

On the Existence of Reduced-Form Representations

Ernesto Rivera Mora

October 26, 2025

Abstract

[Rivera Mora \[2025\]](#) develops a reduced-form approach to analyze information sharing between an informed party and an uninformed party who will engage in future interactions. This reduced-form approach uses belief-dependent utilities to characterize the payoffs of Bayesian equilibria through the agents' hierarchies of beliefs. While [Rivera Mora \[2025\]](#) uses this approach to establish conditions for both the possibility and impossibility of information sharing under a broad class of mechanisms, these results do not address the existence of reduced forms and reduced-form representations. When the belief structure has a fixed domain of uncertainty, some games do not admit reduced-form representations, as certain Bayesian equilibria require types with redundant hierarchies of beliefs. The paper shows that by expanding the agents' domain of uncertainty, reduced forms and reduced-form representations exist for a broad class of games, including all games with finite action spaces. Therefore, belief-dependent utility functions are suitable for capturing the payoffs of future interactions.

1 Introduction

Rivera Mora [2024] (hereafter **RM**) analyzes information-sharing mechanisms in settings in which an informed party has future interactions with an uninformed party. Specifically, the paper studies the extent to which a broad class of mechanisms, called *neutral mechanisms*, can induce information sharing. Neutral mechanisms are defined as those that satisfy four independence conditions: (1) structural independence—action sets are independent of the state, ruling out hard evidence; (2) statistical independence—chance moves are independent of the state, ruling out experiments; (3) outcome independence—outcomes are independent of the state, ruling out state-contingent transfers; and (4) game independence—the mechanism does not change the rules of the after-game, ruling out delegation and arbitration. In other words, a mechanism is considered *neutral* if it does not depend on the realized state and cannot change the fundamental structure of future interactions. This broad class includes cheap talk, extended cheap talk, noisy communication, mediation, money burning, and message-contingent transfer schemes.

To address the reach of neutral mechanisms, **RM** develops a reduced-form approach that characterizes the agents’ payoffs in a given game of incomplete information using belief-dependent preferences. The paper studies a model with two agents: an expert who observes the payoff-relevant state of the world and a layman who does not. In the absence of a way to share information, the agents play a game of incomplete information G in which payoffs depend on actions and the state. The reduced-form approach uses belief-dependent utilities to summarize how information affects the agents’ equilibrium payoffs in G . These belief-dependent utility functions capture how the agents’ hierarchies of beliefs influence their equilibrium payoffs in G .

The reduced-form approach provides three key advantages. First, by abstracting away from the specific details of the mechanism, it yields general results that apply across all neutral mechanisms (cheap talk, mediation, money burning, etc.). Second, it leverages martingale properties of hierarchies of beliefs to identify geometric features of belief distributions consistent with Bayesian updating.¹ Third, by working directly with belief-dependent utilities, the approach makes transparent how belief-dependent payoffs shape agents’ incentives to either share or withhold different pieces of information.

The reduced-form approach used in **RM** transforms the original information sharing problem into an auxiliary mechanism design problem in which there is no after-game and agents have belief-dependent preferences². Through analysis of this auxiliary problem, **RM** uses the martingale properties of hierarchies of beliefs to establish conditions in which neutral mechanisms can fully reveal the state and conditions in which no payoff-relevant information can be shared.

While the reduced-form approach offers different advantages, it also involves some subtleties. First, the information conveyed through some mechanisms may not be commonly known. (For

¹For instance, these geometric features not only include generalizations of *Bayesian plausibility* to higher-order beliefs but also include restrictions on how some statistics correlate with the state (See **Mathevet et al.** [2020], **Rivera Mora** [2024, 2025].)

²**Rivera Mora** [2024] studies optimal mechanism design in settings with belief-dependent preferences

instance, in mediation, the expert may not observe the interaction of the layman with the mediator.) Hence, reduced forms cannot rely solely on finite-order beliefs, as behavior in games of incomplete information is sensitive to all higher-order beliefs [Rubinstein, 1989, Carlsson and Van Damme, 1993, Morris and Shin, 2001]. Second, when games have multiple equilibria, a single reduced form may not suffice to characterize equilibrium payoffs. In such cases, a reduced-form representation—i.e., a set of belief-dependent utility functions—is needed to fully capture all possible equilibrium payoffs. Third, the reduced-form approach requires the existence of reduced forms and reduced-form representations, a requirement addressed by this present paper.

Appendix B shows that there are certain instances in which the existence of a reduced form and a reduced-form representation are guaranteed. The first case is when the equilibrium behavior of a game G is uniquely characterized by the agents' hierarchies of beliefs. Specifically, if each associated Bayesian game has a unique Bayesian equilibrium that is pinned down by the agents' hierarchies, then G has a reduced-form representation. A second case is when G involves an inactive expert (i.e., the expert's action set is a singleton). Each game with an inactive expert has both a reduced form and a reduced-form representation.

Nevertheless, when working with a belief structure whose domain of uncertainty consists solely of payoff-relevant states, some games do not admit a reduced-form representation. This limitation arises because such a set of hierarchies of beliefs is not sufficiently rich to capture certain equilibrium outcomes, particularly those involving correlation between players' actions. The following example illustrates this:

Example 1.1. *There are two agents e and ℓ . Each agent $i \in \{e, \ell\}$ has action set $A_i = \{\underline{a}_i, \bar{a}_i\}$. In this game, there is no private information, so the set of payoff-relevant parameters Θ is a singleton. The game G is shown in Figure 4 of Aumann [1987]:*

	\bar{a}_ℓ	\underline{a}_ℓ
\bar{a}_e	6, 6	2, 7
\underline{a}_e	7, 2	0, 0

Figure 1.1. A game with no reduced form representation (payoff for e , payoff for ℓ)

The game G has three Nash equilibria. The payoff profiles associated with these equilibria are (2, 7), (7, 2), and $(4 + 2/3, 4 + 2/3)$. Additionally, there is a correlated equilibrium with a payoff profile of (5, 5), which lies outside the convex hull of the Nash equilibrium payoffs.

In this example, the absence of uncertainty implies that the set of payoff-relevant parameters of G is a singleton. As a result, the set of hierarchies of beliefs is also a singleton and thus it is not sufficiently rich to capture correlation generated by external signals. Consequently, under such a

belief structure no reduced form can capture the payoffs of the correlated equilibrium of G .³ (See Lemma B.4 for a proof.)

This paper establishes the existence of reduced forms and reduced-form representations for a broad class of games when considering an extended domain of uncertainty. In particular, the existence result applies to all games with finite action sets. By incorporating payoff-irrelevant parameters into the uncertainty space, the paper constructs reduced-form representations that fully capture all equilibrium outcomes, including those with correlated actions. In this sense, the result provides a theoretical foundation for using belief-dependent utilities as a tool to analyze information sharing mechanisms in general applications.

Literature Review This work serves as a complement to RM, which develops a reduced-form approach to study information sharing through a broad class of neutral mechanisms. By using the reduced-form approach, RM characterizes when information sharing is possible based on supermodularity and submodularity conditions between the state and hierarchies of beliefs.⁴ While supermodularity enables complete information sharing through a simple neutral mechanism, submodularity precludes sharing payoff-relevant information.

This paper connects the literature on mechanism design with after-games [Calzolari and Pavan, 2006a,b] with the literature on psychological games. Psychological games were introduced by Geanakoplos, Pearce, and Stacchetti [1989] to model preferences for endogenous beliefs. Battigalli and Dufwenberg [2009] extend this setup to dynamic psychological games. Within this dynamic framework, Battigalli, Corrao, and Dufwenberg [2019] provide a foundation for analyzing belief-dependent motivations in games. Rivera Mora [2024] provides a revelation principle that applies to both settings with psychological preferences and settings with after-games.

This paper contributes to a growing literature that uses belief-dependent utility functions to summarize payoffs of after-games. For instance, belief-dependent utilities have been used to capture bidders' payoffs in auctions with aftermarkets [Dworczak, 2020], to characterize future payoffs of designers with limited intertemporal commitment [Doval and Skreta, 2022, 2024], to study monopolists selling conspicuous goods that signal social status [Rayo and Segal, 2010], to summarize reputational payoffs in cheap talk models [Morris, 2001, Ottaviani and Sørensen, 2006], to summarize payoffs in mechanisms that use cheap talk and signaling [McClellan and Rappoport, 2024], and to bound the value of mediated communication [Corrao and Dai, 2023].

Belief-dependent utility functions are the key tool for most of the information design literature. Most of this literature focuses on decision problems in which only first-order beliefs are needed [Rayo and Segal, 2010, Kamenica and Gentzkow, 2011, Taneva, 2019]. An important exception is Mathevet, Peregó, and Taneva [2020], which studies information design in games and hence requires all hierarchies of beliefs.

³This is reminiscent of the literature on redundant hierarchies. (See Ely and Peski [2006].) As Liu [2009] shows, enriching the underlying space of uncertainty can capture correlation precluded by hierarchies on Θ alone.

⁴Under supermodularity, experts with high states prefer being perceived as having high states. Conversely, under submodularity, experts with low states prefer being perceived as having high states.

This work is closely related to [Ely and Peski \[2006\]](#) and [Liu \[2009\]](#), who show that type spaces with redundant hierarchies of beliefs may lead to different equilibrium predictions. With this literature in mind, the present paper studies a particular framework that allows for a belief structure with payoff redundancies.

Finally, the existence result involves linking Bayesian equilibria across different type structures. To achieve this, the results build upon the pullback property and extension property introduced by [Friedenberg and Meier \[2017\]](#).

2 Framework

This paper extends the framework in [RM](#) in two ways. First, it employs a larger domain of uncertainty that includes both a set of payoff-relevant *states* (Θ) and a set of non-payoff-relevant *labels* (Λ). Second, the definition of a neutral mechanism will now include a *labeling* of its nodes, i.e. a mapping that identifies each terminal node of the mechanism with a unique label. In contrast with the states, the labels have no impact on payoffs and solely serve to capture (payoff) redundancies on the set of hierarchies of beliefs.

Environment Let $\Theta \subsetneq \mathbb{R}$ be a finite set of states. There are two agents: an expert (e) and a layman (ℓ). Write $i \in \{e, \ell\}$ for an agent and $-i$ for the agent in $\{e, \ell\} \setminus \{i\}$. The agents' payoffs depend on the state. The state is drawn from a common prior $\mu \in \Delta(\Theta)$ with full support. The expert observes the realization of the state, while the layman does not. The agents face a simultaneous-move game in which the set of actions for agent i is a metric space A_i . Write $A := A_e \times A_\ell$. The payoff function for agent i is a continuous mapping $\pi_i : \Theta \times A \rightarrow \mathbb{R}$. Write $G = ((A_i, \pi_i) : i \in e, \ell)$ for that game. Finally, there is an infinite and countable set of labels Λ that is not relevant for the agents' payoffs.

Neutral Mechanisms A neutral mechanism is an extensive form played after the expert learns the state but before the agents play the game of interest G . These mechanisms allow agents to share information and exchange transfers. Because the mechanisms are neutral, they do not depend on the realized state and cannot change the fundamental structure of G . Since only neutral mechanisms are analyzed here, they will be referred to simply as mechanisms.

Formally, a mechanism is a profile $\mathcal{M} = ((V, \succ), X, (\mathcal{I}_i, \mathcal{A}_i : i \in e, \ell, c), \rho_c, (\gamma_i : i \in e, \ell), L)$. Here, V is a finite set of nodes and \succ is a precedence relation on V such that (V, \succ) forms a tree. Write $\emptyset \in V$ for the root of the tree and $Z \subsetneq V$ for the set of terminal nodes. (The terminal nodes in Z correspond to the start of G .) The tree (V, \succ) describes an extensive form in which moves can be made by the expert, the layman, and chance (denoted by c). The set X contains all actions in \mathcal{M} . Each agent $i \in \{e, \ell, c\}$ has an action correspondence $\mathcal{A}_i : \mathcal{I}_i \setminus \mathcal{T}_i \rightrightarrows X$ that specifies the actions that are available at each non-terminal information set.

Each $i \in \{e, \ell, c\}$ has an information partition of V , given by $\mathcal{I}_i \subseteq 2^V$. Note, \mathcal{I}_i is a partition of both non-terminal and terminal nodes in the extensive form. Each partition \mathcal{I}_i satisfies three

conditions. First, it satisfies no absentmindedness, i.e., if $\{v, v'\} \subseteq I$ and $v \neq v'$, then $v \not\preceq v'$ and $v' \not\preceq v$. Second, each information partition \mathcal{I}_i is such that i has perfect recall. Third, the mechanism has an observable end, i.e., if $I_i \cap Z \neq \emptyset$, then $I_i \subseteq Z$. Each information set $I_i \in \mathcal{I}_i$ that satisfies $I_i \subseteq Z$ is called a terminal information set; write T_i for an arbitrary terminal information set. The set of terminal information sets for i is $\mathcal{T}_i \subseteq \mathcal{I}_i$. Chance’s behavior is described by a (exogenous) behavior strategy ρ_c that specifies the distribution of chance’s actions at each information set $I_c \in \mathcal{I}_c \setminus \mathcal{T}_c$. (See Appendix A.) The transfer function $\gamma_i : \mathcal{T}_i \rightarrow \mathbb{R}$ associates each terminal information set of $i \in \{e, \ell\}$ with a transfer that i receives.⁵ Lastly, $L : Z \rightarrow \Lambda$ is an injective mapping that identifies each terminal node with a unique label. Call L the labeling of \mathcal{M} .⁶

A designer seeks to construct a mechanism in which the expert shares information prior to the game G being played. While the designer cannot affect the state Θ , she has the freedom to select the mechanism \mathcal{M} .

The supergame A mechanism \mathcal{M} and the game G together induce a supergame, denoted by (\mathcal{M}, G) . The timing of the supergame is as follows: First, Nature chooses state $\theta \in \Theta$, which is then observed by the expert. Next, the agents play \mathcal{M} , after which each agent i observes a terminal information set $T_i \in \mathcal{T}_i$, thereby obtaining a transfer $y_i = \gamma_i(T_i)$. Finally, the agents play G . The payoffs are quasilinear in the outcome of G and the transfer. So, the payoff for i from a state θ , an action profile a , and a transfer y_i is $\pi_i(\theta, a) + y_i$. Observe, the specific labeling of \mathcal{M} does not impact the payoffs of (\mathcal{M}, G) .

Hierarchies of Beliefs This paper extends RM by introducing a larger domain of uncertainty that captures both payoff-relevant and payoff-irrelevant information. The expert’s first-order domain of uncertainty is $D_e^1 := \Lambda$ and the layman’s first-order domain of uncertainty is $D_\ell^1 := \Theta \times \Lambda$.⁷ So, only the expert knows the state, and no agent knows the label. Write H_i for the set of i ’s collectively coherent hierarchies of beliefs under these domains of uncertainty. Following Brandenburger and Dekel [1993], each hierarchy $h_i \in H_i$ has a canonical **extension** h_i^∞ in $\Delta(D_i^1 \times H_{-i})$ that describes the beliefs i has about D_i^1 and the hierarchies H_{-i} . Appendix A provides a formal construction of H_i and a description of each extension h_i^∞ .

2.1 Perfect Bayesian Equilibrium

Fix a mechanism \mathcal{M} and let \mathcal{T}_i be the set of i ’s terminal information sets in \mathcal{M} . An **interim belief mapping** for the expert is a function $\beta_e : \Theta \times \mathcal{I}_e \rightarrow \Delta(V)$ such that $\beta_e(\theta, I_e)(I_e) = 1$ for each $(\theta, I_e) \in \Theta \times \mathcal{I}_e$. So, β_e specifies e ’s beliefs about which node has been reached. Likewise, an **interim belief mapping** for the layman is a function $\beta_\ell : \mathcal{I}_\ell \rightarrow \Delta(\Theta \times V)$ such that $\beta_\ell(I_\ell)(\Theta \times I_\ell) = 1$ for each $I_\ell \in \mathcal{I}_\ell$. So, β_ℓ specifies ℓ ’s beliefs about which node and which state has been reached.

⁵This requirement implicitly assumes that i observes her transfer $\gamma_i(T_i)$.

⁶Observe, this paper differs from Rivera Mora [2025] in that here a mechanism includes a labeling.

⁷In RM, the layman’s first-order domain of uncertainty is $D_\ell^1 = \Theta$ and the expert’s first-order domain of uncertainty is a trivial singleton set.

Observe, interim belief mappings satisfy $\beta_e(\theta, T_e)(Z) = 1$ for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$ and that $\beta_\ell(T_\ell)(Z) = 1$ for each $T_\ell \in \mathcal{T}_\ell$. So, at each terminal node, both agents know that the mechanism has ended.

The interim belief mappings β induce a hierarchy $h_e \in H_e$ for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$ and a hierarchy $h_\ell \in H_\ell$ for each $T_\ell \in \mathcal{T}_\ell$. Write $\delta_e : \Theta \times \mathcal{T}_e \rightarrow H_e$ and $\delta_\ell : \mathcal{T}_\ell \rightarrow H_\ell$ for the **hierarchy mappings** induced by β . So, if the state is θ and the expert observes T_e , the expert has hierarchy $\delta_e(\theta, T_e)$. Likewise, if the layman observes T_ℓ , the layman has hierarchy $\delta_\ell(T_\ell)$. (See Appendix A.)

Fix a mechanism \mathcal{M} . The agents' optimization in (\mathcal{M}, G) is based on their payoffs at each information set $I_i \in \mathcal{I}_i$. A **behavior strategy** for the expert is a mapping σ_e from states Θ and information sets \mathcal{I}_e to distributions of actions available at \mathcal{I}_e . Likewise, a **behavior strategy** for the layman is a mapping σ_ℓ from information sets \mathcal{I}_ℓ to distributions of actions available at \mathcal{I}_ℓ . Each strategy profile (σ_e, σ_ℓ) specifies the agents' actions in G at each terminal node of \mathcal{M} . So, if $T_e \in \mathcal{T}_e$ and $T_\ell \in \mathcal{T}_\ell$, then $\sigma_e(\theta, T_e) \in \Delta(A_e)$ and $\sigma_\ell(T_\ell) \in \Delta(A_\ell)$. (See Appendix A for a full description of the agents' strategies.)

In principle, the interim belief mappings may not be consistent with the strategy profile played. The notion of perfect Bayesian equilibrium requires that they are. Write S_i for the set of pure strategies in \mathcal{M} . For each $(\theta, v) \in \Theta \times V$, write $P(\theta, v \mid s_i, \sigma_{-i})$ for the ex-ante probability that θ occurs and the path goes through v , provided that i plays the pure strategy s_i and $-i$ plays the behavioral strategy σ_{-i} . (See Appendix A.)

Definition 2.1. *The interim beliefs (β_e, β_ℓ) are **consistent** with (σ_e, σ_ℓ) if the following hold:*

(i) *For each $s_e \in S_e$, $(\theta, I_e) \in \Theta \times \mathcal{I}_e$, and $v \in I_e$,*

$$\beta_e(\theta, I_e)(v) \cdot \sum_{v' \in I_e} P(\theta, v' \mid s_e, \sigma_\ell) = P(\theta, v \mid s_e, \sigma_\ell).$$

(ii) *For each $s_\ell \in S_\ell$, $I_\ell \in \mathcal{I}_\ell$, and $(\theta, v) \in \Theta \times I_\ell$,*

$$\beta_\ell(I_\ell)(\theta, v) \cdot \sum_{(\theta', v') \in \Theta \times I_\ell} P(\theta', v' \mid s_\ell, \sigma_e) = P(\theta, v \mid s_\ell, \sigma_e).$$

Consistency requires that e 's (resp. ℓ 's) beliefs are derived by the chain rule of conditional probability.

Fix a strategy profile $\sigma = (\sigma_e, \sigma_\ell)$. Write $U_e(\sigma \mid \theta, I_e, \beta_e)$ for the expert's interim payoffs at θ and I_e under the profile σ and interim beliefs β_e . Likewise, write $U_\ell(\sigma \mid I_\ell, \beta_\ell)$ for the layman's interim payoffs at I_ℓ under the profile σ and interim beliefs β_ℓ . Importantly, interim payoffs include both, the expected transfers from the mechanism \mathcal{M} and the expected payoffs from the game G . (See Appendix A for its computation.)

Definition 2.2. *An assessment $(\sigma_e, \sigma_\ell, \beta_e, \beta_\ell)$ satisfies **sequential rationality** if the following hold:*

(i) *For each $(\theta, I_e) \in \Theta \times \mathcal{I}_e$ and each σ'_e , $U_e(\sigma_e, \sigma_\ell \mid \theta, I_e, \beta_e) \geq U_e(\sigma'_e, \sigma_\ell \mid \theta, I_e, \beta_e)$.*

(ii) For each $I_\ell \in \mathcal{I}_\ell$ and each σ'_ℓ , $U_\ell(\sigma_e, \sigma_\ell | I_\ell, \beta_\ell) \geq U_\ell(\sigma_e, \sigma'_\ell | I_\ell, \beta_\ell)$.

So, an assessment $(\sigma_e, \sigma_\ell, \beta_e, \beta_\ell)$ satisfies sequential rationality if the behavioral strategy σ_i maximizes i 's interim expected payoffs at each information set, provided that the agents face interim beliefs (β_e, β_ℓ) .

Definition 2.3. Call (σ, β) a *perfect Bayesian equilibrium (PBE)* if the assessment (σ, β) satisfies sequential rationality and the belief mappings β are consistent with σ .

3 Preliminaries

Before introducing the reduced-form approach, we present some preliminary properties that are essential for proving the existence result. Since the proof links Bayesian equilibria across different type structures, we will employ the pullback property and extension property of Bayesian games introduced by [Friedenberg and Meier \[2017\]](#) (henceforth **F&M**).⁸ While many type structures lack the extension property, **F&M** shows that this property holds for type structures that are mutually absolutely continuous.⁹ With this in mind, we construct a large type structure that satisfies mutual absolute continuity.

3.1 Hierarchies that admit a Common Prior

Definition 3.1. Fix finite sets $M_e \subsetneq H_e$ and $M_\ell \subsetneq H_\ell$. The set $M_e \times M_\ell$ *admits a common prior* if $M_e \times M_\ell$ is belief-closed and there is some probability measure $\phi \in \Delta(\Theta \times \Lambda \times M_e \times M_\ell)$ such that

- (i) for each $h_i \in M_i$, $\text{marg}_{M_i} \phi(h_i) > 0$, and
- (ii) for each $(\theta, \lambda, h_e, h_\ell) \in \Theta \times \Lambda \times M_e \times M_\ell$,

$$h_e^\infty(\lambda, h_\ell) \cdot \text{marg}_{\Theta \times M_e} \phi(\theta, h_e) = \phi(\theta, \lambda, h_e, h_\ell) = h_\ell^\infty(\theta, \lambda, h_e) \cdot \text{marg}_{M_\ell} \phi(h_\ell). \quad (1)$$

Say $h_i \in H_i$ *admits a common prior*, if there exist a finite set $M_e \subsetneq H_e$ and a finite set $M_\ell \subsetneq H_\ell$ such that $h_i \in M_i$ and $M_e \times M_\ell$ admits a common prior. Write $\text{CP}_i \subsetneq H_i$ for the set of hierarchies of i that admit a common prior. Lemma [C.2](#) shows that $\text{CP} := \text{CP}_e \times \text{CP}_\ell$ and satisfies the mutually absolute continuity property.

Since the set CP is belief closed, it induces a Θ -based type structure $\mathcal{S}_{\text{CP}} = (\text{CP}_e, \text{CP}_\ell, b_e, b_\ell)$. In this type structure, e 's set of types is $\Theta \times \text{CP}_e$ and the set of types for ℓ is CP_ℓ . The belief mappings $b_e : \Theta \times \text{CP}_e \rightarrow \Delta(\text{CP}_\ell)$ and $b_\ell : \text{CP}_\ell \rightarrow \Delta(\Theta \times \text{CP}_e)$ are defined as

$$b_e(\theta, h_e)(h_\ell) := \text{marg}_{H_\ell} h_e^\infty(h_\ell) \quad \text{and} \quad b_\ell(h_\ell)(\theta, h_e) := \text{marg}_{\Theta \times H_e} h_\ell^\infty(\theta, h_e).$$

⁸The results in **F&M** use a framework with a common domain of uncertainty. However, their results extend (mutatis mutandis) to the present case in which only e knows Θ .

⁹For instance, **F&M** shows that the canonical type structure fails the extension property.

So, the mapping b_e describes e 's beliefs about ℓ 's type and the mapping b_ℓ describes ℓ 's beliefs about the state and e 's type.

Write $\text{BG}(\mathcal{S}_{\text{CP}})$ for the Bayesian game induced by G and the type structure \mathcal{S}_{CP} . In this Bayesian game, the **expert's strategy** is a mapping $\varsigma_e : \Theta \times \text{CP}_e \rightarrow \Delta(A_e)$ and the **layman's strategy** is a mapping $\varsigma_\ell : \text{CP}_\ell \rightarrow \Delta(A_\ell)$. The agents' expected payoffs given a profile $(\varsigma_e, \varsigma_\ell)$ are

$$\begin{aligned}\Pi_e(\varsigma_e, \varsigma_\ell \mid \theta, h_e) &:= \int_{H_\ell} \int_{A_e} \int_{A_\ell} \pi_e(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e) d\varsigma_\ell(h_\ell) db_e(\theta, h_e), \quad \text{and} \\ \Pi_\ell(\varsigma_e, \varsigma_\ell \mid h_\ell) &:= \int_{\Theta \times H_e} \int_{A_e} \int_{A_\ell} \pi_e(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e) d\varsigma_\ell(h_\ell) db_\ell(h_\ell).\end{aligned}$$

Call a game G **nice** if the set of Bayesian equilibria of the Bayesian game $\text{BG}(\mathcal{S}_{\text{CP}})$ is non-empty. Lemma C.4 shows that the class of nice games include all games with finite actions sets.¹⁰

3.2 Full-Path Equilibria

The pull-back and the extension property over the type structure \mathcal{S}_{CP} will serve as the key tool in the analysis. With this in mind, we will focus on analyzing PBE that solely induce hierarchies consistent with a common prior.

Fix a mechanism \mathcal{M} and let Z be the terminal nodes of \mathcal{M} . Let $\sigma = (\sigma_e, \sigma_\ell)$ be a strategy profile for the supergame (\mathcal{M}, G) . The profile σ defines a probability distribution $P(\cdot \mid \sigma) \in \Delta Z$ that describes the probability of each terminal node given that the agents play according to σ and the prior of Θ is μ . (See Appendix A.)

Definition 3.2. Fix a mechanism \mathcal{M} and a consistent assessment (σ, β) of (\mathcal{M}, G) . Let Z be the set of terminal nodes of \mathcal{M} . Say (σ, β) has **full path** if $P(z \mid \sigma) > 0$ for each terminal node $z \in Z$.

An assessment (σ, β) has full path if every terminal node is reached with strictly positive probability under the strategy profile. Let $\text{cons}(\mathcal{M})$ denote the set of belief mappings that are consistent with some strategy profile. We say that $\beta \in \text{cons}(\mathcal{M})$ has **full path** if there exists a strategy profile σ such that β is consistent with σ and (σ, β) has full path.

Lemma 3.1. Fix a mechanism \mathcal{M} and $\beta \in \text{cons}(\mathcal{M})$. Let (δ_e, δ_ℓ) be the hierarchy mappings induced by β . Write $M_i \subseteq H_i$ for the range of δ_i . If β has full path, then:

- (i) the mapping $\delta_e : \Theta \times \mathcal{T}_e \rightarrow M_e$ is constant in Θ ,
- (ii) for each $\theta \in \Theta$, the mapping $\delta_e(\theta, \cdot) : \mathcal{T}_e \rightarrow M_e$ is bijective,
- (iii) the mapping $\delta_\ell : \mathcal{T}_\ell \rightarrow M_\ell$ is bijective, and
- (iv) $M_e \times M_\ell$ admits a common prior.

Lemma 3.1 establishes two fundamental properties of full-path equilibria in our framework with extended uncertainty. First, each terminal information set is associated with a unique hierarchy, thereby circumventing the challenges highlighted in Example 1.1. Second, it ensures that all induced

¹⁰Some nice games have infinite action sets. For example, see the game in Section 6.3 of RM.

hierarchies are consistent with a common prior.¹¹ (These properties enable the application of both the pull-back and extension property). Importantly, as shown in Lemma C.5, focusing on full-path equilibria entails no loss of generality.

3.3 The Pull-back and the Extension Property

Write $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ for a mechanism in which the sets of terminal information sets are \mathcal{T}_e and \mathcal{T}_ℓ . Observe, interim beliefs are only important if they arise from the agents' updating in the mechanism.

Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and interim belief mappings $\beta \in \text{cons}(\mathcal{M})$. After the agents finish playing \mathcal{M} , each agent i learns information associated with a terminal information set. Notice, the realized terminal information sets T_e and T_ℓ may not be singletons. So, the expert knows (θ, T_e) but may not know T_ℓ and the layman knows T_ℓ but may not know (θ, T_e) . Thus, we can think of these information sets as reflecting types of the agents. Formally, the expert's set of types is $\Theta \times \mathcal{T}_e$ and the layman's set of types is \mathcal{T}_ℓ . Since \mathcal{M} has an observable end, each profile $(\theta, T_e, T_\ell) \in \Theta \times \mathcal{T}_e \times \mathcal{T}_\ell$ satisfies $\beta_e(\theta, T_e)(Z) = \beta_\ell(T_\ell)(\Theta \times Z) = 1$. Hence, the probability that an expert of type (θ, T_e) assigns to T_ℓ is $\beta_e(\theta, T_e)(T_\ell)$ and the probability that a layman of type T_ℓ assigns to (θ, T_e) is $\beta_\ell(T_\ell)(\{\theta\} \times T_e)$.¹²

Write $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ for the **Bayesian game** induced by $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and β . In this Bayesian game, a strategy for the expert is a mapping $\hat{\sigma}_e : \Theta \times \mathcal{T}_e \rightarrow \Delta(A_e)$ and a strategy for the layman is a mapping $\hat{\sigma}_\ell : \mathcal{T}_\ell \rightarrow \Delta(A_\ell)$. The expected payoffs given a profile $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ are

$$\begin{aligned} \hat{\Pi}_e(\hat{\sigma}_e, \hat{\sigma}_\ell \mid \theta, T_e) &:= \sum_{T_\ell \in \mathcal{T}_\ell} \beta_e(\theta, T_e)(T_\ell) \int_{A_e \times A_\ell} \pi_e(\theta, a_e, a_\ell) d\hat{\sigma}_e(\theta, T_e)(a_e) d\hat{\sigma}_\ell(T_\ell)(a_\ell) \quad \text{and} \\ \hat{\Pi}_\ell(\hat{\sigma}_e, \hat{\sigma}_\ell \mid T_\ell) &:= \sum_{(\theta, T_e) \in \Theta \times \mathcal{T}_e} \beta_\ell(T_\ell)(\{\theta\} \times T_e) \int_{A_e \times A_\ell} \pi_\ell(\theta, a_e, a_\ell) d\hat{\sigma}_e(\theta, T_e)(a_e) d\hat{\sigma}_\ell(T_\ell)(a_\ell) \end{aligned}$$

Lemma 3.1 implies that if (σ, β) has full path, then the type structure in $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ can be *embedded* into the type structure $\text{BG}(\mathcal{S}_{\text{CP}})$ via the (injective) type morphisms $(\delta_e(\theta, \cdot))_{\theta \in \Theta}$ and $\delta_\ell(\cdot)$ in a way that preserves the agents' hierarchies of beliefs. (See Definition 5 in F&M).

Definition 3.3. Fix a mechanism \mathcal{M} and a consistent assessment (σ, β) that has full path on \mathcal{M} . Write (δ_e, δ_ℓ) for the hierarchy mappings induced by β . A Bayesian equilibrium $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ can be **extended** to a Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of $\text{BG}(\mathcal{S}_{\text{CP}})$ if

- (i) for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, $\hat{\sigma}_e(\theta, T_e) = \varsigma_e(\theta, \delta_e(\theta, T_e))$, and
- (ii) for each $T_\ell \in \mathcal{T}_\ell$, $\hat{\sigma}_\ell(T_\ell) = \varsigma_\ell(\delta_\ell(T_\ell))$.

The following results state the pull-back and the extension property in the context of this paper:

¹¹This property may fail for equilibria without full path, as terminal information sets unreached by the strategy profile can induce hierarchies inconsistent with a common prior.

¹²Formally, in a Bayesian game the agents priors are mappings $\hat{\beta}_e : \Theta \times \mathcal{T}_e \rightarrow \Delta(\mathcal{T}_\ell)$ and $\hat{\beta}_\ell : \mathcal{T}_\ell \rightarrow \Delta(\Theta \times \mathcal{T}_e)$. For notational simplicity, I use the notation of (β_e, β_ℓ) instead of the correct notation $(\hat{\beta}_e, \hat{\beta}_\ell)$.

Lemma 3.2. (The pull-back property) Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and $\beta \in \text{cons}(\mathcal{M})$ with full path. If $(\varsigma_e, \varsigma_\ell)$ is a Bayesian equilibrium of $\text{BG}(\mathcal{S}_{\text{CP}})$, then there is some Bayesian equilibrium $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ that can be extended to $(\varsigma_e, \varsigma_\ell)$.

Lemma 3.3. (The extension property) Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and $\beta \in \text{cons}(\mathcal{M})$ with full path. If $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ is a Bayesian equilibrium of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$, and G is nice, then $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ can be extended to some Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of $\text{BG}(\mathcal{S}_{\text{CP}})$.

These properties establish a key connection between equilibrium behavior across different type structures: the “smaller” type structures defined for each of the induced Bayesian games, and the “larger” type structure \mathcal{S}_{CP} . As a result, this correspondence enables the analysis of equilibrium payoffs across hierarchies that admit a common prior, irrespective of the specific induced Bayesian game from which these hierarchies arise.

4 The Reduced-Form Approach

This section introduces a reduced-form approach to characterize agents’ preferences for information. The key idea is to transform the original game into a simpler psychological game where agents have preferences over belief hierarchies.

4.1 Psychological Games

The reduced approach uses auxiliary psychological games to describe the agents preferences for information. In a psychological game, the agents have preferences over their hierarchies of beliefs. A **belief-dependent utility** for the expert is a measurable function $u_e : \Theta \times H_e \rightarrow \mathbb{R}$. Similarly, a **belief-dependent utility** for the layman is a measurable function $u_\ell : H_\ell \rightarrow \mathbb{R}$.¹³

The mechanism \mathcal{M} and belief-dependent utilities (u_e, u_ℓ) induce the **psychological game** $(\mathcal{M}, u_e, u_\ell)$. The timing is as follows: First, Nature chooses state $\theta \in \Theta$, which is then observed by the expert. Next, the agents play \mathcal{M} , after which each agent i observes a terminal information set $T_i \in \mathcal{T}_i$, thereby obtaining a transfer $y_i = \gamma_i(T_i)$. The payoffs of each agent i are quasilinear in the belief-dependent sub-utility u_i and the transfer y_i . So, if the expert observes state θ , has hierarchy h_e , and transfer y_e , the expert’s utility is $u_e(\theta, h_e) + y_e$. If the layman has hierarchy h_ℓ and transfer y_ℓ , the layman’s utility is $u_\ell(h_\ell) + y_\ell$.

A **behavior strategy** for the expert is a mapping ρ_e from states Θ and information sets $\mathcal{I}_e \setminus \mathcal{T}_e$ to distributions of actions available at $\mathcal{I}_e \setminus \mathcal{T}_e$. Likewise, a **behavior strategy** for the layman is a mapping ρ_ℓ from information sets $\mathcal{I}_\ell \setminus \mathcal{T}_\ell$ to distributions of actions available at $\mathcal{I}_\ell \setminus \mathcal{T}_\ell$. (See Appendix A.) For each $(\theta, v) \in \Theta \times V$, write $P(\theta, v \mid s_i, \rho_{-i})$ for the ex-ante probability that θ occurs and the path goes through v , provided that i plays the pure strategy s_i and $-i$ plays the behavioral strategy ρ_{-i} . (See Appendix A.) Write $\mathcal{U}_e(\rho \mid \theta, I_e, \beta)$ for the expert’s interim payoffs at (θ, I_e) , given the interim beliefs β . Lastly write $\mathcal{U}_\ell(\rho \mid I_\ell, \beta)$ for the layman’s interim payoffs at I_ℓ .

¹³Since the layman does not observe the state, the layman’s belief-dependent utility does not directly depend on Θ ; instead it depends on his first-order beliefs h_ℓ^1 about Θ .

Call (ρ, β) a **perfect Bayesian equilibrium (PBE)** if the assessment (ρ, β) satisfies sequential rationality and the belief mappings β are consistent with ρ . This solution concept requires:

Definition 4.1. An assessment $(\rho_e, \rho_\ell, \beta_e, \beta_\ell)$ satisfies **sequential rationality** if the following hold:

- (i) For each $(\theta, I_e) \in \Theta \times \mathcal{I}_e$ and ρ'_e , $\mathcal{U}_e(\rho_e, \rho_\ell \mid \theta, I_e, \beta_e) \geq \mathcal{U}_e(\rho'_e, \rho_\ell \mid \theta, I_e, \beta_e)$.
- (ii) For each $I_\ell \in \mathcal{I}_\ell$ and ρ'_ℓ , $\mathcal{U}_\ell(\rho_e, \rho_\ell \mid I_\ell, \beta_\ell) \geq \mathcal{U}_\ell(\rho_e, \rho'_\ell \mid I_\ell, \beta_\ell)$.

Definition 4.2. The interim beliefs (β_e, β_ℓ) are **consistent** with (ρ_e, ρ_ℓ) if the following hold:

- (i) For each $s_e \in S_e$, $(\theta, I_e) \in \Theta \times \mathcal{I}_e$, and $v \in I_e$,

$$\beta_e(\theta, I_e)(v) \cdot \sum_{v' \in I_e} P(\theta, v' \mid s_e, \rho_\ell) = P(\theta, v \mid s_e, \rho_\ell).$$

- (ii) For each $s_\ell \in S_\ell$, $I_\ell \in \mathcal{I}_\ell$, and $(\theta, v) \in \Theta \times I_\ell$,

$$\beta_\ell(I_\ell)(\theta, v) \cdot \sum_{(\theta', v') \in \Theta \times I_\ell} P(\theta', v' \mid s_\ell, \rho_e) = P(\theta, v \mid s_\ell, \rho_e).$$

4.2 Reduced forms

Definition 4.3. Fix $u_e : \Theta \times H_e \rightarrow \mathbb{R}$ and $u_\ell : H_\ell \rightarrow \mathbb{R}$. Call (u_e, u_ℓ) a **full reduced form** for G if, for each mechanism \mathcal{M} , and each $\beta \in \text{cons}(\mathcal{M})$ with full path, there is some Bayesian equilibrium $\hat{\sigma}$ of the associated Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ such that:

- (i) for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, $u_e(\theta, \delta_e(\theta, T_e)) = \hat{\Pi}_e(\hat{\sigma} \mid \theta, T_e, \beta_e)$, and
- (ii) for each $T_\ell \in \mathcal{T}_\ell$, $u_\ell(\delta_\ell(T_\ell)) = \hat{\Pi}_\ell(\hat{\sigma} \mid T_\ell, \beta_\ell)$.

A pair of belief-dependent utility functions (u_e, u_ℓ) is a full reduced form for G if it captures equilibrium payoffs of each induced Bayesian game (with full path), by making reference to the agents' hierarchies of beliefs. Hence, these reduced forms allow us to analyze mechanisms using belief-dependent utilities rather than explicitly modeling the after-game. Observe that the definition of reduced forms introduced in [RM](#) is stronger than [Definition 4.3](#). While the former requires capturing the payoffs of all PBE of each mechanism \mathcal{M} , the latter only captures PBE with *full path*.¹⁴ Nevertheless, [Lemma C.5](#) shows that restricting the analysis to full-path PBE does not result in any loss of generality. Hence, there is no loss to consider full reduced-forms.

We now link the Bayesian game $\text{BG}(\mathcal{S}_{\text{CP}})$ with the full reduced-forms of G . Fix a Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of $\text{BG}(\mathcal{S}_{\text{CP}})$ and belief-based utility functions $u_e : \Theta \times H_e \rightarrow \mathbb{R}$ and $u_\ell : H_\ell \rightarrow \mathbb{R}$. Say that the pair (u_e, u_ℓ) is **consistent** with $(\varsigma_e, \varsigma_\ell)$ if for each $(\theta, h_e) \in \Theta \times \text{CP}_e$ and each $h_\ell \in \text{CP}_\ell$,

$$u_e(\theta, h_e) = \Pi_e(\varsigma_e, \varsigma_\ell \mid \theta, h_e) \quad \text{and} \quad u_\ell(h_\ell) = \Pi_\ell(\varsigma_e, \varsigma_\ell \mid h_\ell) \quad (2)$$

The pull-back property ([Lemma 3.2](#)) implies the following result:

¹⁴Observe, if a PBE does not have full path, then the beliefs outside of the path of play are arbitrary and hence, may not be consistent with a common prior.

Lemma 4.1. *Fix a BE $(\varsigma_e, \varsigma_\ell)$ of the Bayesian game $\text{BG}(\mathcal{S}_{\text{CP}})$. If $u_e : \Theta \times H_e \rightarrow \mathbb{R}$ and $u_\ell : H_\ell \rightarrow \mathbb{R}$ are consistent with $(\varsigma_e, \varsigma_\ell)$, then (u_e, u_ℓ) is a full reduced form of G .*

Proof. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and some $\beta \in \text{cons}(\mathcal{M})$. Write δ for the hierarchy mappings associated with β .

The Pull-back Property (Lemma 3.2) states that there is some Bayesian equilibrium $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ that can be extended to $(\varsigma_e, \varsigma_\ell)$. So, each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, satisfies

$$u_e(\theta, \delta_e(\theta, T_e)) = \Pi_e(\varsigma_e, \varsigma_\ell \mid \theta, \delta_e(\theta, T_e)) = \Pi_e(\hat{\sigma}_e, \hat{\sigma}_\ell \mid \theta, T_e, \beta_e),$$

where the first equality follows from the fact that u_e is consistent with ς_e and the second from the fact that $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ can be extended to $(\varsigma_e, \varsigma_\ell)$. Likewise, each $T_\ell \in \mathcal{T}_\ell$ satisfies

$$u_\ell(\delta_\ell(T_\ell)) = \Pi_\ell(\varsigma_e, \varsigma_\ell \mid \delta_\ell(T_\ell)) = \Pi_\ell(\hat{\sigma}_e, \hat{\sigma}_\ell \mid T_\ell, \beta_\ell).$$

Therefore, (u_e, u_ℓ) is a full reduced form for G . ■

Lemma 4.1 establishes that each Bayesian equilibrium of $\text{BG}(\mathcal{S}_{\text{CP}})$ generates a full reduced form for G . More precisely, any pair of mappings satisfying Equation (2) constitutes a full reduced form. As a consequence, each nice game G possesses at least one full reduced form:

Corollary 4.1. *If G is nice, then G has some full reduced form.*

4.3 Reduced-form representations

When each induced Bayesian game has a unique Bayesian equilibrium, a single reduced form is sufficient to capture all equilibrium payoffs. However, if some induced Bayesian games have multiple equilibria, a single reduced form may not capture all equilibrium payoffs. To fully characterize the set of equilibrium payoffs, we must consider a collection of reduced forms. This collection is referred to as a reduced-form representation and ensures that all possible equilibrium payoffs are accounted for.

To formally define a reduced-form representation, first we define equivalence of PBE of supergames and psychological games. Write $Y_e(\theta, \mathcal{M}, \sigma)$ for the expected transfer of an expert that observes θ in a mechanism \mathcal{M} under profile σ . Likewise, write $Y_\ell(\mathcal{M}, \sigma)$ for the layman's expected transfer in \mathcal{M} under profile σ .

Definition 4.4. *A PBE (σ, β) of a supergame (\mathcal{M}, G) and a PBE (ρ, β') of a psychological game $(\mathcal{M}', u_e, u_\ell)$ are **equivalent** if the following hold:*

- (i) *For each $\theta \in \Theta$, $U_e(\sigma \mid \theta, \beta_e) = U_e(\rho \mid \theta, \beta')$ and $Y_e(\theta, \mathcal{M}, \sigma) = Y_e(\theta, \mathcal{M}', \rho)$.*
- (ii) *$U_\ell(\sigma \mid \beta_\ell) = U_\ell(\rho \mid \beta')$ and $Y_\ell(\mathcal{M}, \sigma) = Y_\ell(\mathcal{M}', \rho)$.*

So, two PBE are equivalent if they induce the same expected payoffs and the same expected transfers for each agent, and each state observed by the expert.

Definition 4.5. Let RF be a set of full reduced forms for G . Say RF is a **full reduced-form representation** for G if, for each mechanism \mathcal{M} , and each PBE (σ, β) of (\mathcal{M}, G) , there is some mechanism \mathcal{M}' , some $(u_e, u_\ell) \in \text{RF}$, and some PBE of $(\mathcal{M}', (u_e, u_\ell))$ that is equivalent to (σ, β) .

Definition 4.5 formalizes when a set of reduced forms fully captures all equilibrium behavior. A set RF is a full reduced-form representation if for any equilibrium in any mechanism, there exists some reduced form in RF that can replicate the same equilibrium outcomes and transfers.

Theorem 4.1. *If G is nice, then G has a full reduced-form representation.*

4.4 Proof

We introduce some notation before explaining the proof. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$, interim belief mappings $\beta \in \text{cons}(\mathcal{M})$, and belief-based utility functions (u_e, u_ℓ) . A strategy profile (σ_e, σ_ℓ) of the supergame (\mathcal{M}, G) induces the strategy profile $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ of the Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ if the behavioral strategy $\hat{\sigma}_e$ (resp. $\hat{\sigma}_\ell$) is the restriction of σ_e (resp. σ_ℓ) to $\Theta \times \mathcal{T}_e$ (resp. \mathcal{T}_ℓ). Similarly, (σ_e, σ_ℓ) induces the strategy profile (ρ_e, ρ_ℓ) of the psychological game $(\mathcal{M}, u_e, u_\ell)$ if ρ_e (resp. ρ_ℓ) is the restriction of σ_e (resp. σ_ℓ) to $\Theta \times (\mathcal{I}_e \setminus \mathcal{T}_e)$ (resp. $\mathcal{I}_\ell \setminus \mathcal{T}_\ell$).

Proof of Theorem 4.1. Since G is nice, the set of BE of $\text{BG}(\mathcal{S}_{\text{CP}})$ is non-empty. Write RF for the set of belief-based utilities (u_e, u_ℓ) that are consistent with some BE of $\text{BG}(\mathcal{S}_{\text{CP}})$. By Lemma 4.1, each $(u_e, u_\ell) \in \text{RF}$ is a full reduced form for G .

We now show that RF is a full reduced-form representation for G . Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$, and a PBE (σ, β) of the supergame (\mathcal{M}, G) . By Lemma C.5, there is no loss to assume that (σ, β) has full path on \mathcal{M} . It suffices to show that there is some $(u_e, u_\ell) \in \text{RF}$, and a PBE (ρ, β) of $(\mathcal{M}, u_e, u_\ell)$ that is equivalent to (σ, β) .

Write $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ for the BE of the induced Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ that is induced by (σ, β) . By the Extension Property (Lemma 3.3), the profile $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ can be extended to some Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of $\text{BG}(\mathcal{S}_{\text{CP}})$. Let (u_e, u_ℓ) be belief-dependent utility functions that are consistent with $(\varsigma_e, \varsigma_\ell)$ and write $\rho = (\rho_e, \rho_\ell)$ for the restriction of (σ_e, σ_ℓ) on the psychological game $(\mathcal{M}, u_e, u_\ell)$.

It suffices to show that (σ, β) is equivalent to (ρ, β) . Observe, for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$,

$$\hat{\Pi}_e(\hat{\sigma}_e, \hat{\sigma}_\ell \mid \theta, T_e) = \Pi_e(\varsigma_e, \varsigma_\ell \mid \theta, \delta_e(\theta, T_e)) = u_e(\theta, \delta_e(\theta, T_e)), \quad (3)$$

where the first equality follows from the fact that (σ_e, σ_ℓ) can be extended to $(\varsigma_e, \varsigma_\ell)$ and the second from the fact that (u_e, u_ℓ) are consistent with $(\varsigma_e, \varsigma_\ell)$.

Likewise, for each $T_\ell \in \mathcal{T}_\ell$,

$$\hat{\Pi}_\ell(\hat{\sigma}_e, \hat{\sigma}_\ell \mid T_\ell) = \Pi_\ell(\varsigma_e, \varsigma_\ell \mid \delta_\ell(T_\ell)) = u_\ell(\theta, \delta_\ell(h_\ell)). \quad (4)$$

We use equations (3) and (4) to show that (ρ, β) is a PBE of the psychological game $(\mathcal{M}, u_e, u_\ell)$.

Observe, since (σ, β) is a PBE of (\mathcal{M}, G) , the beliefs β are consistent with σ . Since ρ is the restriction of σ to $(\mathcal{M}, u_e, u_\ell)$, the beliefs β are also consistent with ρ .

We first show that (ρ, β) satisfies sequential rationality for the expert. Fix some $(\theta, I_e) \in \Theta \times \mathcal{I}_e$ and some strategy ρ'_e . Observe that

$$\begin{aligned} \mathcal{U}_e(\rho_e, \rho_\ell \mid \theta, I_e, \beta) &= \sum_{T_e \in \mathcal{T}_e} P(T_e \mid \theta, I_e, \rho, \beta_e) \cdot [u_e(\theta, \delta_e(\theta, T_e)) + \gamma_e(T_e)] \\ &= \sum_{T_e \in \mathcal{T}_e} P(T_e \mid \theta, I_e, \rho, \beta_e) \cdot [\Pi_e(\hat{\sigma} \mid \theta, T_e) + \gamma_e(T_e)] \\ &\geq \sum_{T_e \in \mathcal{T}_e} P(T_e \mid \theta, I_e, (\rho'_e, \rho_\ell), \beta_e) \cdot [\Pi_e(\hat{\sigma} \mid \theta, T_e) + \gamma_e(T_e)] \\ &= \mathcal{U}_e(\rho'_e, \rho_\ell \mid \theta, I_e, \beta), \end{aligned}$$

where the second equality follows from Equation (3), and the inequality follows from the fact that (σ, β) is a PBE of (\mathcal{M}, G) and that $\sigma = (\sigma_e, \sigma_\ell)$ induces $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ and (ρ_e, ρ_ℓ) .

Now we show sequential rationality for the layman. Fix any $I_\ell \in \mathcal{I}_\ell$ and alternative strategy ρ'_ℓ . An analogous argument shows that

$$\mathcal{U}_\ell(\rho \mid I_\ell, \beta) \geq \mathcal{U}_\ell(\rho_e, \rho'_\ell \mid I_\ell, \beta).$$

Therefore, (ρ, β) is a PBE of $(\mathcal{M}, u_e, u_\ell)$. Moreover, Equations (3) and (4) imply that (ρ, β) is equivalent to (σ, β) . ■

References

- Robert J Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society*, pages 1–18, 1987.
- Pierpaolo Battigalli and Martin Dufwenberg. Dynamic psychological games. *Journal of Economic Theory*, 144(1):1–35, 2009.
- Pierpaolo Battigalli, Roberto Corrao, and Martin Dufwenberg. Incorporating belief-dependent motivation in games. *Journal of Economic Behavior & Organization*, 167:185–218, 2019.
- Adam Brandenburger and Eddie Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189–198, 1993.
- Giacomo Calzolari and Alessandro Pavan. Monopoly with resale. *The RAND Journal of Economics*, 37(2):362–375, 2006a.
- Giacomo Calzolari and Alessandro Pavan. On the optimality of privacy in sequential contracting. *Journal of Economic theory*, 130(1):168–204, 2006b.

- Hans Carlsson and Eric Van Damme. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, pages 989–1018, 1993.
- Roberto Corrao and Yifan Dai. The bounds of mediated communication. *arXiv preprint arXiv:2303.06244*, 2023.
- Laura Doval and Vasiliki Skreta. Mechanism design with limited commitment. *Econometrica*, 90(4):1463–1500, 2022.
- Laura Doval and Vasiliki Skreta. Optimal mechanism for the sale of a durable good. *Theoretical Economics*, 19(2):865–915, 2024.
- Piotr Dworzak. Mechanism design with aftermarket: Cutoff mechanisms. *Econometrica*, 88(6):2629–2661, 2020.
- Jeffrey C Ely and Marcin Peski. Hierarchies of belief and interim rationalizability. *Theoretical Economics* 1, (1):19–65, 2006.
- Amanda Friedenberg and Martin Meier. The context of the game. *Economic Theory*, 63(2):347–386, 2017.
- John Geanakoplos, David Pearce, and Ennio Stacchetti. Psychological games and sequential rationality. *Games and economic Behavior*, 1(1):60–79, 1989.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Qingmin Liu. On redundant types and bayesian formulation of incomplete information. *Journal of Economic Theory*, 144(5):2115–2145, 2009.
- Laurent Mathevet, Jacopo Perego, and Ina Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020.
- Andrew McClellan and Daniel Rappoport. Signaling good faith by taking stands. *Available at SSRN 4510675*, 2024.
- Stephen Morris. Political correctness. *Journal of political Economy*, 109(2):231–265, 2001.
- Stephen Morris and Hyun Song Shin. Global games: Theory and applications. 2001.
- Marco Ottaviani and Peter Norman Sørensen. Reputational cheap talk. *The Rand journal of economics*, 37(1):155–175, 2006.
- Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5):949–987, 2010.

Ernesto Rivera Mora. Mechanism design with belief-dependent preferences. *Journal of Economic Theory*, 216:105782, 2024.

Ernesto Rivera Mora. Neutral mechanisms: On the feasibility of information sharing. Working Paper, 2025.

Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge". *The American Economic Review*, pages 385–391, 1989.

Ina Taneva. Information design. *American Economic Journal: Microeconomics*, 11(4):151–85, 2019.

A Appendix: Extra Definitions

Hierarchies of beliefs Fix a first-order domain of uncertainty D_i^1 for each agent i . The set of first-order beliefs of agent i is $H_i^1 := \Delta(D_i^1)$. Inductively define the sets D_i^k and H_i^k as follows: Assume the sets D_i^k and H_i^k are defined for k . Write $D_i^{k+1} := D_i^k \times H_{-i}^k$ for the $(k+1)$ -order domain of uncertainty of agent i . Write

$$H_i^{k+1} := \left\{ (h_i^1, \dots, h_i^{k+1}) \in H_i^k \times \Delta(D_i^{k+1}) : \text{marg}_{D_i^k} h_i^{k+1} = h_i^k \right\}$$

for the set of collectively coherent $(k+1)$ -order beliefs of agent i . Note that, if $(h_i^1, \dots, h_i^{k+1}) \in H_i^{k+1}$, then $(h_i^1, \dots, h_i^n) \in H_i^n$ for all $n \leq k$; that is, each $(h_i^1, \dots, h_i^{k+1}) \in H_i^{k+1}$ is coherent. Write

$$H_i = \left\{ (h_i^1, h_i^2, \dots) \in \prod_{k=1}^{\infty} \Delta(D_i^k) : (h_i^1, \dots, h_i^k) \in H_i^k \text{ for each } k \in \mathbb{N} \right\},$$

for the set of i 's **collectively coherent hierarchies of beliefs**. So, $h_i = (h_i^1, h_i^2, \dots) \in H_i$ is a particular hierarchy of beliefs for agent i . Call $H := H_e \times H_\ell$ the **belief structure**.

Following [Brandenburger and Dekel \[1993\]](#), there is a canonical homeomorphism between the spaces H_i and $\Delta(D_i \times H_{-i})$. So, for each hierarchy $h_i \in H_i$, write $h_i^\infty \in \Delta(D_i \times H_{-i})$ for the image of h_i under the canonical homeomorphism. The extension h_i^∞ describes i 's beliefs about both the fundamental uncertainty D_i and the hierarchies of the other agent H_{-i} .

From interim beliefs to hierarchies of beliefs Fix a mechanism \mathcal{M} and interim belief mappings $\beta_e : \Theta \times \mathcal{I}_e \rightarrow \Delta(V)$ and $\beta_\ell : \mathcal{I}_\ell \rightarrow \Delta(\Theta \times V)$. The mappings β_e and β_ℓ induce terminal belief mappings $\hat{\beta}_e : \Theta \times \mathcal{T}_e \rightarrow \Delta(Z)$ and $\hat{\beta}_\ell : \mathcal{T}_\ell \rightarrow \Delta(\Theta \times Z)$. (Recall that for each $(\theta, T_e, T_\ell) \in \Theta \times \mathcal{T}_e \times \mathcal{T}_\ell$, $\beta_e(\theta, T_e)(Z) = 1$ and $\beta_\ell(T_\ell)(\Theta \times Z) = 1$.)

Define mappings $\varphi_e^1 : Z \rightarrow D_e^1$ and $\varphi_\ell^1 : \Theta \times Z \rightarrow D_\ell^1$, so that $\varphi_e^1(z) = L(z) \in D_e^1$ and $\varphi_\ell^1(\theta, z) = (\theta, L(z)) \in D_\ell^1$.¹⁵ Note that φ_i^1 is measurable for each $i \in \{e, \ell\}$. Assume that the measurable

¹⁵Note that [RM](#) uses domain of uncertainty $D_e^1 = \{\triangleleft\}$ and $D_\ell^1 = \Theta$ with $\varphi_e^1(z) = \triangleleft$ and $\varphi_\ell^1(\theta, z) = \theta$.

maps φ_e^k and φ_ℓ^k are defined. Let $\varphi_e^{k+1} : Z \rightarrow D_e^{k+1}$ be defined by $\varphi_e^{k+1}(z) = (\varphi_e^k(z), \varphi_\ell^k(\hat{\beta}_\ell(T_\ell)))$ where $T_\ell \in \mathcal{T}_\ell$ is the unique terminal information set of ℓ such that $z \in T_\ell$. Likewise, let $\varphi_\ell^{k+1} : \Theta \times \mathcal{T}_e \rightarrow D_\ell^{k+1}$ be defined by $\varphi_\ell^{k+1}(\theta, z) = (\varphi_\ell^k(\theta, z), \varphi_e^k(\hat{\beta}_e(\theta, T_e)))$, where $T_e \in \mathcal{T}_e$ is the unique terminal information set of e such that $z \in T_e$. Note, since Θ and Z are both finite, the mappings φ_i^{k+1} are measurable for each $k \in \mathbb{N}$.

Set $\delta_i^k := \varphi_i^k \circ \hat{\beta}_i$. Note that for each T_ℓ , $\text{proj}_{D_e^k} \varphi_e^{k+1}(T_\ell) = \varphi_e^k(T_\ell)$. Thus, for each (θ, T_e) , $\text{marg}_{D_e^k} \delta_e^{k+1}(\theta, T_e) = \delta_e^k(\theta, T_e)$. Similarly, for each (θ, T_e) , $\text{proj}_{D_\ell^k} \varphi_\ell^{k+1}(\theta, T_e) = \varphi_\ell^k(\theta, T_e)$. Consequently, for each T_ℓ , $\text{marg}_{D_\ell^k} \delta_\ell^{k+1}(T_\ell) = \delta_\ell^k(T_\ell)$. Finally, write $\delta_e : \Theta \times \mathcal{T}_e \rightarrow H_e$ and $\delta_\ell : \mathcal{T}_\ell \rightarrow H_\ell$ for $\delta_e(\theta, T_e) = (\delta_e^1(\theta, T_e), \delta_e^2(\theta, T_e), \dots)$ and $\delta_\ell(T_\ell) = (\delta_\ell^1(T_\ell), \delta_\ell^2(T_\ell), \dots)$.

Neutral mechanisms Each mechanism \mathcal{M} has a set of actions X . Each agent $i \in \{e, \ell, c\}$ has an action correspondence $\mathcal{A}_i : \mathcal{I}_i \setminus \mathcal{T}_i \rightrightarrows X$ that specifies the actions that are available at each non-terminal information set. Chance's behavior is described by a exogenous behavioral strategy $\rho_c \in \prod_{I_c \in \mathcal{I}_c \setminus \mathcal{T}_c} \Delta(\mathcal{A}_c(I_c))$. Note, since the set of actions at a given information set can be a singleton, the definition does not require simultaneous moves. Observe, no absentmindedness implies that each $i \in \{e, \ell, c\}$ knows the start of the game, i.e., $\{\emptyset\} \in \mathcal{I}_i$.

Write $S_i := \prod_{I_i \in \mathcal{I}_i \setminus \mathcal{T}_i} \mathcal{A}_i(I_i)$ for the set of pure strategies for i and write $S := S_e \times S_\ell \times S_c$. Write $\psi : S \rightrightarrows V$ for the **path correspondence** of \mathcal{M} . So, $\psi(s)$ denotes the set of nodes of V that constitute the path of play under $s \in S$. Write $\zeta : S \times V \rightarrow Z$ for the **end-node mapping**. So, $\zeta(s | v) \in Z$ is the end node that would be realized if the interaction starts at v and actions are subsequently played according s .

Behavioral strategies of supergames Fix a mechanism \mathcal{M} . To define the strategies of the supergame (\mathcal{M}, G) , it is useful to extend the action correspondence from $\mathcal{A}_i : \mathcal{I}_i \setminus \mathcal{T}_i \rightrightarrows X$ to $\mathcal{A}_i : \mathcal{I}_i \rightrightarrows X \cup A_i$ so that (i) for each terminal information set $T_i \in \mathcal{T}_i$, $\mathcal{A}_i(T_i) = A_i$, and (ii) for each $I_i \in \mathcal{I}_i \setminus \mathcal{T}_i$, the action set $\mathcal{A}_i(I_i)$ corresponds to what it was originally defined in \mathcal{M} .

A **behavioral strategy** for the expert is a mapping $\sigma_e : \Theta \rightarrow \prod_{I_e \in \mathcal{I}_e} \Delta(\mathcal{A}_e(I_e))$. A **behavioral strategy** for the layman is a vector $\sigma_\ell \in \prod_{I_\ell \in \mathcal{I}_\ell} \Delta(\mathcal{A}_\ell(I_\ell))$.

Write $S_i := \prod_{I_i \in \mathcal{I}_i \setminus \mathcal{T}_i} \mathcal{A}_i(I_i)$ for the set of pure strategies for i and write $S := S_e \times S_\ell \times S_c$. Write $\psi : S \rightrightarrows V$ for the **path correspondence** of \mathcal{M} . So, $\psi(s)$ denotes the set of nodes of V that constitute the path of play under $s \in S$. Write $\zeta : S \times V \rightarrow Z$ for the **end-node mapping**. So, $\zeta(s | v) \in Z$ is the end node that would be realized if the interaction starts at v and actions are subsequently played according s .

Each pair (θ, σ_e) (resp. each σ_ℓ) induces a probability distribution over S_e (resp. S_ℓ) given by

$$P(s_e | \theta, \sigma_e) := \prod_{I_e \in \mathcal{I}_e \setminus \mathcal{T}_e} \sigma_e(\theta)(\text{proj}_{I_e} s_e) \quad \text{and} \quad P(s_\ell | \sigma_\ell) := \prod_{I_\ell \in \mathcal{I}_\ell \setminus \mathcal{T}_\ell} \sigma_\ell(\text{proj}_{I_\ell} s_\ell).$$

Likewise, write $P(s_c | \rho_c)$ for the probability distribution over S_c that ρ_c induces.

The profile (σ_e, σ_ℓ) and the prior $\mu \in \Delta(\Theta)$ induce a distribution of paths over $\Theta \times V$. So, for

each $(\theta, v) \in \Theta \times V$, the ex-ante probability that θ occurs and the path goes through v , given that (s_e, σ_ℓ) (resp. (σ_e, s_ℓ)) is played, is

$$P(\theta, v \mid s_e, \sigma_\ell) := \sum_{(s_\ell, s_c) \in S_\ell \times S_c} \mu(\theta) \cdot P(s_\ell \mid \sigma_\ell) \cdot P(s_c \mid \rho_c) \cdot \mathbb{1}[v \in \psi(s_e, s_\ell, s_c)], \quad \text{and}$$

$$P(\theta, v \mid \sigma_e, s_\ell) := \sum_{(s_e, s_c) \in S_e \times S_c} \mu(\theta) \cdot P(s_e \mid \sigma_e) \cdot P(s_c \mid \rho_c) \cdot \mathbb{1}[v \in \psi(s_e, s_\ell, s_c)].$$

For each $s = (s_e, s_\ell, s_c) \in S$, write $P(s \mid \theta, \sigma) := P(s_e \mid \theta, \sigma_e) \cdot P(s_\ell \mid \sigma_\ell) \cdot P(s_c \mid \rho_c)$. So, the probability of reaching $(\theta, v) \in \Theta \times V$ under profile $\sigma = (\sigma_e, \sigma_\ell)$ is

$$P(\theta, v \mid \sigma) := \sum_{s \in S} \mu(\theta) \cdot P(r \mid \theta, \sigma) \cdot \mathbb{1}[v \in \psi(s_e, s_\ell, s_c)],$$

and the probability of reaching v conditional on state θ is $P(v \mid \sigma, \theta) := P(\theta, v \mid \sigma) \cdot \mu(\theta)^{-1}$. Write

$$P(T_e, T_\ell \mid \theta, I_e, \sigma, \beta_e) := \sum_{v \in I_e} \sum_{s \in S} \beta_e(\theta, I_e)(v) \cdot P(r \mid \theta, \sigma) \cdot \mathbb{1}[\zeta(s, v) \in T_e \cap T_\ell]$$

for the probability that the e assigns to (T_e, T_ℓ) given that the state is θ , information set I_e is reached, σ is played and e has interim beliefs β_e . Likewise, write

$$P(\theta, T_e, T_\ell \mid I_\ell, \sigma, \beta_\ell) := \sum_{\theta \in \Theta} \sum_{v \in I_e} \sum_{s \in S} \beta_\ell(I_\ell)(\theta, v) \cdot P(r \mid \theta, \sigma) \cdot \mathbb{1}[\zeta(s, v) \in T_e \cap T_\ell]$$

for the probability that ℓ assigns to (θ, T_e, T_ℓ) given that I_ℓ is reached, σ is played, and ℓ has interim beliefs β_ℓ . The expected transfers induced by σ are

$$Y_e(\theta, \mathcal{M}, \sigma) := \sum_{T_e \in \mathcal{T}_e} \gamma_e(T_e) \cdot P(T_e \mid \sigma, \theta), \quad \text{and}$$

$$Y_\ell(\mathcal{M}, \sigma) := \sum_{T_\ell \in \mathcal{T}_\ell} \gamma_\ell(T_\ell) \cdot P(\Theta \times T_\ell \mid \sigma).$$

Finally, for each $(\theta, v) \in \Theta \times Z$, write $\mathbb{P}(\theta, v) := P(\theta, v \mid \sigma)$. So, each strategy profile σ induces a probability measure $\mathbb{P} \in \Delta(\Theta \times Z)$ that specifies the distribution of states and terminal nodes.

Behavioral strategies of psychological games Fix a psychological game $(\mathcal{M}, u_e, u_\ell)$. The expert's behavior is described by a behavioral strategy $\rho_e : \Theta \rightarrow \prod_{I_e \in \mathcal{I}_e \setminus \mathcal{T}_e} \Delta(\mathcal{A}_e(I_e))$. The layman's behavior is described by a behavioral strategy $\rho_\ell \in \prod_{I_\ell \in \mathcal{I}_\ell \setminus \mathcal{T}_\ell} \Delta(\mathcal{A}_\ell(I_\ell))$.

The probabilities of states and paths $P(\cdot)$, the expected transfers $Y_i(\cdot)$, and the terminal probability distribution $\mathbb{P} \in \Delta(\Theta \times Z)$ that ρ induces are defined in analogous way as above.

B Appendix: A Payoff-Relevant Domain of Uncertainty

This Appendix reviews the reduced-form approach introduced in **RM**. In that framework, **RM** constructs a belief structure $H_e \times H_\ell$ using domains of uncertainty $D_\ell^1 = \{\triangleleft\}$ and $D_e^1 = \Theta$ (where \triangleleft represents a trivial element). So, unlike the extended framework presented in Section 2, this belief structure does not accommodate redundancies over the payoff-relevant state space Θ . Under **RM**'s framework, reduced forms and reduced-form representations are defined as follows:

Definition B.1. *Call (u_e, u_ℓ) a **reduced form** for G if, for each mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and interim belief mappings $\beta = (\beta_e, \beta_\ell) \in \text{cons}(\mathcal{M})$ (that induce (δ_e, δ_ℓ)), there is a Bayesian equilibrium $\hat{\sigma}$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta_e, \beta_\ell)$ such that the following hold:*

- (i) *For each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, $u_e(\theta, \delta_e(\theta, T_e)) = \Pi_e(\hat{\sigma} \mid \theta, T_e, \beta_e)$.*
- (ii) *For each $T_\ell \in \mathcal{T}_\ell$, $u_\ell(\delta_\ell(T_\ell)) = \Pi_\ell(\hat{\sigma} \mid T_\ell, \beta_\ell)$.*

Definition B.2. *Fix a set RF of reduced forms of G . Say RF is a **reduced-form representation** of G if, for each mechanism \mathcal{M} and each PBE (σ, β) of (\mathcal{M}, G) , there is a mechanism \mathcal{M}' , a reduced form $(u_e, u_\ell) \in \text{RF}$, and a PBE of $(\mathcal{M}', u_e, u_\ell)$ that is equivalent to (σ, β) .*

Throughout this appendix, we maintain the belief structure introduced in **RM**, in which the domains of uncertainty are $D_\ell^1 = \{\triangleleft\}$ and $D_e^1 = \Theta$. The appendix provides two different methods for constructing reduced-form representations for two classes of games. The first approach leverages games with unique equilibria, while the second examines games with inactive experts. Additionally, the appendix shows that Example 1.1 has no reduced-form representation within this restricted framework, highlighting the necessity of the extended domain of uncertainty for some games.

B.1 Uniqueness of Equilibria

Definition B.3. *Fix a game $G = ((A_i, \pi_i) : i \in \{e, \ell\})$ and let $\varsigma_e : \Theta \times H_e \rightarrow \Delta(A_e)$ and $\varsigma_\ell : H_\ell \rightarrow \Delta(A_\ell)$ be measurable mappings. Say $(\varsigma_e, \varsigma_\ell)$ is a **universal Bayesian equilibrium** of G if each induced Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ has a unique Bayesian equilibrium $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ such that*

- (i) *for each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, $\hat{\sigma}_e(\theta, T_e) = \varsigma_e(\theta, \delta_e(\theta, T_e))$, and*
- (ii) *for each $T_\ell \in \mathcal{T}_\ell$, $\hat{\sigma}_\ell(T_\ell) = \varsigma_\ell(\delta_\ell(T_\ell))$.*

A universal Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ captures situations where equilibrium behavior is uniquely determined by the state and belief hierarchies, regardless of the specific type structure. Fix belief-based functions (u_e, u_ℓ) and a universal Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$. Say $(\varsigma_e, \varsigma_\ell)$ **induces** (u_e, u_ℓ) if for each tuple $(\theta, h_e, h_\ell) \in \Theta \times H_e \times H_\ell$,

$$u_e(\theta, h_e) = \int_{H_\ell} \int_{A_e} \int_{A_\ell} \pi(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e) d\varsigma_\ell(h_\ell) dh_\ell^\infty, \text{ and}$$

$$u_\ell(h_\ell) = \int_{\Theta \times H_e} \int_{A_e} \int_{A_\ell} \pi(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e) d\varsigma_\ell(h_\ell) dh_e^\infty.$$

Lemma B.1. Fix a is a universal Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ and belief-based utilities (u_e, u_ℓ) . If $(\varsigma_e, \varsigma_\ell)$ induces (u_e, u_ℓ) , then (u_e, u_ℓ) is a reduced form of G . Moreover, $\text{RF} = \{(u_e, u_\ell)\}$ is a reduced-form representation for G .

Proof. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$, belief mappings $\beta \in \text{cons}(\mathcal{M})$. We will show that there is a unique BE $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ such that the following hold:

- (a) For each $(\theta, T_e) \in \Theta \times \mathcal{T}_e$, $u_e(\theta, \delta_e(\theta, T_e)) = \Pi_e(\hat{\sigma} | \theta, T_e, \beta_e)$.
- (b) For each $T_\ell \in \mathcal{T}_\ell$, $u_\ell(\delta_\ell(T_\ell)) = \Pi_\ell(\hat{\sigma} | T_\ell, \beta_\ell)$.

Showing (a) and (b) implies that (u_e, u_ℓ) is a reduced form. Moreover, since $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ is unique they also imply that $\text{RF} = \{(u_e, u_\ell)\}$ is a reduced-form representation.

Notice, since $(\tilde{\sigma}_e, \tilde{\sigma}_\ell)$ is a universal equilibrium of G , the Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ has a unique BE $(\hat{\sigma}_e, \hat{\sigma}_\ell)$. Moreover, $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ satisfies conditions (i) and (ii) from Definition B.3. We show that conditions (i) and (ii) imply (a). Fix $(\theta, T_e) \in \Theta \times \mathcal{T}_e$ and let $h_e = \delta_e(\theta, h_e)$. Conditions (i) and (ii) imply that

$$\begin{aligned} u_e(\theta, \delta_e(\theta, T_e)) &= \int_{H_\ell} \int_{A_e} \int_{A_\ell} \pi(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e) d\varsigma_\ell(h_\ell) dh_e^\infty \\ &= \int_{T_\ell} \int_{A_e} \int_{A_\ell} \pi(\theta, a_e, a_\ell) d\hat{\sigma}_e(\theta, T_e) d\hat{\sigma}_\ell(T_\ell) d\beta_e(\theta, T_e) \\ &= \Pi_e(\hat{\sigma} | \theta, T_e, \beta_e). \end{aligned}$$

Showing (b) follows from an analogous argument. ■

B.2 An Inactive Expert

Say that the **expert is inactive** in $G = ((A_i, \pi_i) : i \in \{e, \ell\})$ if A_e is a singleton. We now show the existence of reduced-form representations for games with inactive experts. Write $A_\ell^*(h_\ell^1) := \arg \max_{a_\ell \in A_\ell} \sum_{\theta \in \Theta} \pi(\theta, a_\ell) \cdot h_\ell^1(\theta)$ for the optimal actions given a belief h_ℓ^1 , and write

$$Q(G) = \left\{ q : \Delta(\Theta) \rightarrow \Delta(A_\ell) : \text{Supp}(q(h_\ell^1)) \subseteq A_\ell^*(h_\ell^1) \text{ for each } h_\ell^1 \in \Delta(\Theta) \right\}$$

for the optimal layman's plans contingent on ℓ 's first-order beliefs. Say that the belief-based utilities (u_e, u_ℓ) are induced by $q \in Q(G)$ if

$$\begin{aligned} u_\ell(h_\ell^1, h_\ell^2, \dots) &= \int_{A_\ell} \pi_\ell(h_\ell^1, a_\ell) dq(h_\ell^1), \text{ and} \\ u_e(\theta, h_e^1, h_e^2, \dots) &= \int_{\Delta(\Theta)} \int_{A_\ell} \pi_e(\theta, a_\ell) dq(h_\ell^1) dh_e^2. \end{aligned}$$

Lemma B.2. Suppose that the expert is inactive in G and $Q(G) \neq \emptyset$. If (u_e, u_ℓ) is induced by $q \in Q(G)$, then (u_e, u_ℓ) is a reduced form for G .

Proof. Fix $q \in Q(G)$ and let (u_e, u_ℓ) be induced by q . Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and some $\beta \in \text{cons}(\mathcal{M})$ that induces (δ_e, δ_ℓ) . Consider the induced Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$. Define the strategy profile $\hat{\sigma}$ of the induced Bayesian game by $\hat{\sigma}(T_\ell) = q(\delta_\ell^1(T_\ell))$. (Since the expert is inactive, $\hat{\sigma}_e$ is trivial.) Observe, since $q \in Q(G)$, it follows that $\hat{\sigma}$ is a Bayesian equilibrium. Moreover, for each $(\theta, T_e, T_\ell) \in \Theta \times \mathcal{T}_e \times \mathcal{T}_\ell$, $\Pi_\ell(\hat{\sigma} \mid T_\ell, \beta_\ell) = u_\ell(\delta_\ell(T_\ell))$ and $\Pi_e(\hat{\sigma} \mid \theta, T_e, \beta_e) = u_e(\theta, \delta_e(\theta, T_e))$. ■

Lemma B.3. *Suppose that the expert is inactive in G and $Q(G) \neq \emptyset$. The set of reduced forms $\text{RF} = \{(u_e, u_\ell) : (u_e, u_\ell) \text{ is induced by some } q \in Q(G)\}$ is a reduced-form representation for G .*

Proof. Fix a mechanism \mathcal{M} and a PBE (σ, β) of (\mathcal{M}, G) . The proof is divided in two steps. The first defines some $q \in \Gamma(G)$ mimics σ_ℓ . The second shows there is some PBE of some psychological game that is equivalent to (σ, β) .

Step 1. Let δ_ℓ^1 be layman's first-order belief mapping that β induces, and let $M \subset \Delta\Theta$ be the range of δ_ℓ^1 . (Recall, $\delta_\ell^1(T_\ell)$ describes ℓ 's first-order posterior belief at $T_\ell \in \mathcal{T}_\ell$.) For each $h_\ell^1 \in M$ write $\mathcal{T}_\ell[h_\ell^1] := \{T_\ell \in \mathcal{T}_\ell : \delta_\ell(T_\ell) = h_\ell^1\}$, for ℓ 's terminal information sets that induce h_ℓ^1 . So, $\{\mathcal{T}_\ell[h_\ell^1] : h_\ell^1 \in M\}$ is a partition of \mathcal{T}_ℓ . Since σ_ℓ is optimal, $\delta_\ell^1(T_\ell) = h_\ell^1$ implies $\text{Supp}(\sigma_\ell(T_\ell)) \subseteq \arg \max_{A_\ell} \pi_\ell(h_\ell^1, a_\ell)$. Thus, there is some $q \in Q(G)$ such that for each posterior $h_\ell^1 \in M$,

$$q(h_\ell^1) \cdot \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} P(T_\ell \mid \sigma) = \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \sigma_\ell(T_\ell) \cdot P(T_\ell \mid \sigma). \quad (5)$$

Notice that for each $\theta \in \Theta$ and each $T_\ell \in \mathcal{T}_\ell[h_\ell^1]$,

$$\mu(\theta) \cdot P(T_\ell \mid \theta, \sigma) = P(\theta, T_\ell \mid \sigma) = h_\ell^1(\theta) \cdot P(T_\ell \mid \sigma), \quad (6)$$

where the first equality follows from definition of P and the second from consistency of σ and δ_ℓ .

Fix $\theta \in \Theta$ and write $P(h_\ell^1 \mid \sigma) := \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} P(T_\ell \mid \sigma)$, and $P(h_\ell^1 \mid \theta, \sigma) := \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} P(T_\ell \mid \theta, \sigma)$. By adding Equation (6) over $T_\ell \in \mathcal{T}_\ell[h_\ell^1]$, it follows that

$$\mu(\theta) \cdot P(h_\ell^1 \mid \theta, \sigma) = h_\ell^1(\theta) \cdot P(h_\ell^1 \mid \sigma). \quad (7)$$

Thus, if $P(h_\ell^1 \mid \theta, \sigma) > 0$, then $h_\ell^1(\theta) > 0$ and $\sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} P(T_\ell \mid \sigma) > 0$. Moreover,

$$\begin{aligned} \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \sigma_\ell(T_\ell) \cdot \frac{P(T_\ell \mid \theta, \sigma)}{P(h_\ell^1 \mid \theta, \sigma)} &= \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \sigma_\ell(T_\ell) \cdot \frac{\mu(\theta) \cdot h_\ell^1(\theta) \cdot P(T_\ell \mid \sigma)}{\mu(\theta) \cdot h_\ell^1(\theta) \cdot P(h_\ell^1 \mid \sigma)} \\ &= \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \sigma_\ell(T_\ell) \cdot \frac{P(T_\ell \mid \sigma)}{P(h_\ell^1 \mid \sigma)} \\ &= q(h_\ell^1), \end{aligned}$$

where the first equality follows from Equations (7) and (6). This implies that

$$\sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \sigma_\ell(T_\ell) \cdot P(T_\ell | \theta, \sigma) = q(h_\ell^1) \cdot \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} P(T_\ell | \theta, \sigma) \quad (8)$$

For each $\alpha \in \Delta A_\ell$, write $\Pi_i(\theta, \alpha) := \int_{A_\ell} \pi_i(\theta, a_\ell) d\alpha$ for i 's expected payoff given state θ and distribution of actions α . Observe Π_i is linear in α . Hence, by Equation (8), for each $\theta, \theta' \in \Theta$,

$$\sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \Pi_i(\theta, q(h_\ell^1)) \cdot P(T_\ell | \theta', \sigma) = \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \Pi_i(\theta, \sigma_\ell(T_\ell)) \cdot P(T_\ell | \theta', \sigma). \quad (9)$$

Step 2. Let $(u_e, u_\ell) \in \text{RF}$ be the belief-dependent utility functions induced by q . It suffices to find a mechanism \mathcal{M}' and a PBE of $(\mathcal{M}', u_e, u_\ell)$ that is equivalent to (σ, β) .

We first construct \mathcal{M}' . In this mechanism, the layman is inactive and the expert selects an element from the set of pure strategies $S_e := \Theta$. The layman does not observe which strategy s_e is selected. Conditional on s_e , chance selects a message in M for the layman. If the expert selects $s_e = \theta$, then the expert receives a transfer $Y_e(\theta, \mathcal{M}, \sigma)$ and chance selects a message $h_\ell^1 \in M$ with probability $P(h_\ell^1 | \theta, \sigma)$ that is sent to the layman. The mechanism provides a fixed transfer $Y_\ell(\mathcal{M}, \sigma)$ to the layman. Consider the strategy $\rho_e : \Theta \rightarrow \Delta(S_e)$ so that $\rho_e(\theta)(\theta) = 1$ for each θ . Let ρ_ℓ the trivial strategy of the layman and write β' for beliefs consistent with $\rho = (\rho_e, \rho_\ell)$. We show that (ρ, β') is PBE of $(\mathcal{M}, u_e, u_\ell)$ that is equivalent to (σ, β) .

To show this, it will be convenient to write two pieces of notation. Write σ'_e for strategy on (\mathcal{M}, G) such that $\sigma'_e(\cdot)(I_e) = \sigma_e(\theta')(I_e)$ for each $I_e \in \Theta \times (\mathcal{I}_e \setminus \mathcal{T}_e)$. Write $V_e(\theta, \theta')$ for the expert's expected utility in $(\mathcal{M}, u_e, u_\ell)$ when the expert selects $s_e = \theta'$, provided that the state is θ , and interim beliefs are β' . To show that (ρ, β') is a PBE it suffices to show that $V_e(\theta, \theta) \geq V_e(\theta, \theta')$ for each $\theta, \theta' \in \Theta$. Observe, since q induces u_e ,

$$\begin{aligned} \mathcal{V}_e(\theta, \theta') &= \sum_{h_\ell^1 \in M} \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \Pi_e(\theta, q(h_\ell^1)) \cdot P(T_\ell | \theta', \sigma) + Y_e(\theta', \mathcal{M}, \sigma) \\ &= \sum_{h_\ell^1 \in M} \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell^1]} \Pi_e(\theta, \sigma_\ell(T_\ell)) \cdot P(T_\ell | \theta', \sigma) + Y_e(\theta', \mathcal{M}, \sigma) \\ &= U_e(\sigma'_e, \sigma_\ell | \emptyset, \beta_e), \end{aligned} \quad (10)$$

where the second equality follows from Equation (9). Observe, in particular, this equation implies that $V_e(\theta, \theta) = U_e(\sigma_e, \sigma_\ell | \emptyset, \beta_e)$. Since (σ, β) is a PBE, $U_e(\sigma_e, \sigma_\ell | \emptyset, \beta_e) \geq U_e(\sigma'_e, \sigma_\ell | \emptyset, \beta_e)$. Thus, $V_e(\theta, \theta') \geq V_e(\theta, \theta)$

Using an analogous argument as above, the expected payoff for the layman under (ρ, β') is

$$\sum_{\theta \in \Theta} \sum_{T_\ell \in \mathcal{T}_\ell} \Pi_\ell(\theta, \sigma_\ell(T_\ell)) \cdot P(\theta, T_\ell | \sigma) + Y_\ell(\mathcal{M}, \sigma) = U_\ell(\sigma | \emptyset, \beta_\ell). \quad (11)$$

Therefore, Equations (10) and (11) imply that (σ, β) is equivalent to (ρ, β) for the expert. \blacksquare

B.3 A Game with no Reduced-Form Representation

Lemma B.4. *The game G from Example 1.1. If the belief structure has a trivial domain of uncertainty $D_e = D_\ell = \{\triangleleft\}$, then G has no reduced-form representation.*

Proof. First notice that there is a supergame (\mathcal{M}, G) and a PBE thereof that provides expected payoffs of $(5, 5)$. To see this, define \mathcal{M} as follows: Chance selects an element from the set $\{(\bar{a}_e, \underline{a}_\ell), (\bar{a}_e, \bar{a}_\ell), (\underline{a}_e, \bar{a}_\ell)\}$ with equal probability. Once chance selects (a_e, a_ℓ) the mechanisms sends a private message each agent i . So, agent i privately observes the component a_i of the vector that chance selected. There are no transfers. Let (σ, β) be a consistent profile such that follow the mechanism's recommendation. Notice that no agent has incentives to deviate and thus, (σ, β) is a PBE of (\mathcal{M}, G) and provides expected payoffs of 5 to each agent.

Since in this game the set of payoff-relevant parameters $\Theta = \{\theta\}$ is a singleton, the belief structure $H_e \times H_\ell$ is also a singleton. Thus, to show that G has not reduced form representation, it suffices to show that (u_e, u_ℓ) with $(u_e(\theta, h_e), u_\ell(h_\ell)) = (5, 5)$ is not a reduced form. We show this by contradiction.

Assume that (u_e, u_ℓ) is a reduced form with $(u_e(\theta, h_e), u_\ell(h_\ell)) = (5, 5)$. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ with $\mathcal{T}_e = \{T_e\}$ and $\mathcal{T}_\ell = \{T_\ell\}$ both singletons, and $\beta \in \text{Cons}(\mathcal{M})$. Consider the Bayesian game $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$. Since (u_e, u_ℓ) is a reduced form, there is a Bayesian equilibrium $\tilde{\sigma}$ of $\text{BG}(\mathcal{T}_e, \mathcal{T}_\ell, \beta)$ such that:

- (i) $\Pi_e(\tilde{\sigma} \mid \theta, T_e, \beta_e) = u_e(\theta, h_e) = 5$.
- (ii) $\Pi_\ell(\tilde{\sigma} \mid T_\ell, \beta_\ell) = u_\ell(h_\ell) = 5$.

Since \mathcal{T}_e and \mathcal{T}_ℓ are both singletons, this implies that G has a Bayesian equilibrium with payoffs $(5, 5)$, leading to a contradiction. \blacksquare

C Appendix: Omitted Proofs

Lemma C.1. *Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and interim beliefs $\beta \in \text{cons}(\mathcal{M})$. Write (δ_e, δ_ℓ) for the hierarchy mappings induced by β .*

- (i) *For each $\theta \in \Theta$, the mapping $\delta_e(\theta, \cdot) : \mathcal{T}_e \rightarrow H_e$ is injective.*
- (ii) *The mapping $\delta_\ell : \mathcal{T}_\ell \rightarrow H_\ell$ is injective.*

Proof. We show Part (i). (Part (ii) follows from an analogous argument.) Let L be the labeling of \mathcal{M} . For each $T_i \subset Z$ write $L(T_i) = \{\lambda \in \Lambda : \lambda = L(z) \text{ for some } z \in T_i\}$. Fix $\theta \in \Theta$ and $T_e, T'_e \in \mathcal{T}_e$ such that $T_e \neq T'_e$. Note that β_e satisfies $\text{Supp}(\beta_e(\theta, T_e)) \subseteq T_e$ and $\text{Supp}(\beta_e(\theta, T'_e)) \subseteq T'_e$. So, if $h_e = \delta_e(\theta, T_e)$ and $h'_e = \delta_e(\theta, T'_e)$, then h_e and h'_e satisfy $\text{Supp}(\text{marg}_\Lambda h_e^\infty) \subseteq L(T_e)$ and $\text{Supp}(\text{marg}_\Lambda h'_e^\infty) \subseteq L(T'_e)$. Moreover, since each labeling is injective, $L(T_e) \cap L(T'_e) = \emptyset$. As a result, $\text{Supp}(\text{marg}_\Lambda h_e^\infty) \neq \text{Supp}(\text{marg}_\Lambda h'_e^\infty)$. This implies $h_e \neq h'_e$. \blacksquare

Lemma C.2. *The set $\text{CP}_e \times \text{CP}_\ell$ is belief closed. Moreover, For each $(h_e, h_\ell) \in \text{CP}_e \times \text{CP}_\ell$, $\text{marg}_{H_\ell} h_e^\infty(h_\ell) > 0$ if and only if $\text{marg}_{H_e} h_\ell^\infty(h_e) > 0$.*

Proof of Lemma C.2. Fix $h_i \in \text{CP}_i$. So, there are some $M_e \subsetneq \text{CP}_e$ and some $M_\ell \subsetneq \text{CP}_\ell$ such that $h_e \in M_e$ and $M_e \times M_\ell$ is a belief-closed set that admits a common prior. Observe that $h_i^\infty(D_i^1 \times \text{CP}_{-i}) \geq h_i^\infty(D_i^1 \times M_{-i}) = 1$. Hence, $\text{CP}_e \times \text{CP}_\ell$ is belief closed.

We now show the second part. Fix $(h_e, h_\ell) \in \text{CP}_e \times \text{CP}_\ell$ and assume $\text{marg}_{H_\ell} h_e^\infty(h_\ell) > 0$. Observe, since $h_e \in \text{CP}_e$, there is some $M_e \subsetneq H_e$ and some $M_\ell \subsetneq H_\ell$ such that $M_e \times M_\ell$ that admits a common prior and $h_e \in M_e$. Thus, there is some $\phi \in \Delta(\Theta \times \Lambda \times M_e \times M_\ell)$ that satisfies Equation (1). Notice, since $M_e \times M_\ell$ admits a common prior, it is belief-closed and thus $h_\ell \in M_\ell$. Moreover,

$$\begin{aligned} \text{marg}_{H_\ell} h_e^\infty(h_\ell) \cdot \text{marg}_{M_e} \phi(h_e) &= \sum_{\lambda \in \Lambda} \sum_{\theta \in \Theta} h_e^\infty(\lambda, h_\ell) \cdot \text{marg}_{\Theta \times M_e} \phi(\theta, h_e) \\ &= \sum_{\lambda \in \Lambda} \sum_{\theta \in \Theta} h_\ell^\infty(\theta, \lambda, h_e) \cdot \text{marg}_{M_\ell} \phi(h_\ell) \\ &= \text{marg}_{H_e} h_\ell^\infty(h_e) \cdot \text{marg}_{M_\ell} \phi(h_\ell), \end{aligned}$$

where the second equality follows from Equation (1). Observe, since $\text{marg}_{M_e} \phi(h_e) > 0$ and $\text{marg}_{M_\ell} \phi(h_\ell) > 0$, it follows that $\text{marg}_{H_\ell} h_e^\infty(h_\ell) > 0$ holds if and only if $\text{marg}_{H_e} h_\ell^\infty(h_e) > 0$. ■

The next result requires some definitions. Fix a set $M_e \times M_\ell \subset \text{CP}_e \times \text{CP}_\ell$. Call $M_e \times M_\ell$ **minimal** if $M_e \times M_\ell$ is belief closed and there is no other subset $M'_e \times M'_\ell \subsetneq M_e \times M_\ell$ that is belief closed. Write $\text{MIN} := \{M_e \times M_\ell \subset \text{CP}_e \times \text{CP}_\ell : M_e \times M_\ell \text{ is minimal}\}$ for the class of minimal sets. Since each minimal set $M_e \times M_\ell$ is belief closed, it induces a Θ -based type structure $(M_e, M_\ell, b_e, b_\ell)$. In this type structure, e 's set of types is $\Theta \times M_e$ and the set of types for ℓ is M_ℓ . The belief mappings $b_e : \Theta \times M_e \rightarrow \Delta(M_\ell)$ and $b_\ell : M_\ell \rightarrow \Delta(\Theta \times M_e)$ satisfy

$$b_e(\theta, h_e)(h_\ell) = \text{marg}_{H_\ell} h_e^\infty(h_\ell), \quad \text{and} \quad b_\ell(h_\ell)(\theta, h_e) = \text{marg}_{\Theta \times H_e} h_\ell^\infty(\theta, h_e).$$

Write $\text{BG}(M_e, M_\ell)$ for the Bayesian game induced by G and the type structure $(M_e, M_\ell, b_e, b_\ell)$.

Lemma C.3. *The set of minimal subsets MIN is a partition of $\text{CP}_e \times \text{CP}_\ell$. Moreover, each $M_e \times M_\ell \in \text{MIN}$ is finite.*

Proof. We divide the proof into two steps. The first shows that MIN is a partition of $\text{CP}_e \times \text{CP}_\ell$. the second shows that each $M_e \times M_\ell \in \text{MIN}$ is finite.

Step 1. We show that each pair $(h_e, h_\ell) \in \text{CP}_e \times \text{CP}_\ell$ belongs to some minimal set. Fix $(h_e, h_\ell) \in \text{CP}_e \times \text{CP}_\ell$. Let \mathcal{B} be the collection of all belief-closed sets containing (h_e, h_ℓ) . Note that \mathcal{B} is non-empty as $\text{CP}_e \times \text{CP}_\ell \in \mathcal{B}$. Let $M = \bigcap_{M' \in \mathcal{B}} M'$. We claim that M is belief-closed and minimal. Note that M is belief-closed: For each $(h'_e, h'_\ell) \in M$, the support of $h'_e{}^\infty$ and $h'_\ell{}^\infty$ is contained in each $M \in \mathcal{B}$, hence in their intersection M . Next, M is minimal: If there were a proper subset $M' \subsetneq M$ that is belief-closed and contains (h_e, h_ℓ) , then $M' \in \mathcal{B}$, contradicting the fact that M is

the intersection of all sets in \mathcal{B} . Therefore, (h_e, h_ℓ) belongs to the minimal set M .

Next, we need to show that no pair (h_e, h_ℓ) belongs to two different minimal sets. Suppose, for contradiction, that $(h_e, h_\ell) \in M$, $(h_e, h_\ell) \in M'$, $M, M' \in \text{MIN}$ and $M \neq M'$. Let $M'' = M \cap M'$. Since both M and M' are belief-closed, M'' is also belief-closed. Moreover, M'' contains (h_e, h_ℓ) and is a proper subset of both M and M' (as $M \neq M'$). This contradicts the fact that M and M' are minimal. Therefore, each pair (h_e, h_ℓ) belongs to one minimal set, and hence, MIN is a partition of $\text{CP}_e \times \text{CP}_\ell$.

Step 2. Fix some $M \in \text{MIN}$ and write $M := M_e \times M_\ell$. Fix some $h_e \in M_e$ and some $h_\ell \in \text{marg}_{H_\ell} \text{Supp}(h_e^\infty)$. (Since M is belief closed, $h_\ell \in M_\ell$.) Write $\bar{M}_e := \text{marg}_{H_e} \text{Supp}(h_e^\infty)$, $\bar{M}_\ell := \text{marg}_{H_\ell} \text{Supp}(h_e^\infty)$, and $\bar{M} := \bar{M}_e \times \bar{M}_\ell$.

Notice, since $(h_e, h_\ell) \in \text{CP}_e \times \text{CP}_\ell$ the extensions h_e^∞ and h_ℓ^∞ have finite supports. Hence, \bar{M} is finite. Thus, since M is minimal and $\bar{M} \subset M$, to show that $M = \bar{M}$ it suffices to show that \bar{M} is belief closed.

To show that \bar{M} is belief-closed it suffices to show that $\bar{M}_\ell = \text{marg}_{H_\ell} \text{Supp}(h_e'^\infty)$ for each $h_e' \in \bar{M}_e$, and $\bar{M}_e = \text{marg}_{H_e} \text{Supp}(h_\ell'^\infty)$ for each $h_\ell' \in \bar{M}_\ell$. Fix $h_e' \in \bar{M}_e$. Observe that, by definition, $\bar{M}_\ell = \text{marg}_{H_\ell} \text{Supp}(h_e^\infty)$. This implies that $\bar{M}_\ell = \text{marg}_{H_\ell} \text{Supp}(h_e'^\infty)$ (See Lemma C.2.) An analogous argument shows that $\bar{M}_e = \text{marg}_{H_e} \text{Supp}(h_\ell'^\infty)$ for each $h_\ell' \in \bar{M}_\ell$. Thus, \bar{M} is belief close, as desired. \blacksquare

Lemma C.4. *Let A_e and A_ℓ be the action sets of G . If A_e and A_ℓ are finite, then G is nice.*

Proof of Lemma C.4. Observe, for each $M = M_e \times M_\ell \in \text{MIN}$, the Bayesian game $\text{BG}(M_e, M_\ell)$ is a Bayesian game with finite action sets and finite set of types. Thus, the Bayesian game $\text{BG}(M_e, M_\ell)$ has some Bayesian equilibrium $\varsigma_e^M, \varsigma_\ell^M$ where $\varsigma_e^M : \Theta \times M_e \rightarrow \Delta(A_e)$ and $\varsigma_\ell^M : M_\ell \rightarrow \Delta(A_\ell)$.

For each $M_e \times M_\ell \in \text{MIN}$, fix a Bayesian equilibrium $(\varsigma_e^M, \varsigma_\ell^M)$ of $\text{BG}(M_e, M_\ell)$. We will use the equilibria $(\varsigma_e^M, \varsigma_\ell^M)_{M \in \text{MIN}}$ to construct a Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of the Bayesian game $\text{BG}(\mathcal{S}_{\text{CP}})$. Observe, since MIN is a partition of $\text{CP}_e \times \text{CP}_\ell$, there is a mapping $\varsigma_e : \Theta \times \text{CP}_e \rightarrow \Delta(A_e)$ and a mapping $\varsigma_\ell : \text{CP}_\ell \rightarrow \Delta(A_\ell)$ such that for each $M \in \text{MIN}$ and each $(h_e, h_\ell) \in M$

$$\varsigma_e(\theta, h_e) = \varsigma_e^M(\theta, h_e) \quad \text{and} \quad \varsigma_\ell(h_\ell) = \varsigma_\ell^M(h_\ell).$$

It suffices to show that $(\varsigma_e, \varsigma_\ell)$ is a BE of $\text{BG}(\mathcal{S}_{\text{CP}})$. To prove this, we need to show that ς_e and ς_ℓ are mutual best responses.

First, consider the expert. Fix a pair $(\theta, h_e) \in \Theta \times \text{CP}_e$. There exists a unique minimal set $M_e \times M_\ell \in \text{MIN}$ such that $h_e \in M_e$. Observe, by definition of (σ_e, σ_ℓ) , for each action $a'_e \in A_e$, we

have:

$$\begin{aligned}
& \int_{\text{CP}_\ell} \int_{A_\ell} \int_{A_e} \pi_e(\theta, a_e, a_\ell) d\varsigma_e(\theta, h_e)(a_e) d\varsigma_\ell(h_\ell)(a_\ell) db_e(\theta, h_e)(h_\ell) \\
&= \int_{M_\ell} \int_{A_\ell} \int_{A_e} \pi_e(\theta, a_e, a_\ell) d\varsigma_e^M(\theta, h_e)(a_e) d\varsigma_\ell^M(h_\ell)(a_\ell) db_e(\theta, h_e)(h_\ell) \\
&\geq \int_{M_\ell} \int_{A_\ell} \int_{A_e} \pi_e(\theta, a'_e, a_\ell) d\varsigma_\ell^M(h_\ell)(a_\ell) db_e(\theta, h_e)(h_\ell) \\
&= \int_{\text{CP}_\ell} \int_{A_\ell} \pi_e(\theta, a'_e, a_\ell) d\varsigma_\ell(h_\ell)(a_\ell) db_e(\theta, h_e)(h_\ell),
\end{aligned}$$

where first and last equalities follow from the fact that $\text{Supp}(b_e(\theta, h_e)) \subseteq M_\ell$, and the inequality follows from the fact that ς_e^M is a best response to ς_ℓ^M in $\text{BG}(M_e, M_\ell)$. Therefore, ς_e is a best response to ς_ℓ in $\text{BG}(\mathcal{S}_{\text{CP}})$.

The argument for the layman is analogous. Since ς_ℓ^M is a best response for ς_e^M in $\text{BG}(M_e, M_\ell)$, $\varsigma_\ell(h_\ell)$ is a best response for $\varsigma_e(\theta, h_e)$ in $\text{BG}(\mathcal{S}_{\text{CP}})$. Therefore, $(\varsigma_e, \varsigma_\ell)$ is a Bayesian equilibrium of $\text{BG}(\mathcal{S}_{\text{CP}})$, as desired. \blacksquare

The following proof requires some notation. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and a consistent assessment (σ, β) thereof. Let Z be the set of terminal nodes of \mathcal{M} and write \mathcal{F} for the discrete σ -algebra of $\Theta \times Z$. The assessment (σ, β) induces a probability distribution $\mathbb{P} \in \Delta(\Theta \times Z)$ given by Call $\mathcal{P} = (\Theta \times Z, \mathcal{F}, \mathbb{P})$ the **terminal probability space** induced by \mathcal{M} and (σ, β) . Write $\mathbf{L}(\theta, z) = L(z)$ for the random variable in \mathcal{P} that identifies the label of each terminal node. Write Θ for the projection of $\Theta \times Z$ onto Θ . Write $T_i[z]$ for the unique element $T_i \in \mathcal{T}_i$ such that $z \in T_i$ and write $\mathbf{Y}_i(\theta, z) = \gamma_i(T_i[z])$. Let δ_i be i 's hierarchy mapping induced by β and write $\mathbf{H}_e(\theta, z) = \delta_e(\theta, T_e[z])$.

Proof of Lemma 3.1. Write $M_\ell \subsetneq H_\ell$ for the range of $\delta_\ell : \mathcal{T}_\ell \rightarrow H_\ell$ and $M_e \subsetneq H_e$ for the range of $\delta_e : \Theta \times \mathcal{T}_e \rightarrow H_e$.

We now show Part (i). Fix $\theta, \theta' \in \Theta$. It suffices to show that $\delta_e(\theta, \cdot) = \delta_e(\theta', \cdot)$ for each $\theta' \in \Theta$. Moreover, by Lemma C.4 in Rivera Mora [2024], it suffices to show that $\beta_e(\theta, \cdot) = \beta_e(\theta', \cdot)$. Fix $T_e \in \mathcal{T}_e$, $s_e \in S_e$, and $z \in T_e$. Note, by definition of $P(\cdot | s_e, \sigma_\ell)$,

$$\begin{aligned}
P(\theta, z | s_e, \sigma_\ell) \cdot \mu(\theta)^{-1} &= \sum_{(s_e, s_c) \in S_e \times S_c} P(s_e | \theta, \sigma_e) \cdot P(s_c | \sigma_c) \cdot \mathbb{1}[z \in \psi(s_e, s_\ell, s_c)] \\
&= \sum_{(s_e, s_c) \in S_e \times S_c} P(s_e | \theta', \sigma_e) \cdot P(s_c | \sigma_c) \cdot \mathbb{1}[z \in \psi(s_e, s_\ell, s_c)] \\
&= P(\theta', z | s_e, \sigma_\ell) \cdot \mu(\theta')^{-1}.
\end{aligned} \tag{12}$$

Thus,

$$\begin{aligned}
\beta_e(\theta, T_e)(z) \cdot P(\{\theta\} \times T_e \mid s_e, \sigma_\ell) \cdot \mu(\theta)^{-1} &= P(\theta, z \mid s_e, \sigma_\ell) \cdot \mu(\theta)^{-1} \\
&= P(\theta', z \mid s_e, \sigma_\ell) \cdot \mu(\theta')^{-1} \\
&= \beta_e(\theta', T_e)(z) \cdot P(\{\theta'\} \times T_e \mid s_e, \sigma_\ell) \cdot \mu(\theta')^{-1}, \quad (13)
\end{aligned}$$

where the first and third equalities follow from the fact that (σ, β) is consistent, and the second from Equation (12). Notice, adding Equation (12) over each $z \in T_e$, implies that

$$P(\{\theta\} \times T_e \mid s_e, \sigma_\ell) \cdot \mu(\theta)^{-1} = P(\{\theta'\} \times T_e \mid s_e, \sigma_\ell) \cdot \mu(\theta')^{-1}. \quad (14)$$

Moreover, since the profile (σ, β) has full path on \mathcal{M} , there is some report $s_e \in \mathcal{S}_e$ such that $P(\{\theta\} \times T_e \mid s_e, \sigma_\ell) > 0$. Hence, by Equations (14) and (13), the interim belief mapping satisfies $\beta_e(\theta, T_e) = \beta_e(\theta', T_e)$.

We now show Part (ii). Fix $\theta \in \Theta$. Observe, since M_e is the range of δ_e , part (ii) implies that M_e is the range of the mapping $\delta_e(\theta, \cdot)$. Thus, the result follows from the fact that the mapping $\delta_e(\theta, \cdot)$ is injective (See Lemma C.1.)

We now show Part (iii). Observe, the result follows from the fact that δ_ℓ is injective (See Lemma C.1) and the fact that M_ℓ is the range of δ_ℓ .

We now show Part (iv). Fix $(\theta, h_e, h_\ell) \in \Theta \times M_e \times M_\ell$ and $\lambda \in \Lambda$ such that $L(z) = \lambda$ for some $z \in Z$. By Part (iii), there is a unique $T_\ell \in \mathcal{T}_\ell$ such that $h_\ell = \delta_\ell(T_\ell)$. By Part (ii), there is a unique $T_e \in \mathcal{T}_e$ such that $\delta_e(\theta, T_e) = h_e$. Thus, $h_e^\infty(\lambda, h_\ell) = \beta_e(\theta, T_e)(\{z\} \cap T_\ell \cap T_e)$. (See Lemma C.4 in Rivera Mora [2024]). We will show that

$$h_e^\infty(\lambda, h_\ell) \cdot \mathbb{P}[\Theta = \theta, \mathbf{H}_e = h_e] = \mathbb{P}[\Theta = \theta, \mathbf{L} = \lambda, \mathbf{H}_e = h_e, \mathbf{H}_\ell = h_\ell] \quad (15)$$

Observe, if $z \in T_\ell \cap T_e$, then

$$\begin{aligned}
h_e^\infty(\lambda, h_\ell) \cdot \mathbb{P}[\Theta = \theta, \mathbf{H}_e = h_e] &= \sum_{T_e \in \mathcal{T}_e[h_e|\theta]} h_e^\infty(\lambda, h_\ell) \cdot \mathbb{P}(\{\theta\} \times T_e) \\
&= \sum_{T_e \in \mathcal{T}_e[h_e|\theta]} \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell]} \beta_e(\theta, T_e)(z) \cdot \mathbb{P}(\{\theta\} \times T_e) \\
&= \sum_{T_e \in \mathcal{T}_e[h_e|\theta]} \sum_{T_\ell \in \mathcal{T}_\ell[h_\ell]} \mathbb{P}(\theta, z) \\
&= \mathbb{P}[\Theta = \theta, \mathbf{L} = \lambda, \mathbf{H}_e = h_e, \mathbf{H}_\ell = h_\ell],
\end{aligned}$$

where the third equality from consistency of (σ, β) . If $z \notin T_\ell \cap T_e$, then both sides of Equation (15) are zero.

An analogous argument shows

$$h_\ell^\infty(\theta, \lambda, h_e) \cdot \mathbb{P}(\mathbf{H}_\ell = h_\ell) = \mathbb{P}(\Theta = \theta, \mathbf{L} = \lambda, \mathbf{H}_e = h_e, \mathbf{H}_\ell = h_\ell). \quad (16)$$

Write $\phi(\theta, \lambda, h_e, h_\ell) := \mathbb{P}(\Theta = \theta, \mathbf{L} = \lambda, \mathbf{H}_e = h_e, \mathbf{H}_\ell = h_\ell)$. So, by Equations (15) and (16), $M_e \times M_\ell$ admits the common prior ϕ . Thus, $M_e \subset \text{CP}_e$ and $M_\ell \subset \text{CP}_\ell$. ■

To show that there is no loss in restricting our analysis to PBE with full path, we need to compare equilibria across different mechanisms. With this in mind, we introduce the notion of equivalence:

Definition C.1. A PBE (σ, β) of a supergame (\mathcal{M}, G) and a PBE (σ', β') of a supergame (\mathcal{M}', G) are **equivalent** if the following hold:

- (i) For each $\theta \in \Theta$, $U_e(\sigma \mid \theta, \{\emptyset\}, \beta_e) = U_e(\sigma' \mid \theta, \{\emptyset\}, \beta'_e)$ and $Y_e(\theta, \mathcal{M}, \sigma) = Y_e(\theta, \mathcal{M}', \sigma')$.
- (ii) $U_\ell(\sigma \mid \{\emptyset\}, \beta_\ell) = U_\ell(\sigma' \mid \{\emptyset\}, \beta'_\ell)$ and $Y_\ell(\mathcal{M}, \sigma) = Y_\ell(\mathcal{M}', \sigma')$.

Two PBE equilibria are equivalent if they induce the same agents' expected payoffs and the same expected transfers. So, in this sense the designer is indifferent between equivalent PBE. The following result shows that there is no loss by focusing on PBE with full path.

Lemma C.5. Fix a mechanism \mathcal{M} and a PBE (σ, β) of (\mathcal{M}, G) . There exist a mechanism \mathcal{M}' and a PBE (σ', β') of (\mathcal{M}', G) such that

- (i) (σ', β') has full path, and
- (ii) (σ', β') is equivalent to (σ, β) .

Proof of Lemma C.5. Fix a mechanism \mathcal{M} and an PBE (σ, β) of (\mathcal{M}, G) . The proof requires introducing some notation. Let Z be the terminal nodes of \mathcal{M} , let S_i be the pure strategies of $i \in \{e, \ell, c\}$ in \mathcal{M} , and let $\zeta : S \times V \rightarrow Z$ be the terminal-node mapping of \mathcal{M} . Write

$$Z^+ := \{z \in Z : P(z \mid \sigma) > 0\}$$

for the terminal nodes in \mathcal{M} that are reached under positive probability by σ . Write

$$S_e^+ := \{s_e \in S_e : P(s_e \mid \sigma_e, \theta) > 0 \text{ for some } \theta \in \Theta\}$$

for the expert's set of pure strategies that have positive probability under the strategy σ_e . Likewise, write $S_\ell^+ := \{s_\ell \in S_\ell : P(s_\ell \mid \sigma_\ell) > 0\}$ and $S_c^+ := \{s_c \in S_c : P(s_c \mid \sigma_c) > 0\}$.

The proof is divided in 4 steps. The first step constructs \mathcal{M}' ; the second constructs (σ', β') ; the third shows (σ', β') is a PBE of (\mathcal{M}', G) ; the last shows that (σ', β') has full path and that (σ', β') is equivalent to (σ, β) .

Step 1. We construct \mathcal{M}' . Set $V' := \{\emptyset\} \cup Z^+$ for the set of nodes of \mathcal{M}' . So, V' only contains the root and the terminal nodes of the original mechanism. Let \succsim' be the restriction of \succsim to the set $V' \subseteq V$ and attach this precedence relation to V' to form the tree (V', \succsim') . Set the action set

of \mathcal{M}' as $X' := S_e^+ \cup S_\ell^+ \cup S_c^+$. The information sets for i are $\mathcal{I}'_i := \{I_i \cap V' : I_i \in \mathcal{I}_i\}$. So, in this sense, the agents have the same information in \mathcal{M}' as if they were in mechanism \mathcal{M} conditional of being in V' .

Let \mathcal{T}'_i be the terminal information sets of i . We now describe the action correspondence $\mathcal{A}'_i : \mathcal{I}'_i \setminus \mathcal{T}'_i \rightrightarrows X$ of agent i . Set $\mathcal{A}'_i(\{\emptyset\}) := S_i^+$, so each agent i selects one of the pure strategies of the original mechanism that was played under positive probability. (Note, the root $\{\emptyset\}$ is the unique element of $\mathcal{I}'_i \setminus \mathcal{T}'_i$.) Set the strategy of chance as the strategy σ'_c such that $\sigma'_c(\{\emptyset\})(s_c) := P(s_c | \sigma_c)$ for each $s_c \in S_c^+$. So, the probability that chance chooses s_c in \mathcal{M}' is the same as in the original mechanism \mathcal{M} .

We now describe how the actions in \mathcal{M}' influence the reached terminal node. Write $S^+ := S_e^+ \times S_\ell^+ \times S_c^+$. Note, in the original mechanism \mathcal{M} , each pure strategy profile $s \in S^+$ leads to some terminal node $z = \zeta(s | \emptyset) \in Z^+$. Thus, set the terminal-node mapping $\zeta' : X' \times V' \rightarrow Z'$ of \mathcal{M}' as $\zeta'(s | \emptyset) = \zeta(s | \emptyset)$ for each $s \in S^+$. So, in this sense, each pure strategy profile s in \mathcal{M}' identifies the terminal node reached in \mathcal{M}' : the terminal node that would be realized on \mathcal{M} under profile s .

Finally we now describe the transfers and labels in \mathcal{M}' . Write $\gamma'_i = \gamma_i \circ \chi_i$ and set $L'(z) = L(z)$ for each $z \in Z^+$. In this sense, the transfers and the labels are assigned in the same way as the original mechanism \mathcal{M} . Finally, set $\mathcal{M}' = ((V', \succsim'), X', (\mathcal{I}'_i, \mathcal{A}'_i : i \in \{e, \ell, c\}), \sigma'_c, (\gamma'_i : i \in \{e, \ell\}), L')$.

Step 2. We now construct the assessment (σ', β') . To do so, first note that for each $I'_i \in \mathcal{I}'_i$ there is a unique $I_i \in \mathcal{I}_i$ such that $I'_i = I_i \cap V'$. Write $\chi_i(I'_i)$ for such terminal information set. So, the mapping $\chi_i : \mathcal{I}'_i \rightarrow \mathcal{I}_i$ identifies a each information set in \mathcal{M}' with the correspondent information set in the original mechanism \mathcal{M} . Set σ'_e as the expert's strategy such that

- (i) for each $(\theta, s_e) \in \Theta \times S_e^+$, $\sigma'_e(\theta, \emptyset)(s_e) = P(s_e | \theta, \sigma_e)$ and
- (ii) for each $(\theta, T'_e) \in \Theta \times \mathcal{T}'_e$, $\sigma'_e(\theta, T'_e) = \sigma_e(\theta, \chi_e(T'_e))$.

In this sense, the strategy σ'_e in the reduced mechanism \mathcal{M}' mimics the behavior of σ_e at \mathcal{M} (see (i)) and the behavior of σ_e after \mathcal{M} has concluded. (See (ii)). Likewise, set σ'_ℓ as the layman's strategy such that

- (i) for each $s_\ell \in S_\ell^+$, $\sigma'_\ell(\emptyset)(s_\ell) = P(s_\ell | \sigma_\ell)$, and
- (ii) for each $T'_\ell \in \mathcal{T}'_\ell$, $\sigma'_\ell(T'_\ell) = \sigma_\ell(\chi_\ell(T'_\ell))$.

We now construct β'_e . Fix (θ, I'_e) and write $I_e = \chi_e(I'_e)$. Observe, since β is consistent with σ , the belief $\beta_e(\theta, I_e)$ assigns positive probability only to nodes in $I_e \cap V'$. (The nodes in I_e that are reachable by σ). With this in mind, let $\beta'_e(\theta, I'_e) \in \Delta(V')$ be such that $\beta'_e(\theta, I'_e)(v) = \beta_e(\theta, I_e)(v)$ for each $v \in V'$. We now define β'_ℓ . Fix $I'_\ell \in \mathcal{I}'_\ell$ and write $I_\ell = \chi_\ell(I'_\ell)$. Observe, since (σ, β) is a consistent assessment, the belief $\beta_\ell(I_\ell)$ assigns positive probability only to terminal nodes in $I_\ell \cap V'$. With this in mind, let $\beta'_\ell(I'_\ell) \in \Delta(\Theta \times V')$ be such that $\beta'_\ell(I'_\ell)(\theta, v) = \beta_\ell(I_\ell)(\theta, v)$ for each $(\theta, v) \in \Theta \times V'$.

Step 3. We show β'_e is consistent with $(\sigma'_e, \sigma'_\ell)$. Fix $s_e \in S_e^+$ and write $P(\theta, v | s_e, \sigma'_\ell)$ for the ex-ante distribution of $(\theta, v) \in \Theta \times V'$ (in mechanism \mathcal{M}') under (s_e, σ'_ℓ) . Observe, by construction, $P(s_\ell | \sigma'_\ell) = P(s_\ell | \sigma_\ell)$ for each $s_\ell \in S_\ell^+$. So, $P(\theta, v | s_e, \sigma'_\ell) = P(\theta, v | s_e, \sigma_\ell)$. Thus, for each

$(\theta, I'_e) \in \Theta \times \mathcal{I}'_e$ and each $v \in I'_e$,

$$\begin{aligned} \beta'_e(\theta, I'_e)(v) \cdot P(\{\theta\} \times I'_e \mid s_e, \sigma'_\ell) &= \beta_e(\theta, \chi(I'_e))(v) \cdot P(\{\theta\} \times \chi(I'_e) \mid s_e, \sigma_\ell) \\ &= P(\theta, v \mid s_e, \sigma_\ell) \\ &= P(\theta, v \mid s_e, \sigma'_\ell), \end{aligned}$$

where the first equality follows from definition of β'_e and the second from consistency of (σ, β) . An analogous argument shows that β'_ℓ is consistent with $(\sigma'_e, \sigma'_\ell)$.

We now show that (σ', β') satisfies sequential rationality. Fix a strategy $\tilde{\sigma}'_e$ for e in (\mathcal{M}', G) . Write $\tilde{\sigma}_e$ for a strategy of e in (\mathcal{M}, G) such that $P(s_e \mid \tilde{\sigma}_e) = P(s_e \mid \tilde{\sigma}'_e)$ for each $s_e \in S_e^+$ and $\tilde{\sigma}_e(\theta, \bar{T}'_e) = \tilde{\sigma}_e(\theta, T'_e)$. So, $\tilde{\sigma}_e$ selects pure strategies in \mathcal{M} and actions in G with the same probability as $\tilde{\sigma}'_e$. Write $\mathcal{T}' = \mathcal{T}'_e \times \mathcal{T}'_\ell$. Fix $(\theta, I'_e) \in \Theta \times \mathcal{I}'_e$ and write $I_e = \chi_e(I'_e)$. For each $(T'_e, T'_\ell) \in \mathcal{T}'$, write $\chi(T'_e, T'_\ell) := (\chi_e(T'_e), \chi_\ell(T'_\ell))$. Notice that

$$\begin{aligned} U_e(\sigma'_e, \sigma'_\ell \mid \theta, I'_e, \beta'_e) &= \sum_{(T'_e, T'_\ell) \in \mathcal{T}'} [\Pi_e(\sigma'_e, \sigma'_\ell \mid \theta, T'_e, T'_\ell) + \gamma'_e(T'_e)] \cdot P(T'_e, T'_\ell \mid \sigma', \theta, I'_e, \beta_e) \\ &= \sum_{(T'_e, T'_\ell) \in \mathcal{T}'} [\Pi_e(\sigma_e, \sigma_\ell \mid \theta, \chi(T'_e, T'_\ell)) + \gamma_e(\chi_e(T'_e))] \cdot P(\chi(T'_e, T'_\ell) \mid \sigma, \theta, I_e, \beta_e) \\ &= U_e(\sigma_e, \sigma_\ell \mid \theta, I_e, \beta_e), \end{aligned} \tag{17}$$

where the second equality follows from the fact that $\gamma'_e = \gamma_e \circ \chi_e$ and the last equality from the fact that (σ_e, σ_ℓ) assigns positive probability only to the information sets T_i such that $T_i = \chi_i(T'_i)$ for some $T'_i \in \mathcal{T}'_i$.

An analogous argument shows that $U_e(\tilde{\sigma}'_e, \sigma'_\ell \mid \theta, I'_e, \beta'_e) = U_e(\tilde{\sigma}_e, \sigma_\ell \mid \theta, I_e, \beta_e)$. Therefore,

$$U_e(\sigma'_e, \sigma'_\ell \mid \theta, I'_e, \beta'_e) = U_e(\sigma_e, \sigma_\ell \mid \theta, I_e, \beta_e) \geq U_e(\tilde{\sigma}_e, \sigma_\ell \mid \theta, I_e, \beta_e) = U_e(\tilde{\sigma}'_e, \sigma'_\ell \mid \theta, I'_e, \beta'_e),$$

where the inequality follows from the fact that (σ, β) is sequentially rational. An analogous argument shows that (σ, β) is sequentially rational for ℓ . We conclude that (σ, β) is a PBE.

Step 4. Note that by construction, $P(z \mid \sigma') = P(z \mid \sigma) > 0$ for each $z \in Z^+$. So, (σ, β) has full path in \mathcal{M}' . Fix $\theta \in \Theta$. Observe, since $\gamma'_i = \gamma_i \circ \chi_i$, it follows that $Y_e(\theta, \mathcal{M}, \sigma) = Y_e(\theta, \mathcal{M}', \sigma')$ and $Y_\ell(\mathcal{M}, \sigma) = Y_\ell(\mathcal{M}', \sigma')$. In addition, Equation (17) implies $U_e(\sigma \mid \theta, \{\emptyset\}, \beta_e) = U_e(\sigma' \mid \theta, \{\emptyset\}, \beta_e)$. Moreover, an analogous argument shows that $U_\ell(\sigma \mid \{\emptyset\}, \beta_e) = U_\ell(\sigma' \mid \{\emptyset\}, \beta'_e)$. Hence (σ', β') is equivalent to (σ, β) . ■

Proof of Lemma 3.2. Fix a mechanism $\mathcal{M} = (\cdot, \mathcal{T}_e, \mathcal{T}_\ell)$ and a consistent assessment (σ, β) with full path in \mathcal{M} . The result follows from Proposition 1 in F&M. ■

Proof of Lemma 3.3. Write (δ_e, δ_ℓ) for the hierarchy mappings induced by β . Observe that there are some sets $M_e \subsetneq \text{CP}_e$ and $M_\ell \subsetneq \text{CP}_\ell$ that satisfies conditions (ii)-(iv) of Lemma 3.1. Note, since $M_e \times M_\ell$ admits a common prior, it is belief closed. So, the set $(\text{CP}_e \setminus M_e) \times (\text{CP}_\ell \setminus M_\ell)$ is also

belief closed. (See Lemma C.2.) Hence, the type structure $(\mathcal{T}_e, \mathcal{T}_\ell, \beta_e, \beta_\ell)$ induces a *decomposition* of the type structure \mathcal{S}_{CP} via the injective type morphisms (δ_e, δ_ℓ) . (See Definition 14 in F&M.) Moreover, since G is nice, the Bayesian game $\text{BG}(\mathcal{S}_{\text{CP}})$ has a Bayesian equilibrium. Therefore, $(\hat{\sigma}_e, \hat{\sigma}_\ell)$ can be extended to some Bayesian equilibrium $(\varsigma_e, \varsigma_\ell)$ of $\text{BG}(\mathcal{S}_{\text{CP}})$. (See Proposition 4 in F&M.) ■