure—one in which we can maintain all possible predictions associated with the players' actual type structures.

One possibility, often suggested in the literature, is to use the canonical construction of the universal type structure (e.g., the bottom-up construction of Mertens-Zamir [?, 1985] and/or the top-down construction of Brandenburger-Dekel [?, 1993]). After all, these constructions 'contain' any type structure as a belief-closed subset. But the examples in Sections 4.1-4.2 suggest that this will not do. In particular, there are two problems. The first example suggests that the universal construction may be too large, and the second suggests that the universal construction may be too small.

Begin with the example in Section 4.1. There, the structure Λ_2 is a belief-closed subset of the universal type structure. This suggests that, for a particular game such as Γ_1 , the universal structure may be too large—if the goal is to obtain one large structure that contains all possible predictions. For this purpose, we may want to 'remove' certain types from the large construction. In particular, we will want to do so if those types are associated with a type structure (such as Λ_0) that has no equilibrium relative to the given game (just as (Γ_0, Λ_0) has no equilibrium). But notice that this large structure will be based not only on the parameter set, but also on the particular game being studied. (For a Θ_0 -based game Γ , distinct from Γ_0 , (Γ, Λ_0) may very well have an equilibrium. The types in Λ_0 should not be excluded for the game Γ .) As such, the construction is quite different from the universal construction. Recall, the universal structure depends on the parameter set—not on the game itself.

Now, turn to the example in Section 4.2. What should the large structure look like here? Certainly, the large structure should have a prediction associated with each type in Λ_2 . After all, there is a Bayesian equilibrium associated with the game (Γ_1, Λ_2) , and we would like to understand how each type in Λ_2 would behave under an equilibrium analysis. Note, to maintain the hierarchies of beliefs associated with types in $T_0^a \times T_0^b \subseteq T_2^a \times T_2^b$, in the large structure, some types must assign probability $(1 - p)$ to $T_1^a \times T_1^b \subseteq T_2^a \times T_2^b$.

But, we would also like to maintain each prediction associated with (Γ_1, Λ_1) . It would seem that, for this, the types $T_0^a \times T_0^b \subseteq T_2^a \times T_2^b$ must assign zero probability to the types $T_1^a \times T_1^b$.

Given that our goal is to (i) have a prediction for each type associated with Λ_2 , and (ii) maintain each of the predictions associated with Λ_1 , we must have two copies of the types $T_1^a \times T_1^b$ in our structure: one copy that gets positive probability under $T_0^a \times T_0^b$ and another copy that gets zero probability under $T_0^a \times T_0^b$. Figure 4.5 illustrates this structure.

Note, we begin by looking at predictions associated with non-redundant structures. We, the super-analyst, then ask for a large type structure—i.e., one that can capture all of these predictions. What we see is that the structure may need to contain two belief-closed subsets that are ‘identical’ (and so induce the same hierarchies of beliefs). So, the output is a particular form of a redundant structure.




Figure 4.5

The canonical constructions of universal structures are all non-redundant. So, if our goal is to retain all possible predictions, then the universal structure may be too small.

It is important to note that the redundancies introduced here are distinct from the redundancies already mentioned in the literature. Section 7d discusses this further.

(ii) There is another possible route. Instead of asking for an analysis relative to a particular ‘large’ type structure, we can change the method of analysis. We can define the concept of Bayesian Equilibrium relative to belief-closed subsets of a given type structure. Then, we can analyze a

Bayesian Game (Γ, Λ) relative to each belief-closed subset of Λ . As such, our analysis may not be “completed” when we find the Bayesian Equilibria associated with the given type structure.

5 Restrictions on Type Structures

Now, we turn to exploring conditions under which the Extension Property is satisfied. We begin with conditions on the type structure.

Let’s begin by reviewing the failure of the Equilibrium Extension Property in Section 4.2. There, we had structures Λ_1 and Λ_2 . The structure Λ_1 can be embedded into Λ_2 via the identity maps. So, referring to Figure 5.1, the structure Λ_1 can be viewed as a belief-closed subset of the structure Λ_2 . Types in this subset impose an equilibrium restriction on types outside of this subset—i.e., on types in $T_0^a = T_2^a \setminus T_1^a$ and $T_0^b = T_2^b \setminus T_1^b$. This is because these latter types assign positive probability to types associated with Λ_1 . This problem would not arise if types in T_0^a (resp. T_0^b) assigned probability 1 to types in T_0^b (resp. T_0^a).




Figure 5.1

Suppose we instead have a type structure, viz. Λ_* , that can be viewed as the union of two type structures. For a given game, can we extend an equilibrium associated with one of these structures to an equilibrium associated with Λ_* ? The answer will be yes if and only if there exists an equilibrium associated with the other structure.

We formalize this idea and then state a general result. But is it interesting to study the case where Λ_* can be viewed as the union of two type structures? We go on to show that this arises naturally in one special case—where the analyst’s structure satisfies a common prior assumption.

5.1 Decomposing Type Structures

We now formalize the idea that a type structure Λ_* can be viewed as the union of some structure Λ and some ‘remaining structure,’ which we’ll call the difference structure.

Definition 5.1 Fix Θ -based structures Λ and Λ_* , so that Λ can be embedded into Λ_* via (h^a, h^b) . Say Λ **induces a decomposition of Λ_*** (via (h^a, h^b)) if $(T_*^a \setminus h^a(T^a)) \times (T_*^b \setminus h^b(T^b))$ is closed (in $T_*^a \times T_*^b$) and forms a belief-closed subset of $T_*^a \times T_*^b$.

Note, by definition, Λ induces a decomposition of Λ_* only if Λ can be embedded into Λ_* . When Λ induces a decomposition of Λ_* , Λ_* can be viewed as the union of two type structures. The first is the structure induced by Λ . (Refer back to Section 3 for the definition.) The second we'll call the **difference structure**. This consists of

$$(\Lambda_* \setminus \Lambda) = \langle \Theta; T_*^a \setminus h^a(T^a), T_*^b \setminus h^b(T^b); \kappa_*^a, \kappa_*^b \rangle.$$

Here, $T_*^a \setminus h^a(T^a)$ (resp. $T_*^b \setminus h^b(T^b)$) is a closed subset of the Polish space T_*^a (resp. T_*^b), and so Polish. The map $\kappa_*^a : (T_*^a \setminus h^a(T^a)) \rightarrow \Delta(\Theta \times (T_*^b \setminus h^b(T^b)))$ (resp. $\kappa_*^b : (T_*^b \setminus h^b(T^b)) \rightarrow \Delta(\Theta \times (T_*^a \setminus h^a(T^a)))$) is defined so that $\kappa_*^a(t_*^a)(E) = \lambda_*^a(t_*^a)(E)$ (resp. $\kappa_*^b(t_*^b)(E) = \lambda_*^b(t_*^b)(E)$) for each event E in $\Theta \times (T_*^b \setminus h^b(T^b))$ (resp. $\Theta \times (T_*^a \setminus h^a(T^a))$).⁸ The difference structure is indeed a type structure, in the sense of Definition 2.1. Lemma A1 in Appendix A establishes this formally.

Here is the result.⁹

Lemma 5.1 Fix Θ -based structures Λ and Λ_* , so that Λ induces a decomposition of Λ_* . Fix, also, a Θ -based game Γ so that (Γ, Λ) has an equilibrium. Then, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium of the difference game $(\Gamma, (\Lambda_* \setminus \Lambda))$.

As a consequence of Lemma 5.1 and the Pull-Back Property, we have the following:

Proposition 5.1 Fix Θ -based structures Λ and Λ_* , so that Λ induces a decomposition of Λ_* . Fix, also, a Θ -based game Γ , so that (Γ, Λ) has an equilibrium. Then, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the game (Γ, Λ_*) .

5.2 Mutual Absolute Continuity and The Common Prior Assumption

Now, we turn to the question: Is it of interest to consider the case where Λ induces a decomposition of Λ_* ? We will show that there is a notable case in which Λ does induce a decomposition of Λ_* .

As an intermediate step, consider the following restriction on the type structure.

Definition 5.2 Say a Θ -based interactive type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ is **countable** if T^a and T^b are countable and endowed with the discrete topology.

Definition 5.3 (Stuart [?, 1997]) Say a Θ -based interactive type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ is **mutually absolutely continuous** if Λ is countable and

⁸Note, if E is Borel in $\Theta \times (T_*^b \setminus h^b(T^b))$, then it is Borel in $\Theta \times T_*^b$. As such, the map κ_*^a is well-defined.

⁹Proofs for this section can be found in Appendix A.

- (i) for each $t^a \in T^a$, $\lambda^a(t^a)(\Theta \times \{t^b\}) > 0$ implies $\lambda^b(t^b)(\Theta \times \{t^a\}) > 0$, and
(ii) for each $t^b \in T^b$, $\lambda^b(t^b)(\Theta \times \{t^a\}) > 0$ implies $\lambda^a(t^a)(\Theta \times \{t^b\}) > 0$.¹⁰

The condition of mutual absolute continuity has an important relationship to the Equilibrium Extension Property. Suppose we can embed a structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ into a structure $\Lambda_* = \langle \Theta; T_*^a, T_*^b; \lambda_*^a, \lambda_*^b \rangle$. Further, assume that Λ_* is mutually absolutely continuous. Then, a type of Ann, viz. t_*^a , assigns strictly positive probability to a type of Bob, viz. t_*^b , if and only if t_*^b also assigns strictly positive probability to t_*^a . Now, consider a type t_*^a that is not contained in the structure induced by Λ . Can the type t_*^a assign strictly positive probability to a type of Bob in the structure induced by Λ ? No. The structure induced by Λ is a belief-closed subset. So, types in this structure cannot assign positive probability to the type t_*^a , which is what mutual absolute continuity would require. As such, the type t_*^a must assign probability one to types in the difference structure. That is, Λ induces a decomposition of Λ_* .

Indeed, we have the following result.

Proposition 5.2 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* and so that Λ_* satisfies mutual absolute continuity. Fix, also, a Θ -based game Γ , so that (Γ, Λ) has an equilibrium. Then, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the game (Γ, Λ_*) .*

Economic models often impose the condition of a common prior assumption (CPA). The idea of the CPA is that differences in beliefs reflect only differences in information. That is, if an outside observer looks at the situation, he would be able to understand the different beliefs (i.e., associated with different types) as reflecting some underlying belief, common to both players. Each type of each player reflects the conditional of this belief on certain information.

In this situation, what does Ann think Bob thinks about Ann? Can a type of Ann consider it possible that Bob considers that type of Ann impossible? The answer would seem to be no. In particular, this appears to require that Ann considers it possible that Bob has learned certain information that is inconsistent with the information she herself learned. This suggests that, if a type structure satisfies the CPA, then it is mutually absolutely continuous. Indeed, this is the case, and so the CPA has important implications for the Equilibrium Extension Property.

Let us formalize this idea. Fix a Θ -based type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$. Write $[t^a]$ for the event $\Theta \times \{t^a\} \times T^b$, and likewise with a and b interchanged. Given a measure $\mu \in \Delta(\Theta \times T^a \times T^b)$ with $\mu([t^a]) > 0$, write $\mu(\cdot | [t^a])$ for conditional of μ on $[t^a]$ and write $\text{marg}_{\Theta \times T^b} \mu$ for the marginal of μ on $\Theta \times T^b$.

Definition 5.4 *Fix a Θ -based interactive type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$. Call $\mu \in \Delta(\Theta \times T^a \times T^b)$ a **common prior (for Λ)** if Λ is countable and, for all $t^a \in T^a$,*

¹⁰As Stuart [?, 1997] points out, this is not quite mutual absolute continuity, in the sense of probability theory.

(i) $\mu([t^a]) > 0$,

(ii) $\lambda^a(t^a) = \text{marg}_{\Theta \times T^b} \mu(\cdot || [t^a])$,

and, likewise, with a and b reversed. Say the structure Λ **admits a common prior** if there is a common prior for Λ .

Lemma 5.2 Fix a Θ -based interactive type structure $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$, where Λ admits a common prior. Then, Λ is mutually absolutely continuous.

Proof. Now, let μ be a common prior for Λ . Note that

$$\lambda^a(t^a)(\Theta \times \{t^b\}) = \frac{\mu(\Theta \times \{t^a\} \times \{t^b\})}{\mu(\Theta \times \{t^a\} \times T^b)}.$$

So, $\lambda^a(t^a)(\Theta \times \{t^b\}) > 0$ if and only if $\mu(\Theta \times \{t^a\} \times \{t^b\}) > 0$. But an analogous argument gives that $\lambda^b(t^b)(\Theta \times \{t^a\}) > 0$ if and only if $\mu(\Theta \times \{t^a\} \times \{t^b\}) > 0$. This establishes the result. ■

As a corollary of Lemma 5.2:

Corollary thesection.1 Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* and so that Λ_* admits a common prior. Fix, also, a Θ -based game Γ , so that (Γ, Λ) has an equilibrium. Then, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the game (Γ, Λ_*) .

Note, Corollary thesection.1 involves a restriction on the type structure. Namely, it requires that the structure Λ be embedded into the (countable) structure Λ_* , where Λ_* admits a common prior. But this does not imply that $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property. Rather, it says that $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for a given game Γ , provided that (Γ, Λ_*) has a Bayesian Equilibrium. So, the result involves a restriction on the game, too. In this sense, we have a limited Equilibrium Extension Property.

Again, recall, Corollary thesection.1 is false absent the common prior restriction (or absent the conditions in Propositions 5.1-5.2). In the example of Section 4.2, $\langle \Lambda_1, \Lambda_2 \rangle$ fails the Equilibrium Extension Property for Γ_1 , despite the fact that there is an equilibrium of the Bayesian Game (Γ_1, Λ_2) .

6 Restrictions on Games

In Section 5, we considered restrictions on the type structure. It was shown that, when the type structure satisfies certain conditions, we do get an Equilibrium Extension Property, albeit in a limited sense. Now, we show that, when the game itself satisfies certain conditions, we again get a limited Extension Property.

Refer back to the negative results in Section 4. Both results began the construction with a Θ_0 -based game $\Gamma_0 = \langle \Theta_0; C_0^a, C_0^b; \pi_0^a, \pi_0^b \rangle$ and a Θ_0 -based structure $\Lambda_0 = \langle \Theta_0; T_0^a, T_0^b; \lambda_0^a, \lambda_0^b \rangle$, so that (Γ_0, Λ_0) has no Bayesian Equilibrium. We can go further, and take Θ_0 , C_0^a , and C_0^b to be compact metrizable, and T_0^a and T_0^b to be singletons. (See Section 7c.) In this case, Glicksberg's Theorem [?, 1952] says that one of the payoff functions—i.e., π_0^a or π_0^b —must be discontinuous. With this, each of the constructed Θ_1 -based games (i.e., the games called Γ_1 in Sections 4.1-4.2) also has a discontinuous payoff function.

What if the game studied is continuous? For some such games, there is indeed an Extension Property.

Definition 6.1 *Say a Θ -based game, viz. $\Gamma = \langle \Theta; C^a, C^b; \pi^a, \pi^b \rangle$, is **compact and continuous** if C^a and C^b are each compact and π^a and π^b are each continuous.*

Proposition 6.1 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* via (h^a, h^b) and $T_*^a \setminus h^a(T^a), T_*^b \setminus h^b(T^b)$ is finite. Suppose Γ is compact and continuous. Then, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ .*

The proof can be found in Appendix B. Here, we give the idea. Doing so will illuminate the role of the various conditions, including the restriction on the type structures.

Suppose Λ can be embedded into Λ_* via (h^a, h^b) . Fix an equilibrium (s^a, s^b) of the Bayesian Game (Γ, Λ) . We want to show that there is a Bayesian Equilibrium of the game (Γ, Λ_*) , viz. (s_*^a, s_*^b) , that extends the equilibrium (s^a, s^b) , i.e., that satisfies $s^a = s_*^a \circ h^a$ and $s^b = s_*^b \circ h^b$.

We will begin by constructing a certain game of complete information, viz. G , that depends on the game Γ and the equilibrium (s^a, s^b) . There will be many players in this game, each corresponding to a type in $T_*^a \setminus h^a(T^a)$ and $T_*^b \setminus h^b(T^b)$. As such, there are a finite number of players in this game. Each such player gets to make a choice from C^a or C^b , as in Γ . The payoff functions will be constructed in a specific way. In particular, they will depend on Γ and the equilibrium (s^a, s^b) .

The complete information game G is compact and continuous. Compactness follows from the fact that the underlying game is compact. Continuity uses the fact that the underlying game is continuous—but does not follow immediately from this fact. (Recall, the payoff functions depend on the equilibrium and any given equilibrium may be discontinuous.)

Now we have a compact and continuous complete information game G , with a finite number of players. As such, there is a mixed strategy equilibrium of G .

Finally, we return to the Bayesian game (Γ, Λ_*) . We consider the strategies that extend the equilibrium (s^a, s^b) of (Γ, Λ) . We show that, in a certain sense, these strategies correspond to the mixed strategies of the complete information game G . As such, we can use the fact that there is a mixed strategy equilibrium of G to show that there is an equilibrium of (Γ, Λ_*) that extends the equilibrium (s^a, s^b) of (Γ, Λ) .

7 Discussion

This section discusses the relationship to the literature, and further discusses some of the results in the paper.

a. Misspecifying the Type Structure: The idea that the analyst can misspecify the type structure is not new to this paper. It can also be found in the robustness and interim rationalizability literatures. Let us review the canonical questions in these areas:

(i) What if the analyst misspecifies players' actual higher-order beliefs?

(ii) What if the analyst misspecifies the players' parameter set?

For (i): The sensitivity of an analysis to players' higher-order beliefs has a long history in game theory. This insight goes back to Geanakoplos-Polemarchakis [?, 1982], followed by Monderer-Samet [?, 1989], Rubinstein [?, 1989], and Carlsson-van Damme [?, 1993]. Kajii-Morris [?, 1997] shows how misspecification of these beliefs matter for a Bayesian Equilibrium analysis.

The idea is that players may think that only parameters in E (a subset of Θ) are possible, they may think that others think the same, etc..., up to some m^{th} -order belief. In this case, we will say (informally) that the event E satisfies mutual belief up to level m . The analyst looks at this situation and incorrectly deduces that the event E satisfies mutual belief at all levels. But, in fact, Ann's $(m + 1)^{\text{th}}$ -order belief considers the possibility that E is not mutually believed up to level m . As such, the analyst misspecifies players' hierarchies of beliefs—even if only by a little bit.

Here, we do not change players' hierarchies of beliefs. Return to our examples. Note that types t_1^a and t_1^b induce the exact same hierarchies of beliefs in both Λ_1 and Λ_2 . The only difference between the two situations is the context within which they lie.

For (ii): This question is discussed in papers such as Battigalli-Siniscalchi [?, 2003], Ely-Peski [?, 2006] and Dekel-Fudenberg-Morris [3, 2007].

The idea is that players may observe signals external to the game as specified. By conditioning their choices on these external signals, new choices may be consistent with equilibrium play. Formally, this is the idea that the analyst's parameter set is Θ , while the players' parameter set is $\Theta \times \Sigma$, where Σ is a payoff-irrelevant signal space. Liu [?, 2004] shows that this is equivalent to the case in which the analyst uses a non-redundant Θ -based structure, but the players, in fact, use a redundant Θ -based structure (as formalized in Battigalli-Siniscalchi [?, 2003], Ely-Peski [?, 2006] and Dekel-Fudenberg-Morris [3, 2007]).

Here, we restrict attention to non-redundant structures. As such, we implicitly assume that the analyst 'understands' the players' parameter set (or the signals the players may see). The reason is to separate these two difficulties: By doing so, we see that, even if the analyst correctly specifies the signals the players may observe, the analyst's predictions may still differ from the players' predictions. This can happen if the analyst fails to 'understand' the context of the game.

b. The Context of the Game: There are two distinct views of a game. Under the first view, the game itself is a complete description of all interactions past, present, and future. See, for instance, the discussion in Kohlberg-Mertens [?, 1986, pp. 1005]. Under the second view, it is impractical to write down “the big game.” Instead, the game studied represents a snapshot of the strategic situation. This is a game-theoretic analog to Savage’s [?, 1954] Small Worlds view in decision theory.

Our position is that each of these views is of interest—both deserve to be studied. Here, we focus on the second view, where there is a history prior to the given game. As such, it seems natural to consider the case where the history influences which hierarchies of beliefs players can hold. That is, it seems natural to consider the case in which the history determines the context of the game.

c. The Construction: The negative results in Section 4 constructed Bayesian Games based on a “germ game” that had no equilibrium. The particular game used is irrelevant to the analysis. What is important, however, is that such a game exists. To see that we can find such a game, recall that Sion-Wolfe [?, 1957] provides an example of a complete information game in which there is no mixed strategy equilibrium. Turn this into a Bayesian Game, with a singleton parameter set and singleton type sets. This game has no Bayesian Equilibrium.

Both negative results make use of this non-existence “germ game.” As we pointed out, the example in Section 4.1 can be viewed as a corollary of non-existence of equilibrium in Bayesian Games. But, importantly, the example in Section 4.2 cannot. In particular, that game does have an equilibrium. Moreover, we may not be able to construct this extension failure for other solution concepts that fail existence in some “germ game.”

To see this last claim, think of the correlated rationalizability solution concept. There are games for which there are no rationalizable strategies. (See Example 2 in Dufwenberg-Stegeman [4, 2002].) Yet, we have the following result.

Result: Fix Θ -based structures, Λ and Λ_* , so that Λ can be properly embedded into Λ_* . Fix also a Θ -based game Γ . If the rationalizable strategies are non-empty in both the game (Γ, Λ) and (Γ, Λ_*) , then (Λ, Λ_*) satisfies rationalizable extension and pull-back properties.¹¹

As such, we cannot use the example in Dufwenberg-Stegeman as a “germ game” to resurrect the second extension failure. (Note, carefully, in Section 4.2, both (Γ_1, Λ_1) and (Γ_1, Λ_2) did have an equilibrium, and we still failed extension.)

In sum: The second extension failure not only makes use of a “germ game.” It also makes use of the fact that, in a Bayesian Equilibrium, each type has a correct belief.

¹¹The proof can be found in the Online Appendix. Dekel-Fudenberg-Morris [3, 2007] show an analogous result, when the parameter and action sets are finite.

d. Large Type Structures: In Section 4.3, we said that, if we want a type structure that captures all possible predictions, this structure may need to be redundant, i.e., may need to contain two types that induce the same hierarchies of beliefs. Redundant structures have also played a role in other papers. Refer to part a(ii) above and note that redundancy may also be appropriate if the analyst misspecifies the players' parameter sets.

It is important to note that these two types of redundancies are quite different. In the first case, the redundancies are introduced by adding a belief-closed subset that is “identical” to the one already present. (Refer to Figure 4.5 on this point.) This would not be appropriate for the second case. To see this, refer to Example 1 in Dekel-Fudenberg-Morris [3, 2007] and consider the type structure there with singleton type sets. Adding an identical belief-closed subset does not change the analysis, and so does not capture all possible predictions, for the case where the analyst misspecified the parameter set. On the other hand, Dekel-Fudenberg-Morris [3, 2007] provides an example of a different redundant type structure that would be appropriate to capture those predictions.

e. The Common Prior Assumption: Definition 5.4 states that the CPA reflects two requirements, a common prior requirement and a positivity requirement.

Consider the sets $[t^a] = \Theta \times \{t^a\} \times T^b$ and note that these sets form a partition of $\Theta \times T^a \times T^b$. Write τ^a for the subalgebra generated by this partition. Given a measure $\mu \in \Delta(\Theta \times T^a \times T^b)$ and an event E in $\Theta \times T^a \times T^b$, write $\mu(E, \cdot || \tau^a)$ for a version of μ -conditional probability of E given τ^a . (Note, since the conditioning events for Ann and Bob are distinct, the versions of conditional probability will also be distinct.) The **common prior requirement** is: There exists a measure $\mu \in \Delta(\Theta \times T^a \times T^b)$ and a version of μ -conditional probability of E given τ^a so that, for any type t^a and any event E in $[t^a]$, $\lambda^a(t^a)(\text{proj}_{\Theta \times T^b} E) = \mu(E, [t^a] || \tau^a)$. (Note, $\mu(E, \cdot || \tau^a)$ is constant on $[t^a]$.) **Positivity** requires that, in addition, $\mu([t^a]) > 0$, for each type $t^a \in T^a$.

The positivity requirement is important for Corollary thesection.1. To see this, return to the type structure Λ_2 , in Section 4.2. Per part c above, we can take Θ_0 , T_0^a , and T_0^b to be singletons, so that $\Theta_1 = \{\theta_0, \theta_1\}$, $T_2^a = \{t_0^a, t_1^a\}$, and $T_2^b = \{t_0^b, t_1^b\}$. This structure satisfies the common prior requirement. In particular, we can choose the common prior μ so that $\mu(\theta_1, t_1^a, t_1^b) = 1$.

Let us verify that this measure μ is indeed a common prior. To see this, fix an event E in $\Theta_1 \times T_2^a \times T_2^b$. For any state $(\cdot, t_0^a, \cdot) \in [t_0^a]$, let $\mu(E, (\cdot, t_0^a, \cdot) || \tau^a)$ equal $1-p$ (resp p) if $(\theta_1, t_0^a, t_1^b) \in E$ and $(\theta_0, t_0^a, t_0^b) \notin E$ (resp. $(\theta_1, t_0^a, t_1^b) \notin E$ and $(\theta_0, t_0^a, t_0^b) \in E$). Let $\mu(E, (\cdot, t_0^a, \cdot) || \tau^a)$ equal 1 if $(\theta_0, t_0^a, t_0^b), (\theta_1, t_0^a, t_1^b) \in E$. Otherwise, let $\mu(E, (\cdot, t_0^a, \cdot) || \tau^a)$ equal 0. Next, consider a state $(\cdot, t_1^a, \cdot) \in [t_1^a]$. For any such state, let $\mu(E, (\cdot, t_1^a, \cdot) || \tau^a)$ equal 1 if $(\theta_1, t_1^a, t_1^b) \in E$. Otherwise, set $\mu(E, (\cdot, t_1^a, \cdot) || \tau^a)$ equal to 0. It is readily verified that this map satisfies the requirements of a (regular and proper) version of μ -conditional probability.

So, we do indeed have a common prior for Λ_2 . This prior is not positive—it assigns zero probability to $[t_0^a]$ and $[t_0^b]$. Of course, Corollary thesection.1 tells us it cannot be. Put differently,

we can now see that the common prior requirement alone does not suffice for Corollary thesection.1. We also need the positivity requirement.

The need for the positivity requirement is important from the perspective of generalizing Corollary thesection.1. In particular, if T^a is uncountably infinite, there is no probability measure that assigns strictly positive probability to each event $[t^a]$. (See Remark ??.) This suggests a limitation to Corollary thesection.1. (Alternatively, this might suggest that other tools are needed to study the case of uncountably infinite spaces—i.e., lexicographic probability systems [?, 1991], conditional probability systems [?, 1955], or non-standard probabilities.)

There is an interesting connection to be made at the conceptual level. Does a non-positive common prior fit with the CPA? Arguably not. Recall, the idea of the CPA is that differences in probabilities only reflect differences in information. As a consequence, the only personalistic features of probability come from informational differences. But, there may be many (regular and proper) versions of conditional probability. Given this, the common prior requirement (as specified above) need not pin down the beliefs (i.e., each $\lambda^a(t^a), \lambda^b(t^b)$). Indeed, in the example above, there are many Θ -based structures Λ corresponding to the common prior μ . In fact, choosing distinct probabilities p gives just such structures.

f. Restrictions on the Game: Section 6 showed an Equilibrium Extension Property for games that are compact and continuous. Of course, many games of interest, e.g., auction and voting games, fail the continuity property, despite the fact that they do have an equilibrium.

While Proposition 6 gives a positive result, it does not say that the Equilibrium Extension Property is satisfied whenever the game is compact and continuous. The result also imposes a condition on the type structure—it requires that there are a finite number of types in $T_*^a \setminus h^a(T^a), T_*^b \setminus h^b(T^b)$. We conjecture, but do not know, that the result does not obtain absent this condition.

Appendix A Proofs for Section 5

Lemma A1 *Suppose Λ induces a decomposition of Λ_* . Consider the difference structure, viz.*

$$(\Lambda_* \setminus \Lambda) = \langle \Theta; T_*^a \setminus h^a(T^a), T_*^b \setminus h^b(T^b); \kappa_*^a, \kappa_*^b \rangle.$$

Then, $(\Lambda_ \setminus \Lambda)$ is an interactive type structure, in the sense of Definition 2.1.*

Proof. First, $(T_*^a \setminus h^a(T^a))$ and $(T_*^b \setminus h^b(T^b))$ are Polish. To see this, note that $(T_*^a \setminus h^a(T^a)) \times (T_*^b \setminus h^b(T^b))$ is closed in $T_*^a \times T_*^b$, and so $(T_*^a \setminus h^a(T^a))$ (resp. $(T_*^b \setminus h^b(T^b))$) is closed in T_*^a (resp. T_*^b). Since $(T_*^a \setminus h^a(T^a))$ (resp. $(T_*^b \setminus h^b(T^b))$) is a closed subsets of a Polish space, it is Polish (Aliprantis-Border [?, 1999; page 73]).

Now, we turn to show that κ_*^a and κ_*^b are measurable. We will show this for κ_*^a ; an analogous argument can be made for κ_*^b .

Fix an event E in $\Delta(\Theta \times (T_*^b \setminus h^b(T^b)))$. We want to show that $(\kappa_*^a)^{-1}(E)$ is an event in $T_*^a \setminus h^a(T^a)$. To show this, it suffices to find an event F in $\Delta(\Theta \times T_*^b)$ so that $(\kappa_*^a)^{-1}(E) = (\lambda_*^a)^{-1}(F)$. If so, then the measurability of κ_*^a follows from the measurability of λ_*^a and the fact that we endow $T_*^a \setminus h^a(T^a)$ with the induced topology.

Let F be the set of measures $\mu \in \Delta(\Theta \times T_*^b)$ so that there is some measure $\nu \in E$, with $\mu(G) = \nu(G)$ for every event G in $\Theta \times (T_*^b \setminus h^b(T^b))$. Then F is an event in $\Delta(\Theta \times T_*^b)$. (This is an immediate consequence of Lemma 14.16 in Aliprantis-Border [?, 1999].)

To show that $(\kappa_*^a)^{-1}(E) = (\lambda_*^a)^{-1}(F)$: Fix a type $t_*^a \in T_*^a \setminus h^a(T^a)$ with $\kappa_*^a(t_*^a) \in E$. Then, certainly, there is some measure $\mu \in F$ so that $\mu(G) = \kappa_*^a(t_*^a)(G)$ for each event G in $\Theta \times (T_*^b \setminus h^b(T^b))$. By construction, $\mu = \lambda_*^a(t_*^a)$, so that $t_*^a \in (\lambda_*^a)^{-1}(F)$. Conversely, fix a type $t_*^a \in T_*^a$ with $\lambda_*^a(t_*^a) \in F$. Then, there is some measure $\nu \in E$ with $\lambda_*^a(t_*^a)(G) = \nu(G)$ for every event G in $\Theta \times (T_*^b \setminus h^b(T^b))$. In particular, $\lambda_*^a(t_*^a)(\Theta \times (T_*^b \setminus h^b(T^b))) = \nu(\Theta \times (T_*^b \setminus h^b(T^b))) = 1$, and so $t_*^a \in T_*^a \setminus h^a(T^a)$. By construction, $\kappa_*^a(t_*^a) = \nu$. With this, $t_*^a \in (\kappa_*^a)^{-1}(E)$, as required. ■

The next lemma will be of use throughout the appendices.

Lemma A2 *Let Λ and Λ_* be two Θ -based interactive type structures, so that Λ can be embedded into Λ_* via (h^a, h^b) . Then, for each $t^a \in T^a$, $\lambda^a(t^a)(E) = \lambda_*^a(h^a(t^a))((\text{id} \times h^b)(E))$ for every event E in $\Theta \times T^b$.*

Proof. Fix some $t^a \in T^a$ and some event E in $\Theta \times T^b$. Property 3.2 gives that $(\text{id} \times h^b)(E)$ is an event in $\Theta \times T_*^b$, and $E = (\text{id} \times h^b)^{-1}((\text{id} \times h^b)(E))$. Now, using the fact that (h^a, h^b) is a type morphism,

$$\lambda_*^a(h^a(t^a))((\text{id} \times h^b)(E)) = \lambda^a(t^a)((\text{id} \times h^b)^{-1}((\text{id} \times h^b)(E))) = \lambda^a(t^a)(E),$$

as required. ■

Proof of Lemma 5.1. Suppose Λ induces a decomposition of Λ_* via (h^a, h^b) . Write

$$\Lambda(h^a, h^b) = \langle \Theta; h^a(T^a), h^b(T^b); \kappa^a, \kappa^b \rangle$$

for the structure induced by Λ , and write

$$(\Lambda_* \setminus \Lambda) = \langle \Theta; T_*^a \setminus h^a(T^a), T_*^b \setminus h^b(T^b); \kappa_*^a, \kappa_*^b \rangle$$

for the difference structure. By Lemma A1, the difference structure is a well-defined type structure.

Fix a Θ -based game Γ and an equilibrium (s^a, s^b) for the Bayesian Game (Γ, Λ) . Suppose there exists an equilibrium for the difference game $(\Gamma, (\Lambda_* \setminus \Lambda))$, viz. $(s_{\nabla}^a, s_{\nabla}^b)$. Construct a strategy, viz. s_*^a , for Ann in (Γ, Λ_*) , as follows. For each $t^a \in T^a$, let $s_*^a(h^a(t^a)) = s^a(t^a)$. (This is well-defined, since h^a is injective.) For each $t_*^a \in T_*^a \setminus h^a(T^a)$, let $s_*^a(t_*^a) = s_{\nabla}^a(t_*^a)$. Likewise for Bob.

Note that s_*^a is measurable, and so a strategy of (Γ, Λ_*) : To see this, fix an event E in $\Delta(C^a)$. Note $(s_*^a)^{-1}(E) = (s_{\nabla}^a)^{-1}(E) \cup h^a((s^a)^{-1}(E))$. Since h^a is bimeasurable (Property 3.2) and s_*^a, s_{∇}^a are measurable, $(s_*^a)^{-1}(E)$ is the union of two measurable sets and so measurable.

We will now show that (s_*^a, s_*^b) is an equilibrium for the game (Γ, Λ_*) .

First, fix a type $h^a(t^a) \in h^a(T^a)$. Given a strategy $r_*^a : T_*^a \rightarrow \Delta(C^a)$, write r^a for the strategy $r^a : T^a \rightarrow \Delta(C^a)$ with $r^a(u^a) = r_*^a(h^a(u^a))$ for all $u^a \in T^a$. Likewise for b . Then, using Lemma A2, for any strategies r_*^a and r_*^b ,

$$\int_{\Theta \times T_*^b} \pi^a(\theta, r_*^a(h^a(t^a)), r_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)) = \int_{\Theta \times T^b} \pi^a(\theta, r^a(t^a), r^b(t^b)) d\lambda^a(t^a).$$

So, using the fact that (s^a, s^b) is a Bayesian Equilibrium of (Γ, Λ) ,

$$\int_{\Theta \times T_*^b} \pi^a(\theta, s_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)) \geq \int_{\Theta \times T_*^b} \pi^a(\theta, r_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)), \quad (\text{A1})$$

for all strategies r_*^a of the game (Γ, Λ_*) .

Turn to a type $t_*^a \in T_*^a \setminus h^a(T^a)$. Given a strategy $r_*^a : T_*^a \rightarrow \Delta(C^a)$, write r_{∇}^a for the strategy $r_{\nabla}^a : T_*^a \setminus h^a(T^a) \rightarrow \Delta(C^a)$ with $r_{\nabla}^a(t_*^a) = r_*^a(t_*^a)$ for all $t_*^a \in T_*^a \setminus h^a(T^a)$. Likewise for b . Certainly,

$$\int_{\Theta \times T_*^b} \pi^a(\theta, r_*^a(t_*^a), r_*^b(t_*^b)) d\lambda_*^a(t_*^a) = \int_{\Theta \times T_*^b \setminus h^b(T^b)} \pi^a(\theta, r_{\nabla}^a(t_*^a), r_{\nabla}^b(t_*^b)) d\lambda_*^a(t_*^a).$$

So, using the fact that $(s_{\nabla}^a, s_{\nabla}^b)$ is a Bayesian Equilibrium of $(\Gamma, \Lambda_* \setminus \Lambda)$,

$$\int_{\Theta \times T_*^b} \pi^a(\theta, s_*^a(t_*^a), s_*^b(t_*^b)) d\lambda_*^a(t_*^a) \geq \int_{\Theta \times T_*^b} \pi^a(\theta, r_{\nabla}^a(t_*^a), s_*^b(t_*^b)) d\lambda_*^a(t_*^a), \quad (\text{A2})$$

for all strategies r_*^a of the game (Γ, Λ_*) .

Taking Equations A1-A2, together with analogous statements for player b , gives that (s_*^a, s_*^b) is an equilibrium for the game (Γ, Λ_*) . The converse follows immediately from the Pull-Back Property, i.e., Proposition 3.1. ■

Proof of Proposition 5.1. If $\langle \Lambda, \Lambda_* \rangle$ satisfies the Extension Property for Γ , then it is immediate that there is an equilibrium for the game (Γ, Λ_*) . Conversely, suppose there is an equilibrium for the game (Γ, Λ_*) . By the Pull-Back Property (Proposition 3.1), there is an equilibrium for the difference game $(\Gamma, (\Lambda_* \setminus \Lambda))$. Now, using Lemma 5.1, $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ . ■

Lemma A3 *Fix non-redundant Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* via (h^a, h^b) . If Λ_* is mutually absolutely continuous, then Λ induces a decomposition of Λ_* via (h^a, h^b) .*

Proof. By Lemma 3.1, we can take $T_*^a \setminus h^a(T^a)$ to be non-empty. In particular, fix $t_*^a \in T_*^a \setminus h^a(T^a)$. Recall, since Λ_* is mutually absolutely continuous, it is countable. As such, we can find some $t_*^b \in T_*^b$ with $\lambda_*^a(t_*^a)(\Theta \times \{t_*^b\}) > 0$. Again using the fact that Λ_* is mutually absolutely continuous, $\lambda_*^b(t_*^b)(\Theta \times \{t_*^a\}) > 0$. So, by Lemma A2, $t_*^b \in T_*^b \setminus h^b(T^b)$. That is, $\lambda_*^a(t_*^a)(\Theta \times (T_*^b \setminus h^b(T^b))) = 1$. In addition, we have established that $T_*^b \setminus h^b(T^b)$ is non-empty. Reversing the argument for $t_*^b \in T_*^b \setminus h^b(T^b)$, we get also that $\lambda_*^b(t_*^b)(\Theta \times (T_*^a \setminus h^a(T^a))) = 1$. This establishes the result. ■

Proof of Proposition 5.2. Immediate from Lemma A3 and Proposition 5.1. ■

Appendix B Proofs for Section 6

This appendix is devoted to proving Proposition 6.1. Fix two (non-redundant) Θ -based type structures $\Lambda = \langle \Theta; T^a, T^b; \lambda^a, \lambda^b \rangle$ and $\Lambda_* = \langle \Theta; T_*^a, T_*^b; \lambda_*^a, \lambda_*^b \rangle$. Suppose, further, that Λ can be properly embedded into Λ_* via (h^a, h^b) , so that $T_*^a \setminus h^a(T^a)$ and $T_*^b \setminus h^b(T^b)$ are finite (and possibly empty). By Lemma 3.1, it is without loss of generality to assume that $T_*^a \setminus h^a(T^a)$ is non-empty. Write $T_*^a \setminus h^a(T^a) = \{1, \dots, M\}$. The set $T_*^b \setminus h^b(T^b)$ may or may not be empty. We will begin by assuming that $T_*^b \setminus h^b(T^b)$ is non-empty, and can be written as $T_*^b \setminus h^b(T^b) = \{1, \dots, L\}$. We later discuss the case in which $T_*^b \setminus h^b(T^b)$ is empty.

Consider a Θ -based compact and continuous game $\Gamma = \langle \Theta; C^a, C^b; \pi^a, \pi^b \rangle$. Fix a Bayesian Equilibrium of the game (Γ, Λ) , viz. (s^a, s^b) . We want to show that there is a Bayesian Equilibrium of the game (Γ, Λ_*) , viz. (s_*^a, s_*^b) , with $s^a = s_*^a \circ h^a$ and $s^b = s_*^b \circ h^b$.

Section 6 gives the idea of the proof. In particular, we begin by constructing the game of complete information, namely G . The game has $M + L$ players, corresponding to $T_*^a \setminus h^a(T^a) = \{1, \dots, M\}$ and $T_*^b \setminus h^b(T^b) = \{1, \dots, L\}$. Write m for a player in $\{1, \dots, M\}$ and l for a player in $\{1, \dots, L\}$. The choice set for a player $m \in T_*^a \setminus h^a(T^a)$ is C^a , and the choice set for a player $l \in T_*^b \setminus h^b(T^b)$ is C^b .

We now define the payoff functions. In particular, we will have certain maps $\nu^m : C^a \rightarrow \mathbb{R}$ and $w^m : C^a \times [C^b]^L \rightarrow \mathbb{R}$, which we will soon specify. Then, we will define the payoff function for player m as $u^m : [C^a]^M \times [C^b]^L \rightarrow \mathbb{R}$, where

$$u^m(c_1^a, \dots, c_m^a, \dots, c_M^a, c_1^b, \dots, c_L^b) := v^m(c_m^a) + w^m(c_m^a, c_1^b, \dots, c_L^b).$$

The payoff function for player l is defined analogously.

Now, we turn to specifying v^m and w^m . Recall, $h^b : T^b \rightarrow T_*^b$ is injective and bimeasurable (Properties 3.1-3.2). As such, we can define a bijective and bimeasurable map $g^b : h^b(T^b) \rightarrow T^b$, so that $g^b(h^b(t^b)) = t^b$. Now, we define $v^m : C^a \rightarrow \mathbb{R}$ so that

$$v^m(c^a) := \int_{\Theta \times h^b(T^b)} \pi^a(\theta, c^a, s^b(g^b(t_*^b))) d\lambda_*^a(m).$$

We also define $w^m : C^a \times [C^b]^L \rightarrow \mathbb{R}$ so that

$$w^m(c^a, c_1^b, \dots, c_L^b) := \int_{\Theta \times (T_*^b \setminus h^b(T^b))} \pi^a(\theta, c^a, \vec{c}^b(t_*^b)) d\lambda_*^a(m),$$

where $\vec{c}^b(t_*^b) = c_l^b$, if $t_*^b = l$.

To show that the payoff function u^m is continuous, we show that v^m and w^m are continuous. The proof will make use of two technical results.

Lemma B1 *Fix metrizable spaces Ω_1, Ω_2 . Let $f : \Omega_1 \times \Omega_2 \rightarrow \mathbb{R}$ be a bounded continuous function and define $F : \Omega_1 \rightarrow \mathbb{R}$ so that*

$$F(\omega_1) = \int_{E_2} f(\omega_1, \omega_2) d\mu,$$

where E_2 is some event in Ω_2 and $\mu \in \Delta(\Omega_2)$. Then, F is a bounded continuous function.

Proof. The fact that F is bounded follows directly from the fact that f is bounded and $\mu(E_2) \leq 1$. We focus on showing that F is continuous. For this, fix a sequence $(\omega_{1,n} : n = 1, 2, \dots)$ contained in Ω_1 and suppose $\omega_{1,n} \rightarrow \omega_{1,*}$. To show that F is continuous, it suffices to show that $F(\omega_{1,n}) \rightarrow F(\omega_{1,*})$.

Write $f_*(\cdot) : \Omega_2 \rightarrow \mathbb{R}$ for the $\omega_{1,*}$ -section of the map f . Also, for each n , write $f_n(\cdot) : \Omega_2 \rightarrow \mathbb{R}$ for the $\omega_{1,n}$ -section of the map f . Let's point to two properties of these maps: First, since f is continuous, each of f_*, f_1, f_2, \dots is measurable (Aliprantis-Border [?, 1999; Theorem 4.47]). Second, since f is bounded, f_* is bounded and the sequence $(f_n : n = 1, 2, \dots)$ is uniformly bounded. Given this, it suffices to show that $f_n \rightarrow f_*$. If so, then, by the Bounded Convergence Theorem, $F(\omega_{1,n}) \rightarrow F(\omega_{1,*})$. (See Doob [?, 1993; pages 83-84].)

To show that $f_n \rightarrow f_*$: Note that $\omega_{1,n} \rightarrow \omega_{1,*}$. So, for any ω_2 , $(\omega_{1,n}, \omega_2) \rightarrow (\omega_{1,*}, \omega_2)$. Given that f is continuous, it follows that $f_n \rightarrow f_*$, as desired. ■

Lemma B2 Fix topological spaces $\Omega_1, \dots, \Omega_N, \Phi$. Let J be an integer between 1 and N , and write proj for the projection of $\Omega_1 \times \dots \times \Omega_N$ onto $\Omega_1 \times \dots \times \Omega_J$. If $f : \Omega_1 \times \dots \times \Omega_J \rightarrow \Phi$ is continuous, then $f \circ \text{proj}$ is also continuous.

Proof. Note that proj is a continuous function. Since $f \circ \text{proj}$ is the composite of two continuous functions, it is continuous. ■

Lemma B3 The function v^m is continuous.

Proof. Fix a sequence $(c_n^a : n \in \mathbb{N})$ in C^a that converges to $c^a \in C^a$. It suffices to show that, for each $\varepsilon > 0$, there exists some \bar{n} with $|v^m(c_n^a) - v^m(c^a)| < \varepsilon$ for all $n > \bar{n}$. For this, it will be useful to recall that π^a is bounded, and so we can write $\pi^a(\Theta \times C^a \times C^b) \subseteq [-\rho, \rho]$ where $\rho > 0$.

Consider the map $s^b \circ g^b$ and note that this is a measurable map from a metrizable space to a Polish space. So, by Lusin's Theorem (see Kechris [6, 1995; Theorem 17.12]), there exists a closed set $F \subseteq h^b(T^b)$ with

$$\lambda_*^a(m)(\Theta \times (h^b(T^b) \setminus F)) < \frac{\varepsilon}{4\rho}$$

and so that $(s^b \circ g^b)|_F$ is continuous. As such, the map $\Theta \times C^a \times F \rightarrow \Theta \times C^a \times \Delta(C^b)$, given by $(\theta, c^a, t_*^b) \mapsto (\theta, c^a, s^b(g^b(t_*^b)))$, is also continuous. With this and the fact that π^a is continuous, we can define a continuous mapping from $\Theta \times C^a \times F$ to \mathbb{R} by $(\theta, c^a, t_*^b) \mapsto \pi^a(\theta, c^a, s^b(g^b(t_*^b)))$.

Note, $\Theta \times C^a \times F$ is metrizable. As such, using Lemma B1, there exists an \bar{n} such that, for all $n > \bar{n}$,

$$\left| \int_{\Theta \times F} \pi^a(\theta, c_n^a, s^b(g^b(t_*^b))) d\lambda_*^a(m) - \int_{\Theta \times F} \pi^a(\theta, c^a, s^b(g^b(t_*^b))) d\lambda_*^a(m) \right| < \frac{\varepsilon}{2}.$$

Now note:

$$\begin{aligned} |v^m(c_n^a) - v^m(c^a)| &= \left| \int_{\Theta \times h^b(T^b)} [\pi^a(\theta, c_n^a, s^b(g^b(t_*^b))) - \pi^a(\theta, c^a, s^b(g^b(t_*^b)))] d\lambda_*^a(m) \right| \\ &\leq \left| \int_{\Theta \times F} [\pi^a(\theta, c_n^a, s^b(g^b(t_*^b))) - \pi^a(\theta, c^a, s^b(g^b(t_*^b)))] d\lambda_*^a(m) \right| \\ &\quad + \left| \int_{\Theta \times (h^b(T^b) \setminus F)} [\pi^a(\theta, c_n^a, s^b(g^b(t_*^b))) - \pi^a(\theta, c^a, s^b(g^b(t_*^b)))] d\lambda_*^a(m) \right| \\ &< \frac{\varepsilon}{2} + 2\rho\left(\frac{\varepsilon}{4\rho}\right) \\ &= \varepsilon, \end{aligned}$$

as desired. ■

Lemma B4 The map w^m is continuous.

Proof. For $l \in \{1, \dots, L\}$, define $w^{m,l} : C^a \times C^b \rightarrow \mathbb{R}$ by

$$w^{m,l}(c^a, c^b) = \int_{\Theta \times \{l\}} \pi^a(\theta, c^a, c^b) d\lambda_*^a(m).$$

Since π^a is continuous, Lemma B1 gives that each $w^{m,l}$ is continuous. Note:

$$w^m(c^a, c_1^b, \dots, c_L^b) = w^{m,1}(c^a, c_1^b) + \dots + w^{m,l}(c^a, c_l^b) + \dots + w^{m,L}(c^a, c_L^b).$$

So, using Lemma B2 and the fact that a finite sum of continuous functions is continuous, we have that w^m is continuous. ■

Lemma B5 *The payoff functions u^m are each continuous.*

Proof. Use Lemmata B3, B4, and B2, to get that u^m is the sum of two continuous functions. So, u^m is continuous. ■

We extend the payoff functions u^m to $[\Delta(C^a)]^M \times [\Delta(C^b)]^L$ in the usual way. And, likewise, for the payoff functions u^l .

Lemma B6 *There exists some mixed choice equilibrium for the game G .*

Proof. Note, C^a and C^b are compact metrizable. By Lemma B5, each u^m and u^l is continuous. So, the result follows from Glicksberg's Theorem [?, 1952]. ■

Let \bar{S}_*^a be the set of all strategies r_*^a of (Γ, Λ_*) satisfying $r_*^a \circ h^a = s^a$. Note, this is indeed well-defined since h^a is injective. Likewise define \bar{S}_*^b .

A standard argument establishes the next remark. (Specifically, it involves two steps. First, it changes the order of integration. Second, it uses the fact that h^a is bimeasurable to get: Any map $r_*^a : T_*^a \rightarrow \Delta(C^a)$ satisfying $(r_*^a | h^a(T^a)) = s^a \circ g^a$ is measurable if and only if $r_*^a | (T_*^a \setminus h^a(T^a))$ is measurable.)

Remark B1 *Fix some $m \in T_*^a \setminus h^a(T^a)$. For any $(r_*^a, r_*^b) \in \bar{S}_*^a \times \bar{S}_*^b$,*

$$\Pi^a(m, r_*^a, r_*^b) = u^m(r_*^a(1), \dots, r_*^a(m), \dots, r_*^a(M), r_*^b(1), \dots, r_*^b(L)).$$

Conversely, given some $(\sigma_1^a, \dots, \sigma_M^a) \in [\Delta(C^a)]^M$ and some $(\sigma_1^b, \dots, \sigma_L^b) \in [\Delta(C^b)]^L$, there exists a unique $(r_^a, r_*^b) \in \bar{S}_*^a \times \bar{S}_*^b$ with $(r_*^a(1), \dots, r_*^a(M)) = (\sigma_1^a, \dots, \sigma_M^a)$ and $(r_*^b(1), \dots, r_*^b(L)) = (\sigma_1^b, \dots, \sigma_L^b)$. In this case,*

$$\Pi^a(m, r_*^a, r_*^b) = u^m(\sigma_1^a, \dots, \sigma_m^a, \dots, \sigma_M^a, \sigma_1^b, \dots, \sigma_L^b).$$

Lemma B7 *Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* via (h^a, h^b) and $T_*^a \setminus h^a(T^a)$, $T_*^b \setminus h^b(T^b)$ are each finite and non-empty. The pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for Γ .*

Proof. Consider the game G constructed above. By Lemma B6, there exists a mixed strategy profile, viz. $(\sigma_1^a, \dots, \sigma_M^a, \sigma_1^b, \dots, \sigma_L^b)$, that is an equilibrium for this game. Now, by Remark B1, we can find a strategy profile $(s_*^a, s_*^b) \in \bar{S}_*^a \times \bar{S}_*^b$ so that $(s_*^a(1), \dots, s_*^a(M)) = (\sigma_1^a, \dots, \sigma_M^a)$ and $(s_*^b(1), \dots, s_*^b(L)) = (\sigma_1^b, \dots, \sigma_L^b)$. We will show that (s_*^a, s_*^b) is a Bayesian Equilibrium for the game (Γ, Λ_*) .

First, fix some type $h^a(t^a) \in h^a(T^a)$. Given any strategy $q_*^a : T_*^a \rightarrow \Delta(C^a)$, we have

$$\begin{aligned} \int_{\Theta \times T_*^b} \pi^a(\theta, q_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)) &= \int_{\Theta \times h^b(T^b)} \pi^a(\theta, q_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)) \\ &= \int_{\Theta \times T^b} \pi^a(\theta, q_*^a(h^a(t^a)), s_*^b(h^b(t^b))) d\lambda^a(t^a), \end{aligned}$$

where we repeatedly use Lemma A2. Using this and the fact that $(s^a, s^b) = (s_*^a \circ h^a, s_*^b \circ h^b)$ is a Bayesian equilibrium for the game (Γ, Λ) , we have that

$$\begin{aligned} \int_{\Theta \times T_*^b} \pi^a(\theta, s_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)) &= \int_{\Theta \times T^b} \pi^a(\theta, s_*^a(h^a(t^a)), s_*^b(h^b(t^b))) d\lambda^a(t^a) \\ &\geq \int_{\Theta \times T^b} \pi^a(\theta, q_*^a(h^a(t^a)), s_*^b(h^b(t^b))) d\lambda^a(t^a) \\ &= \int_{\Theta \times T_*^b} \pi^a(\theta, q_*^a(h^a(t^a)), s_*^b(t_*^b)) d\lambda_*^a(h^a(t^a)), \end{aligned}$$

for all strategies q_*^a of the game (Γ, Λ_*) .

Next, fix some type $m \in T_*^a \setminus h^a(T^a)$. Note, for any strategy q_*^a of the game, there exists a strategy $p_*^a \in \bar{S}_*^a$ with $q_*^a(m) = p_*^a(m)$ and, so, $\Pi^a(m, q_*^a, s_*^b) = \Pi^a(m, p_*^a, s_*^b)$. By Remark B1 and the fact that $(\sigma_1^a, \dots, \sigma_M^a, \sigma_1^b, \dots, \sigma_L^b)$ is an equilibrium, we have that $\Pi^a(m, s_*^a, s_*^b) \geq \Pi^a(m, q_*^a, s_*^b)$, for all strategies q_*^a of the game (Γ, Λ_*) .

We can repeat the argument, reversing a and b . This gives that, for each type $t_*^b \in T_*^b$,

$$\Pi^b(t_*^b, s_*^a, s_*^b) \geq \Pi^b(t_*^b, s_*^a, q_*^b),$$

for all strategies q_*^b of the game (Γ, Λ_*) .

As such, (s_*^a, s_*^b) is a Bayesian equilibrium of (Γ, Λ_*) . Moreover, $s_*^a \circ h^a = s^a$ and $s_*^b \circ h^b = s^b$, as required. ■

Now, consider the case where $T_*^b \setminus h^b(T^b) = \emptyset$. Here, the game G has M players. We define each u^m so that

$$u^m(c_1^a, \dots, c_m^a, \dots, c_M^a, c_1^b, \dots, c_L^b) := v^m(c_m^a).$$

Then, we can repeat the above argument to get the following Lemma.

Lemma B8 Fix Θ -based structures Λ and Λ_* , so that Λ can be properly embedded into Λ_* via (h^a, h^b) , where $T_*^b \setminus h^b(T^b) = \emptyset$. The pair $\langle \Lambda, \Lambda_* \rangle$ satisfies the Equilibrium Extension Property for

Γ.

Proof of Proposition 6.1. Immediate from Lemmata B7-B8. ■