

Symmetric Bayes Correlated Equilibrium: An Exact Characterization under Latent Information*

Sen Lu[‡]

April 2026

Abstract

This paper asks whether unrestricted Bayes correlated equilibrium (BCE) admits an exact lower-dimensional symmetric reduction in static games with latent information. I introduce Symmetric Bayes Correlated Equilibrium (SBCE), a reduced equilibrium object for symmetric environments with exchangeable retained information and symmetric conditionally independent additional information. Fixing the common payoff-relevant environment, the main result shows that the set of outcome distributions induced by symmetric Bayes Nash equilibrium, as one ranges over all admissible additional information structures in this class, coincides exactly with the set of outcome distributions induced by SBCE in the corresponding reduced game. Thus SBCE is the exact reduced equilibrium object for this symmetric latent information class, not merely a tractable restriction of BCE. The class nests complete information, private information, and intermediate noisy signal environments, including economically important settings such as entry and auctions. The result provides a single-game foundation for robust analysis of strategic behavior under weak informational assumptions in symmetric games.

Keywords: Bayes correlated equilibrium, Symmetric Bayes Correlated Equilibrium, latent information, symmetric games, exact reduction

JEL Codes: C72, D82, D83

1 Introduction

Static games with incomplete information depend critically on what players know about payoff-relevant uncertainty. In many economic environments, however, the underlying information structure is not observed by the analyst and is difficult to specify credibly. Existing work therefore typically proceeds either by imposing a benchmark informational

*I am grateful to Jeremy Fox, Xun Tang, Mallesh Pai, Isabelle Perrigne, Matthew Thirkettle, and seminar participants at Rice, IIOC, CRA and ESWC Seoul for helpful comments. All errors are mine.

[‡]Sen Lu: School of Insurance, Central University of Finance and Economics. senluxem@gmail.com

assumption, such as complete or private information, or by using Bayes correlated equilibrium (BCE) as a reduced-form characterization of equilibrium outcomes under latent information (Bergemann and Morris, 2016). While BCE delivers a fully general characterization, its unrestricted form becomes unwieldy in symmetric environments with a moderate or large number of players. This raises the central question of the paper: within a sufficiently rich symmetric class, does unrestricted BCE admit an exact lower-dimensional symmetric reduction?

This paper gives an affirmative answer for a class of symmetric static games with exchangeable retained information and symmetric conditionally independent additional information. On the original game side, symmetry requires not only common payoff primitives and exchangeable retained information, but also symmetric equilibrium behavior: players use a common decision rule, and strategies depend on opponent-indexed additional information only up to relabeling. I introduce *Symmetric Bayes Correlated Equilibrium* (SBCE), the reduced equilibrium object associated with this class. Fix a common payoff-relevant environment, consisting of the payoff primitives and the marginal distribution of retained information and public states. The main theorem shows that the set of outcome distributions induced by symmetric Bayes Nash equilibrium, as one ranges over all admissible additional information structures in this class, coincides exactly with the set of outcome distributions induced by SBCE in the corresponding reduced game. SBCE is therefore the exact reduced equilibrium object for this symmetric latent information class, rather than merely a tractable restriction of BCE.

The class is broad, nesting complete information, private information, and intermediate noisy signal environments. It also aligns closely with the anonymous and exchangeable structures that have become standard in many-player empirical work, where tractability typically requires payoffs and behavior to depend on a player's own payoff-relevant type and on an aggregate description of opponents rather than on their identities (Menzel, 2016). The conditional independence restriction is substantive, but exactness here is not driven by an artificially narrow special case. The theorem shows that a broad and empirically central class of symmetric games admits an exact lower-dimensional reduced-form characterization when the common payoff-relevant environment is held fixed and latent additional information structures are allowed to vary within the class.

This exact characterization matters at the conceptual level and also provides the foundation for subsequent analysis. It identifies the correct reduced equilibrium object for symmetric games with latent information, rather than a computational shortcut whose relation to the original game is merely approximate. More broadly, it fits into a line of work showing that analysts can often characterize behavior through reduced equilibrium restrictions, rather than attempting to recover the full, and typically infinite-dimensional, information structure itself (Bergemann and Morris, 2016; Bergemann et al., 2022). The paper identifies the reduced object, establishes the exact equivalence, and clarifies the scope of the class for which exactness obtains.

A natural next question is whether the same symmetric latent information class is preserved under linked game constructions for joint prediction and counterfactual analysis under fixed information, in the spirit of Bergemann et al. (2022). That extension goes beyond the single game analysis of this paper, because it requires additional linked game objects and results. A companion paper develops that extension for the same symmetric class. The focus here is the single game result on which that broader linked game analysis builds.

This paper is most directly related to Bergemann and Morris (2016), who identify BCE as the fully general reduced-form object for equilibrium analysis under latent information and characterize the associated ordering of information structures through individual sufficiency. Within a symmetric latent information class, SBCE is the exact lower-dimensional reduced object.

The paper is also related to the broader literature on correlated equilibrium in games with incomplete information. Following the complete information notion of correlated equilibrium in Aumann (1974, 1987), a number of correlated equilibrium concepts for games with incomplete information have been studied, notably by Forges (1993, 2006) and Liu (2015). That literature asks what informational restrictions should be imposed on the recommendation device, and thus what forms of correlation among recommendations, types, and states are admissible in the reduced object. The approach here is different: rather than refining the informational restrictions imposed on a general correlated-equilibrium object, the paper restricts attention to a symmetric latent-information class. Within that class, the reduced object on the reduced-game side is more structured than an unrestricted BCE:

it is a symmetric independent recommendation rule. The contribution is not a new general correlated equilibrium notion, but an exact characterization result for that class.

Finally, the paper is related to subsequent work that uses BCE-based reasoning to study prediction, counterfactual analysis, and econometric identification under weak assumptions on information.¹ Bergemann et al. (2022) study joint predictions and counterfactuals under latent but fixed information through linked games, while recent econometric work uses BCE to study identification, inference, and estimation in games, auctions, and discrete choice environments (Magnolfi and Roncoroni, 2023; Syrgkanis et al., 2021; Gualdani and Sinha, 2024). The paper gives an exact characterization for a symmetric latent information class: it identifies the reduced equilibrium object and establishes the equivalence on which those further developments can build.

The remainder of the paper is organized as follows. Section 2 introduces the class of symmetric static games studied in the paper and formalizes the maintained assumptions on symmetry and latent information. Section 3 defines Symmetric Bayes Correlated Equilibrium as the reduced equilibrium object of interest. Section 4 states and proves the main exact equivalence theorem. Section 5 illustrates the maintained class through canonical economic environments. Section 6 discusses the scope of the result and its relation to the broader literature on games with latent information. Section 7 concludes the paper.

2 Model

We begin with the original static game and its information structure. Each player’s information has two components: retained payoff-relevant information and an additional latent information component. The retained component is carried into the reduced game, whereas the additional component captures further information that may be present in the original game but is not directly represented in the reduced object. The maintained class studied in the paper is defined by symmetric payoffs, exchangeability of retained information, and symmetric conditional independence of the additional information layer.

¹Related work has also extended BCE-based robustness analysis to dynamic environments. See, for example, Hara et al. (2025), who study dynamic games with unknown information structure through the notion of Markov correlated equilibrium.

2.1 Players, actions, states, and information

There is a finite set of players, indexed by $i = 1, \dots, N$, a finite set of states Θ , and a common finite action set A . Let

$$\mathcal{A} := A^N$$

denote the set of action profiles.

The state $\theta \in \Theta$ is publicly observed before actions are chosen and is therefore common knowledge. Each player i privately observes a pair (t_i, s_i) , where $t_i \in T$ is retained payoff-relevant information and $s_i \in S_i$ is an additional information component. Let

$$t = (t_1, \dots, t_N) \in T^N$$

denote the retained information profile, and let

$$s = (s_1, \dots, s_N) \in \mathcal{S}, \quad \mathcal{S} := \prod_{i=1}^N S_i,$$

denote the additional information profile and its space.

The space S_i may be finite- or infinite-dimensional. No finite-dimensional restriction is imposed on this component. In particular, one may think of s_i as an opponent-indexed collection $s_i = (s_{i,k})_{k \neq i}$, where each component $s_{i,k}$ may itself be vector-valued or belong to a general measurable space. This representation is used only to index player i 's additional information and entails no loss of generality.

Payoffs are given by functions

$$u_i : \mathcal{A} \times T^N \times \Theta \rightarrow \mathbb{R}.$$

Thus player i 's payoff may depend on the action profile, the retained information profile, and the payoff-relevant state.

Uncertainty is described by a common-prior information structure

$$I = (T, (S_i)_{i=1}^N, \pi),$$

where

$$\pi \in \Delta(T^N \times \mathcal{S} \times \Theta)$$

is the joint distribution of retained information, additional information, and states. Ex ante, π describes the common prior over (t, s, θ) . Ex post, after (t, s, θ) is realized, the

state θ is publicly observed and common knowledge, whereas (t_i, s_i) remains private to player i .

A game is summarized by the pair $(\mathcal{G}, I)$, where

$$\mathcal{G} = ((u_i)_{i=1}^N, A, \Theta)$$

collects the payoff-relevant primitives and I describes the information structure.

2.2 Symmetry

We now impose the symmetry restriction that is central to the analysis. For any permutation φ of $\{1, \dots, N\}$, define

$$\varphi(a) := (a_{\varphi(1)}, \dots, a_{\varphi(N)}), \quad \varphi(t) := (t_{\varphi(1)}, \dots, t_{\varphi(N)}).$$

The game is symmetric if players are ex ante identical up to relabeling. Because retained information may affect payoffs directly, symmetry is imposed jointly on player labels, action profiles, and retained information profiles. Formally, for every permutation φ of $\{1, \dots, N\}$, every player i , every action profile $a \in \mathcal{A}$, every retained information profile $t \in T^N$, and every state $\theta \in \Theta$,

$$u_i(a, t, \theta) = u_{\varphi(i)}(\varphi(a), \varphi(t), \theta).$$

Equivalently, the payoff structure is invariant under a consistent relabeling of players, actions, and retained information.

2.3 Exchangeable retained information and symmetric conditionally independent additional information

We next impose structure on the informational environment introduced above. The retained component t_i is the payoff-relevant information carried into the reduced game, whereas the additional component s_i captures further information that may be present in the original game but is not directly represented in the reduced object.

Let π^1 denote the marginal of π on $T^N \times \Theta$. For any permutation φ of $\{1, \dots, N\}$ and any retained information profile $t = (t_1, \dots, t_N) \in T^N$, recall that $\varphi(t) := (t_{\varphi(1)}, \dots, t_{\varphi(N)})$.

Exchangeable retained information. The retained information structure is *exchangeable* if

$$\pi^1(\varphi(t), \theta) = \pi^1(t, \theta) \quad \text{for all } (t, \theta) \in T^N \times \Theta$$

and every permutation φ of $\{1, \dots, N\}$.

Thus, the ex ante distribution of retained payoff-relevant information is invariant under relabeling of players.

We next impose restrictions on the additional information component s_i . For each player i , S_i is a measurable space and may be infinite dimensional. In particular, one may think of $s_i = (s_{i,k})_{k \neq i}$, where $s_{i,k}$ denotes the part of player i 's additional information associated with opponent k . This opponent-indexed representation is used only to express conditional independence and symmetry and entails no loss of generality.

For each player i , let ϕ_{-i} denote a permutation of player i 's opponents. This permutation acts naturally on the opponents' retained information profile t_{-i} and, under the opponent-indexed representation above, on player i 's additional information s_i by relabeling opponent indices.

Homogeneous conditional independence of additional information. The additional information structure is *homogeneous conditionally independent* if there exists a common stochastic kernel

$$\rho(\cdot \mid t_i, t_{-i}, \theta) \in \Delta(S_i)$$

such that, for every retained information profile $t \in T^N$, every state $\theta \in \Theta$, and every measurable rectangle

$$B = \prod_{i=1}^N B_i \subseteq \prod_{i=1}^N S_i,$$

one has

$$\pi(B \mid t, \theta) = \prod_{i=1}^N \rho(B_i \mid t_i, t_{-i}, \theta).$$

Thus, conditional on (t, θ) , the additional information components are mutually independent across players and generated by a common kernel. The maintained restriction is therefore stronger than mere conditional independence: it combines factorization across players with a common conditional law.

Symmetric additional information structure. The additional information structure is *symmetric* if it is homogeneous conditionally independent and the common kernel ρ is invariant to permutations of opponents. Formally, for every player i , every own retained information value $t_i \in T$, every opponents' retained information profile $t_{-i} \in T^{N-1}$, every state $\theta \in \Theta$, every permutation ϕ_{-i} of player i 's opponents, and every measurable set $B_i \subseteq S_i$,

$$\rho(B_i \mid t_i, t_{-i}, \theta) = \rho(\phi_{-i}(B_i) \mid t_i, \phi_{-i}(t_{-i}), \theta),$$

where $\phi_{-i}(B_i)$ denotes the image of B_i under the relabeling of opponent-indexed coordinates induced by ϕ_{-i} .

This restriction requires that the conditional law of player i 's additional information depends on opponents' retained information only through its configuration up to relabeling. Equivalently, players with the same own retained information and the same opponents' retained information profile up to permutation are assigned additional information according to the same conditional law.

Accordingly, the maintained informational restrictions are as follows: the retained payoff-relevant component t is exchangeable across players, while the additional information component s is conditionally independent across players given (t, θ) and generated by a common signal kernel that is symmetric with respect to permutations of opponents.

Both the equilibrium object for the original game and the recommendation object for the reduced game are formulated in terms of retained information, while allowing the original game to contain additional information through the component s .

This maintained class accommodates several benchmark cases. Independent private information is obtained when the additional information component is degenerate and carries no information beyond the retained component. Complete information is obtained when the additional information component fully reveals opponents' retained information. Intermediate cases are also allowed: players may observe noisy signals about opponents' retained information, and the precision of those signals may depend on their own retained information. Thus, symmetry does not require identical ex post informational precision across players. It requires only a common ex ante signal generation rule, while allowing realized retained information to induce heterogeneous informational advantages ex post.

2.4 Symmetric Bayes Nash equilibrium

In the original game, player i observes the full information triple

$$(t_i, s_i, \theta) \in T \times S_i \times \Theta$$

before choosing an action. A (mixed) strategy for player i is a measurable mapping

$$\sigma_i : T \times S_i \times \Theta \rightarrow \Delta(A).$$

A strategy profile is denoted by

$$\sigma = (\sigma_1, \dots, \sigma_N).$$

Given $(t, s, \theta) \in T^N \times S \times \Theta$, the profile σ induces the product distribution

$$\sigma(a \mid t, s, \theta) := \prod_{i=1}^N \sigma_i(a_i \mid t_i, s_i, \theta), \quad a = (a_1, \dots, a_N) \in \mathcal{A}.$$

We next formalize symmetry of the strategy profile. For each player i , let ϕ_{-i} denote a permutation of player i 's opponents, acting on the opponent-indexed coordinates of s_i . A strategy profile σ is *symmetric* if it satisfies the following two conditions.

- (i) *Opponent-relabeling invariance within each player.* For every player i , every own retained information value $t_i \in T$, every additional information realization $s_i \in S_i$, every state $\theta \in \Theta$, every action $a_i \in A$, and every permutation ϕ_{-i} of player i 's opponents,

$$\sigma_i(a_i \mid t_i, s_i, \theta) = \sigma_i(a_i \mid t_i, \phi_{-i}(s_i), \theta).$$

- (ii) *Common decision rule across players.* For any two players i and j , let $\psi_{ij} : S_i \rightarrow S_j$ denote the natural relabeling of opponent-indexed coordinates induced by the relabeling of player identities from i to j . Then, for every $t_i \in T$, every $s_i \in S_i$, every $\theta \in \Theta$, and every $a \in A$,

$$\sigma_i(a \mid t_i, s_i, \theta) = \sigma_j(a \mid t_i, \psi_{ij}(s_i), \theta).$$

Equivalently, all players use the same conditional law of action, up to relabeling of player identities and corresponding opponent-indexed coordinates.

Player i 's expected payoff under σ is

$$U_i(\sigma) = \sum_{\theta \in \Theta} \sum_{t \in T^N} \int_{s \in S} \sum_{a \in \mathcal{A}} u_i(a, t, \theta) \sigma(a \mid t, s, \theta) \pi(ds, t, \theta).$$

A symmetric strategy profile σ is a *symmetric Bayes Nash equilibrium* if, for every player i and every measurable deviation

$$\tilde{\sigma}_i : T \times S_i \times \Theta \rightarrow \Delta(A),$$

one has

$$U_i(\sigma) \geq U_i(\sigma_{-i}, \tilde{\sigma}_i).$$

2.5 SBNE-induced distributions

We now separate the common payoff-relevant environment from the additional information structure. Let

$$\mathcal{E} := (\mathcal{G}, \pi^1),$$

where

$$\mathcal{G} = ((u_i)_{i=1}^N, \mathcal{A}, \Theta)$$

collects the payoff-relevant game primitives, and

$$\pi^1 \in \Delta(T^N \times \Theta)$$

is the marginal distribution of retained information and public states.

Let $\mathcal{I}(\mathcal{E})$ denote the class of admissible original information structures consistent with $\mathcal{E}$ and the maintained assumptions of Section 2.3. Thus the common payoff-relevant environment is held fixed, while the additional information structure is allowed to vary within $\mathcal{I}(\mathcal{E})$.

Given $\mathcal{E}$, an admissible information structure $I \in \mathcal{I}(\mathcal{E})$, and a symmetric Bayes Nash equilibrium σ of the corresponding original game, define the induced joint distribution over action profiles, retained information profiles, and states by

$$\mu_{\sigma, I}(a, t, \theta) := \int_{s \in S} \sigma(a \mid t, s, \theta) \pi(ds, t, \theta), \quad (a, t, \theta) \in \mathcal{A} \times T^N \times \Theta.$$

We refer to $\mu_{\sigma, I}$ as an *SBNE-induced distribution*, and let

$$\mathcal{M}^{SBNE}(\mathcal{E}) := \left\{ \begin{array}{l} \mu_{\sigma, I} \in \Delta(\mathcal{A} \times T^N \times \Theta) : \\ I \in \mathcal{I}(\mathcal{E}), \sigma \text{ is a symmetric Bayes Nash equilibrium under } I \end{array} \right\}$$

denote the set of all such induced distributions.

The next section introduces the reduced solution concept that exactly characterizes the set $\mathcal{M}^{SBNE}(\mathcal{E})$ within the maintained class.

3 Symmetric Bayes Correlated Equilibrium

We now turn to the reduced equilibrium object for a fixed common payoff-relevant environment. Throughout this section, fix the same environment

$$\mathcal{E} := (\mathcal{G}, \pi^1)$$

introduced in Section 2.5. On the reduced game side, $\mathcal{E}$ is held fixed and the additional information component is absent. The reduced object is defined directly from retained information and the fixed environment.

3.1 Symmetric independent recommendation rules

For each player i and each retained information profile $t = (t_1, \dots, t_N) \in T^N$, let

$$n_{-i}(t_{-i}) \in \mathbb{Z}_+^{|T|}$$

denote the vector of opponents' retained information counts, whose τ -th component is

$$n_{-i}^\tau(t_{-i}) := \sum_{j \neq i} \mathbf{1}\{t_j = \tau\}, \quad \tau \in T.$$

Thus, $n_{-i}(t_{-i})$ is an anonymous summary of opponents' retained information.

Let

$$\mathcal{H} := \left\{ m \in \mathbb{Z}_+^{|T|} : \sum_{\tau \in T} m^\tau = N - 1 \right\}$$

denote the set of feasible opponents' count vectors. A *symmetric independent recommendation rule* is a stochastic kernel

$$\nu(\cdot \mid t_i, m, \theta) \in \Delta(A), \quad (t_i, m, \theta) \in T \times \mathcal{H} \times \Theta,$$

common to all players.

Given $t \in T^N$ and $\theta \in \Theta$, the rule ν induces the conditional distribution

$$\nu(a \mid t, \theta) = \prod_{i=1}^N \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta), \quad a \in \mathcal{A}.$$

Hence recommendations are conditionally independent across players given (t, θ) and depend on opponents' retained information only through its anonymous count vector.

Given $\mathcal{E}$ and ν , define the induced joint distribution over action profiles, retained information, and states by

$$\mu_{\nu, \mathcal{E}}(a, t, \theta) := \pi^1(t, \theta) \nu(a \mid t, \theta) \quad \text{for } (a, t, \theta) \in \mathcal{A} \times T^N \times \Theta.$$

We refer to $\mu_{\nu, \mathcal{E}}$ as the reduced game distribution induced by $\mathcal{E}$ and ν .

3.2 Bayes consistency

For a fixed environment $\mathcal{E} = (\mathcal{G}, \pi^1)$, Bayes consistency requires that a recommendation rule preserve the exogenous joint distribution of retained information and states encoded in π^1 . Formally, a joint distribution

$$\mu \in \Delta(\mathcal{A} \times T^N \times \Theta)$$

satisfies *Bayes consistency* if its marginal on $T^N \times \Theta$ is π^1 , that is,

$$\sum_{a \in \mathcal{A}} \mu(a, t, \theta) = \pi^1(t, \theta), \quad \forall (t, \theta) \in T^N \times \Theta.$$

Every distribution induced by a symmetric independent recommendation rule satisfies Bayes consistency by construction, and so Bayes consistency is not imposed separately in the definition of SBCE.

3.3 Obedience

Obedience is imposed conditional on the information available to a player at the recommendation stage. In the present environment, player i observes retained information $t_i \in T$, the recommended action $a_i \in A$, and the publicly observed state $\theta \in \Theta$. Accordingly, a joint distribution

$$\mu \in \Delta(\mathcal{A} \times T^N \times \Theta)$$

satisfies *obedience* if, for every player i , every retained information value $t_i \in T$, every recommended action $a_i \in A$, every state $\theta \in \Theta$, and every alternative action $a'_i \in A$ such that

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu(a_i, a_{-i}, t_i, t_{-i}, \theta) > 0,$$

one has

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu(a_i, a_{-i}, t_i, t_{-i}, \theta) \left[u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta) \right] \geq 0.$$

These inequalities require that, conditional on (t_i, a_i, θ) , player i does not gain from deviating to a'_i . We keep this full playerwise statement of obedience.

3.4 Definition of SBCE

We now define the reduced equilibrium object.

A *symmetric Bayes correlated equilibrium* (SBCE) recommendation rule for a fixed environment $\mathcal{E}$ is a symmetric independent recommendation rule ν whose induced joint distribution $\mu_{\nu, \mathcal{E}}$ satisfies obedience. Equivalently, an *SBCE-induced distribution* is a distribution

$$\mu \in \Delta(\mathcal{A} \times T^N \times \Theta)$$

induced by a symmetric independent recommendation rule for $\mathcal{E}$ and satisfying obedience.

Let

$$\mathcal{M}^{SBCE}(\mathcal{E}) := \{ \mu_{\nu, \mathcal{E}} \in \Delta(\mathcal{A} \times T^N \times \Theta) : \nu \text{ is an SBCE recommendation rule for } \mathcal{E} \}$$

denote the set of SBCE-induced distributions associated with $\mathcal{E}$.

Every element of $\mathcal{M}^{SBCE}(\mathcal{E})$ is therefore induced by a symmetric independent recommendation rule, is Bayes consistent with π^1 , and is obedient. The main theorem in the next section shows that

$$\mathcal{M}^{SBCE}(\mathcal{E}) = \mathcal{M}^{SBNE}(\mathcal{E}).$$

4 Characterization of SBNE Outcomes by SBCE

Throughout this section, fix a common payoff-relevant environment

$$\mathcal{E} = (\mathcal{G}, \pi^1),$$

where $\mathcal{G} = ((u_i)_{i=1}^N, \mathcal{A}, \Theta)$ collects the payoff relevant game primitives and π^1 is the marginal distribution of retained signals and public states. Assume that $\mathcal{G}$ is symmetric and that π^1 is exchangeable, so that $\mathcal{E}$ belongs to the maintained symmetric class.

On the original game side, the additional information structure is allowed to vary over the admissible class $\mathcal{I}(\mathcal{E})$ defined in Section 2.5. The question is whether the set of symmetric equilibrium outcomes generated as the additional information structure varies over $\mathcal{I}(\mathcal{E})$ is exactly captured by SBCE. The main result shows that it is: that set coincides exactly with the set of SBCE-induced distributions in the corresponding reduced game.

Theorem 1 (Characterization of SBNE outcomes). *For every fixed common payoff-relevant environment $\mathcal{E} = (\mathcal{G}, \pi^1)$, the set of SBNE-induced distributions in the original game coincides with the set of SBCE-induced distributions in the reduced game:*

$$\mathcal{M}^{SBNE}(\mathcal{E}) = \mathcal{M}^{SBCE}(\mathcal{E}).$$

The theorem shows that SBCE provides an exact reduced characterization of symmetric equilibrium outcomes for a fixed common payoff-relevant environment. Equivalently, a distribution

$$\mu \in \Delta(\mathcal{A} \times T^N \times \Theta)$$

belongs to $\mathcal{M}^{SBNE}(\mathcal{E})$ if and only if it belongs to $\mathcal{M}^{SBCE}(\mathcal{E})$.

The equivalence is exact in both directions. On the original game side, $\mathcal{M}^{SBNE}(\mathcal{E})$ is obtained by holding fixed the common environment $\mathcal{E}$ and ranging over all admissible additional information structures in $\mathcal{I}(\mathcal{E})$. On the reduced game side, $\mathcal{M}^{SBCE}(\mathcal{E})$ is defined directly from the same fixed environment $\mathcal{E}$. Thus SBCE is not merely a tractable approximation, a relaxation, or an outer bound. It exactly characterizes the original symmetric equilibrium outcome set within the maintained class.

The nontrivial content of the theorem is that the additional information component s need not be tracked explicitly in the reduced-form characterization of outcomes, even though it may affect behavior in the original game. For the class studied here, its equilibrium effect can be absorbed into a symmetric recommendation rule defined on the retained information and the fixed environment $\mathcal{E}$. This reduction rests on two features of the model: first, the retained component (t, θ) , together with the payoff primitives in $\mathcal{G}$, is sufficient to index the common payoff-relevant environment; second, the additional information component enters through a symmetric conditionally independent structure, so its effect can be represented through symmetric independent recommendation rules that de-

pend on opponents' retained information only through the anonymous configuration of that information.

4.1 Supporting lemmas for $\mathcal{M}^{SBCE}(\mathcal{E}) \subseteq \mathcal{M}^{SBNE}(\mathcal{E})$

Lemma 1 (Canonical additional information structure). *Let ν be a symmetric independent recommendation rule for the fixed environment $\mathcal{E}$. For each player i , let*

$$S_i^* := A^{N-1},$$

viewed as an opponent-indexed collection and equipped with the natural action of permutations of player i 's opponents by coordinate relabeling. Let

$$\kappa_i : A \rightarrow S_i^*$$

be the measurable injective map given by

$$\kappa_i(a) := (a, \dots, a) \in A^{N-1}.$$

For each $(t, \theta) \in T^N \times \Theta$, define the conditional law of player i 's additional information by

$$\eta_i^\nu(B_i \mid t, \theta) := \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in B_i\} \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta), \quad B_i \subseteq S_i^* \text{ measurable},$$

and let

$$\Pi^\nu(\cdot \mid t, \theta) := \bigotimes_{i=1}^N \eta_i^\nu(\cdot \mid t, \theta) \in \Delta\left(\prod_{i=1}^N S_i^*\right).$$

Then the family $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ defines an admissible original information structure

$$I^\nu \in \mathcal{I}(\mathcal{E}).$$

In particular, the additional information layer of I^ν is conditionally independent given (t, θ) and symmetric across players.

Lemma 2 (Implementation of the recommendation rule). *Let ν be a symmetric independent recommendation rule for the fixed environment $\mathcal{E}$, and let $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ be the canonical additional information structure constructed in Lemma 1. Fix an arbitrary reference distribution $q^\circ \in \Delta(A)$. For each player i , define the canonical strategy*

$$\sigma_i^\nu(\cdot \mid t_i, s_i, \theta) := \begin{cases} \delta_{\kappa_i^{-1}(s_i)}, & \text{if } s_i \in \kappa_i(A), \\ q^\circ, & \text{if } s_i \notin \kappa_i(A), \end{cases} \quad (t_i, s_i, \theta) \in T \times S_i^* \times \Theta,$$

where δ_a denotes the degenerate distribution on A that assigns probability one to a . Let

$$\sigma^\nu(a \mid t, s, \theta) := \prod_{i=1}^N \sigma_i^\nu(a_i \mid t_i, s_i, \theta)$$

denote the induced conditional distribution over action profiles. Then, for every

$$(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta,$$

one has

$$\int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta) = \nu(a \mid t, \theta),$$

where

$$\nu(a \mid t, \theta) := \prod_{i=1}^N \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta).$$

Consequently, the joint distribution on $\mathcal{A} \times T^N \times \Theta$ induced by the original game construction $(\pi^1, \Pi^\nu, \sigma^\nu)$ is precisely $\mu_{\nu, \mathcal{E}}$.

Lemma 3 (Optimality of the canonical strategy). *Let ν be an SBCE recommendation rule for the fixed environment $\mathcal{E}$. Let $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ be the canonical additional information structure constructed in Lemma 1, and let σ^ν be the canonical strategy profile defined in Lemma 2. Then σ^ν is a symmetric Bayes Nash equilibrium under the admissible original information structure*

$$I^\nu \in \mathcal{I}(\mathcal{E}).$$

4.2 Supporting lemmas for $\mathcal{M}^{SBNE}(\mathcal{E}) \subseteq \mathcal{M}^{SBCE}(\mathcal{E})$

Lemma 4 (Anonymous aggregation and induced recommendation rule). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, and let σ be a symmetric Bayes Nash equilibrium under I . For each player i , define*

$$\gamma_i(t_i, m, \theta) := \sum_{t_{-i} \in T^{N-1}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta), \quad (t_i, m, \theta) \in T \times \mathcal{H} \times \Theta.$$

Then there exists a common anonymous marginal distribution

$$\gamma \in \Delta(T \times \mathcal{H} \times \Theta)$$

such that

$$\gamma(t_i, m, \theta) = \gamma_i(t_i, m, \theta)$$

for every player i .

Fix an arbitrary reference distribution $q^\circ \in \Delta(A)$. Using the anonymous marginal distribution γ and the equilibrium strategy profile σ , define for each player i

$$\nu_i^\sigma(a_i \mid t_i, m, \theta) := \frac{\sum_{t_{-i} \in T^{N-1}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta) \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)}{\gamma(t_i, m, \theta)},$$

whenever $\gamma(t_i, m, \theta) > 0$, and define

$$\nu_i^\sigma(\cdot \mid t_i, m, \theta) := q^\circ$$

whenever $\gamma(t_i, m, \theta) = 0$.

Then there exists a common rule

$$\nu^\sigma(\cdot \mid t_i, m, \theta) \in \Delta(A)$$

such that

$$\nu_i^\sigma(\cdot \mid t_i, m, \theta) = \nu^\sigma(\cdot \mid t_i, m, \theta) \quad \text{for every player } i \text{ and every } (t_i, m, \theta) \in T \times \mathcal{H} \times \Theta,$$

and the induced conditional distribution over action profiles is given by

$$\nu^\sigma(a \mid t, \theta) := \prod_{i=1}^N \nu_i^\sigma(a_i \mid t_i, n_{-i}(t_{-i}), \theta), \quad a \in \mathcal{A}.$$

Thus ν^σ defines a symmetric independent recommendation rule in the sense of Section 3.1.

Lemma 5 (Reproduction of the induced distribution). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, let σ be a symmetric Bayes Nash equilibrium under I , and let ν^σ be the symmetric independent recommendation rule constructed in Lemma 4. Then, for every $(t, \theta) \in T^N \times \Theta$ such that $\pi^1(t, \theta) > 0$,*

$$\nu^\sigma(a \mid t, \theta) = \frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} \quad \text{for all } a \in \mathcal{A}.$$

Consequently, the joint distribution on $\mathcal{A} \times T^N \times \Theta$ induced by the reduced game construction (π^1, ν^σ) is precisely $\mu_{\sigma, I}$.

Lemma 6 (Obedience of the induced recommendation rule). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, and let σ be a symmetric Bayes Nash equilibrium under I . Let*

$$\nu^\sigma(\cdot \mid t_i, m, \theta) \in \Delta(A)$$

be the induced recommendation rule constructed in Lemma 4. Then the induced joint distribution $\mu_{\nu^\sigma, \mathcal{E}}$ satisfies obedience.

4.3 Proof of Theorem 1

Proof. We prove the two set inclusions separately.

Step 1: $\mathcal{M}^{SBCE}(\mathcal{E}) \subseteq \mathcal{M}^{SBNE}(\mathcal{E})$. Take an arbitrary $\mu \in \mathcal{M}^{SBCE}(\mathcal{E})$. By definition of $\mathcal{M}^{SBCE}(\mathcal{E})$, there exists an SBCE recommendation rule ν for the fixed environment $\mathcal{E}$ such that $\mu = \mu_{\nu, \mathcal{E}}$.

By Lemma 1, the rule ν induces an admissible original information structure $I^\nu \in \mathcal{I}(\mathcal{E})$. By Lemma 2, the canonical strategy profile σ^ν satisfies $\mu_{\sigma^\nu, I^\nu} = \mu_{\nu, \mathcal{E}}$. By Lemma 3, σ^ν is a symmetric Bayes Nash equilibrium under I^ν . Hence $\mu = \mu_{\sigma^\nu, I^\nu} \in \mathcal{M}^{SBNE}(\mathcal{E})$. Since $\mu \in \mathcal{M}^{SBCE}(\mathcal{E})$ was arbitrary, it follows that

$$\mathcal{M}^{SBCE}(\mathcal{E}) \subseteq \mathcal{M}^{SBNE}(\mathcal{E}).$$

Step 2: $\mathcal{M}^{SBNE}(\mathcal{E}) \subseteq \mathcal{M}^{SBCE}(\mathcal{E})$. Take an arbitrary $\mu \in \mathcal{M}^{SBNE}(\mathcal{E})$. By definition of $\mathcal{M}^{SBNE}(\mathcal{E})$, there exist an admissible original information structure $I \in \mathcal{I}(\mathcal{E})$ and a symmetric Bayes Nash equilibrium σ under I such that $\mu = \mu_{\sigma, I}$.

By Lemma 4, the pair (I, σ) induces a symmetric independent recommendation rule ν^σ . By Lemma 6, the induced joint distribution $\mu_{\nu^\sigma, \mathcal{E}}$ satisfies obedience. Then, by the definition of SBCE in Section 3.4, it follows that ν^σ is an SBCE recommendation rule for the fixed environment $\mathcal{E}$.

By Lemma 5, $\mu_{\nu^\sigma, \mathcal{E}} = \mu_{\sigma, I} = \mu$. Therefore $\mu \in \mathcal{M}^{SBCE}(\mathcal{E})$. Since $\mu \in \mathcal{M}^{SBNE}(\mathcal{E})$ was arbitrary, it follows that

$$\mathcal{M}^{SBNE}(\mathcal{E}) \subseteq \mathcal{M}^{SBCE}(\mathcal{E}).$$

Combining the two inclusions yields

$$\mathcal{M}^{SBNE}(\mathcal{E}) = \mathcal{M}^{SBCE}(\mathcal{E}).$$

□

5 Canonical Environments within the Maintained Class

This section illustrates the economic breadth of the maintained class for which the exact characterization is obtained. Although Theorem 1 is derived under symmetry and

conditional-independence restrictions, the class is not a narrow technical corner case. It nests standard benchmark information environments, including complete information, private information, and intermediate noisy-signal structures, and it includes canonical symmetric settings familiar in economic applications. The section therefore clarifies the substantive content of the environment for which SBCE provides an exact reduced-form characterization.

Entry. Consider a symmetric entry environment with fixed payoff-relevant primitives $\mathcal{G}^e$ and an admissible original information structure $I^e \in \mathcal{I}(\mathcal{E}^e)$. Each firm chooses whether to enter a market. Let firm i 's retained information be $t_i \in T$, its additional information be $s_i \in S_i$, and let $\theta \in \Theta$ denote the public state.

The maintained class allows richer payoff dependence than formulations in which profits depend only on the action profile and firm i 's own private type. In particular, entry payoffs may depend symmetrically on the full retained-information profile $t = (t_i, t_{-i})$, so that opponents' retained information may affect firm i 's payoff through a symmetric function of the profile. The retained information t_i may itself be multidimensional, and some components may govern the precision of the additional-information layer s_i about opponents' retained information.

Thus, the theorem applies to symmetric entry environments with rich retained heterogeneity and additional noisy information, not only to complete- or private-information cases.

Auctions. Consider next a symmetric auction environment with fixed payoff-relevant primitives $\mathcal{G}^a$ and an admissible original information structure $I^a \in \mathcal{I}(\mathcal{E}^a)$, with a common bid set A . Let player i 's retained information be $t_i \in T$ and its additional information be $s_i \in S_i$, and let $\theta \in \Theta$ denote the public state. The retained information t_i may include a bidder-specific payoff-relevant component, such as a value for the object. In this environment, the richness of the maintained class lies in allowing payoffs to depend on anonymous summaries of rivals' bids. A player's payoff may depend on its own bid, its own retained information, and a statistic such as the highest rival bid, the second-highest bid, or another anonymous threshold summary of opponents' bids.

The maintained class therefore accommodates symmetric auction environments in which strategic interaction operates through anonymous bid summaries. It covers complete information and private information environments, but also intermediate environments in which bidders possess additional noisy information about opponents' retained information through s_i , provided that the auction environment is admissible and the maintained symmetry and conditional independence restrictions are satisfied. The theorem is not tied to entry environments in which strategic interaction is summarized by counts; it also applies to canonical auction settings with anonymous but strategically nontrivial payoff-relevant interaction.

Taken together, the two environments highlight complementary sources of richness within the maintained class. In entry environments, richness appears on the retained-information side, since payoffs may depend on the retained-information profile. In auction environments, richness appears on the action side, since payoffs may depend on anonymous payoff-relevant summaries of opponents' bids. Both environments may also feature multidimensional retained information and additional noisy information about opponents' retained information through s_i . The point is not simply that the theorem covers two familiar applications. It is that the maintained class is broad within its intended symmetric framework and economically recognizable.

6 Scope of the Exact Reduction

Section 5 clarifies the economic content of the maintained class. Section 6 turns to the scope of the exactness claim established in Theorem 1. Unrestricted Bayes correlated equilibrium (BCE) remains the fully general reduced-form object for static games with latent information. The contribution of this paper is not to replace that general object, nor to refine the equilibrium discipline of an arbitrary correlated-equilibrium notion. The paper restricts attention to a maintained class of admissible original environments, namely symmetric games with exchangeable retained information and symmetric conditionally independent additional information, and shows that, within that class, the exact reduced-form object is SBCE.

This distinction is central to the scope of the theorem: what is restricted is the maintained architecture of the original game, not the logic of equilibrium characterization on

the reduced side. Exactness therefore has substantive content. The reduced equilibrium object identified here is not a convenient proxy, an approximation, or an outer bound. Within the maintained class, it is the correct reduced equilibrium object: the outcome distributions induced by symmetric Bayes Nash equilibrium coincide exactly with those induced by Symmetric Bayes Correlated Equilibrium in the reduced game.

An exact lower-dimensional reduction matters for a further reason. In symmetric environments, the issue is not only that the unrestricted BCE object may be large, but that it describes equilibrium restrictions at a level of generality finer than is needed for the maintained class. Once exactness is established, the analyst can work directly with the lower-dimensional symmetric object without sacrificing fidelity to the original game. The reduction therefore does not merely simplify the characterization of outcomes; it identifies the level of robust equilibrium analysis appropriate for the class.

At the same time, the theorem should not be read as saying that symmetry alone is sufficient for exact reduction. The result depends jointly on restrictions on payoffs, on the retained information structure, and on the additional information layer. In particular, exchangeability of retained information and symmetric conditional independence of the latent additional information component are essential. When these ingredients are combined, symmetry does more than reduce dimensionality: it determines the exact reduced equilibrium object for the class.

Bergemann and Morris (2016) identify BCE as the fully general reduced-form object for equilibrium outcomes under latent information. Within a maintained symmetric class, the exact reduced-form object is SBCE rather than the unrestricted BCE.

The present paper is therefore best understood as a single-game foundation. It identifies the reduced equilibrium object, establishes the exact equivalence, and clarifies the environments for which exactness obtains. Extensions to linked games, joint predictions, and fixed-information counterfactual analysis require additional structure and results and are developed separately. More broadly, the paper shows that in games with latent information, symmetry may do more than simplify analysis: under an appropriate informational architecture, it can identify the exact reduced equilibrium object.

7 Conclusion

This paper studies symmetric static games with latent information and identifies Symmetric Bayes Correlated Equilibrium (SBCE) as the exact reduced equilibrium object for a maintained symmetric class. Fixing the common payoff-relevant environment, the main theorem shows that the set of outcome distributions induced by symmetric Bayes Nash equilibrium, as the admissible additional information structure varies within the class, coincides exactly with the set induced by SBCE in the corresponding reduced game.

The maintained class is broad enough to include complete information, private information, and intermediate noisy signal environments, as well as economically recognizable settings such as entry and auctions. At the same time, unrestricted BCE remains the appropriate benchmark outside that class. Once attention is restricted to the maintained symmetric class, SBCE provides the exact reduced-form characterization of equilibrium outcomes.

A Proof of Supporting Lemmas

Lemma 1 (Canonical additional information structure). *Let ν be a symmetric independent recommendation rule for the fixed environment $\mathcal{E}$. For each player i , let*

$$S_i^* := A^{N-1},$$

viewed as an opponent-indexed collection and equipped with the natural action of permutations of player i 's opponents by coordinate relabeling. Let

$$\kappa_i : A \rightarrow S_i^*$$

be the measurable injective map given by

$$\kappa_i(a) := (a, \dots, a) \in A^{N-1}.$$

For each $(t, \theta) \in T^N \times \Theta$, define the conditional law of player i 's additional information by

$$\eta_i^\nu(B_i \mid t, \theta) := \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in B_i\} \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta), \quad B_i \subseteq S_i^* \text{ measurable},$$

and let

$$\Pi^\nu(\cdot \mid t, \theta) := \bigotimes_{i=1}^N \eta_i^\nu(\cdot \mid t, \theta) \in \Delta\left(\prod_{i=1}^N S_i^*\right).$$

Then the family $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ defines an admissible original information structure

$$I^\nu \in \mathcal{I}(\mathcal{E}).$$

In particular, the additional information layer of I^ν is conditionally independent given (t, θ) and symmetric across players.

Proof. Fix a symmetric independent recommendation rule $\nu(\cdot \mid t_i, m, \theta) \in \Delta(A)$. For each player i , let S_i^* , κ_i , η_i^ν , and Π^ν be as defined in the statement.

Since A is finite and $\nu(\cdot \mid t_i, n_{-i}(t_{-i}), \theta) \in \Delta(A)$, it follows that $\eta_i^\nu(\cdot \mid t, \theta) \in \Delta(S_i^*)$ for every player i . Equivalently, for each $(t, \theta) \in T^N \times \Theta$, $\eta_i^\nu(\cdot \mid t, \theta)$ is the pushforward of $\nu(\cdot \mid t_i, n_{-i}(t_{-i}), \theta)$ under κ_i . Hence

$$\Pi^\nu(\cdot \mid t, \theta) := \bigotimes_{i=1}^N \eta_i^\nu(\cdot \mid t, \theta) \in \Delta\left(\prod_{i=1}^N S_i^*\right)$$

is well defined.

Define the joint information structure by

$$\pi^\nu(dt, ds, d\theta) := \pi^1(dt, d\theta) \Pi^\nu(ds \mid t, \theta).$$

Then the marginal of π^ν on $T^N \times \Theta$ is π^1 , so the fixed environment $\mathcal{E} = (\mathcal{G}, \pi^1)$ is preserved. Moreover, for every measurable rectangle

$$\prod_{i=1}^N B_i \subseteq \prod_{i=1}^N S_i^*,$$

one has

$$\Pi^\nu\left(\prod_{i=1}^N B_i \mid t, \theta\right) = \prod_{i=1}^N \eta_i^\nu(B_i \mid t, \theta),$$

so the additional information layer is conditionally independent across players given (t, θ) . Moreover, the playerwise conditional laws $\eta_i^\nu(\cdot \mid t, \theta)$ are all induced by the same underlying recommendation rule $\nu(\cdot \mid t_i, n_{-i}(t_{-i}), \theta)$, so this construction satisfies the common-kernel requirement in Section 2.3.

Now define the common kernel $\rho^\nu(\cdot \mid t_i, t_{-i}, \theta) := \eta_i^\nu(\cdot \mid t, \theta)$. Under the canonical convention $S_i^* = A^{N-1}$ for all i , this definition depends on i only through (t_i, t_{-i}, θ) , because the expression

$$\eta_i^\nu(B \mid t, \theta) = \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in B\} \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta)$$

depends on i only through (t_i, t_{-i}, θ) .

To verify symmetry, fix a player i , a permutation ϕ_{-i} of player i 's opponents, and a measurable set $B_i \subseteq S_i^*$. Because the recommendation rule ν is symmetric independent, the conditional law $\nu(\cdot \mid t_i, n_{-i}(t_{-i}), \theta)$ depends on opponents' retained information only through the anonymous count vector $n_{-i}(t_{-i})$. Since relabeling opponents does not change this count vector, one has

$$n_{-i}(t_{-i}) = n_{-i}(\phi_{-i}(t_{-i})).$$

Using the definition of η_i^ν , it follows that

$$\begin{aligned} \eta_i^\nu(\phi_{-i}(B_i) \mid t_i, \phi_{-i}(t_{-i}), \theta) &= \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in \phi_{-i}(B_i)\} \nu(a_i \mid t_i, n_{-i}(\phi_{-i}(t_{-i})), \theta) \\ &= \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in \phi_{-i}(B_i)\} \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta). \end{aligned}$$

Because $\kappa_i(a_i) = (a_i, \dots, a_i)$ is constant across opponent-indexed coordinates, it is invariant under relabeling of those coordinates. Hence

$$\phi_{-i}(\kappa_i(a_i)) = \kappa_i(a_i) \quad \text{for every } a_i \in A,$$

and therefore

$$\mathbf{1}\{\kappa_i(a_i) \in \phi_{-i}(B_i)\} = \mathbf{1}\{\kappa_i(a_i) \in B_i\}.$$

Substituting this into the preceding display yields

$$\eta_i^\nu(\phi_{-i}(B_i) \mid t_i, \phi_{-i}(t_{-i}), \theta) = \sum_{a_i \in A} \mathbf{1}\{\kappa_i(a_i) \in B_i\} \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta) = \eta_i^\nu(B_i \mid t_i, t_{-i}, \theta).$$

By the definition

$$\rho^\nu(\cdot \mid t_i, t_{-i}, \theta) := \eta_i^\nu(\cdot \mid t, \theta),$$

it follows that

$$\rho^\nu(\phi_{-i}(B_i) \mid t_i, \phi_{-i}(t_{-i}), \theta) = \rho^\nu(B_i \mid t_i, t_{-i}, \theta).$$

Thus the common kernel ρ^ν is invariant to permutations of opponents, in the sense required by the definition of a symmetric additional information structure.

Define the joint distribution $\pi^\nu \in \Delta(T^N \times S^* \times \Theta)$ by

$$\pi^\nu(t, s, \theta) := \pi^1(t, \theta) \Pi^\nu(s \mid t, \theta), \quad (t, s, \theta) \in T^N \times S^* \times \Theta,$$

where $S^* := \prod_{i=1}^N S_i^*$. Then π^ν has marginal π^1 on $T^N \times \Theta$ by construction. Moreover, its additional-information layer is conditionally independent given (t, θ) , generated by the common kernel ρ^ν , and symmetric across players in the sense of Section 2.3. Therefore

$$I^\nu := (T, (S_i^*)_{i=1}^N, \pi^\nu) \in \mathcal{I}(\mathcal{E}).$$

□

Lemma 2 (Implementation of the recommendation rule). *Let ν be a symmetric independent recommendation rule for the fixed environment $\mathcal{E}$, and let $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ be the canonical additional information structure constructed in Lemma 1. Fix an arbitrary reference distribution $q^\circ \in \Delta(A)$. For each player i , define the canonical strategy*

$$\sigma_i^\nu(\cdot \mid t_i, s_i, \theta) := \begin{cases} \delta_{\kappa_i^{-1}(s_i)}, & \text{if } s_i \in \kappa_i(A), \\ q^\circ, & \text{if } s_i \notin \kappa_i(A), \end{cases} \quad (t_i, s_i, \theta) \in T \times S_i^* \times \Theta,$$

where δ_a denotes the degenerate distribution on A that assigns probability one to a . Let

$$\sigma^\nu(a \mid t, s, \theta) := \prod_{i=1}^N \sigma_i^\nu(a_i \mid t_i, s_i, \theta)$$

denote the induced conditional distribution over action profiles. Then for every $(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta$, one has

$$\int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta) = \nu(a \mid t, \theta),$$

where

$$\nu(a \mid t, \theta) := \prod_{i=1}^N \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta).$$

Consequently, the joint distribution on $\mathcal{A} \times T^N \times \Theta$ induced by the original game construction $(\pi^1, \Pi^\nu, \sigma^\nu)$ is precisely $\mu_{\nu, \mathcal{E}}$.

Proof. Fix $(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta$. By construction, for each player i and each $(t_i, s_i, \theta) \in T \times S_i^* \times \Theta$,

$$\sigma_i^\nu(a_i \mid t_i, s_i, \theta) = \mathbf{1}\{s_i = \kappa_i(a_i)\} \quad \text{whenever } s_i \in \kappa_i(A).$$

If $s_i \notin \kappa_i(A)$, then $\sigma_i^\nu(\cdot \mid t_i, s_i, \theta) = q^\circ$ by definition.

Hence, on the event $\{s_i \in \kappa_i(A) \text{ for all } i\}$,

$$\sigma^\nu(a \mid t, s, \theta) = \prod_{i=1}^N \mathbf{1}\{s_i = \kappa_i(a_i)\}.$$

Under the canonical structure of Lemma 1, each marginal $\eta_i^\nu(\cdot \mid t, \theta)$ is supported on $\kappa_i(A)$, so the complement of this event is $\Pi^\nu(\cdot \mid t, \theta)$ -null. Therefore,

$$\begin{aligned} \int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta) &= \int_{s \in \prod_{i=1}^N S_i^*} \prod_{i=1}^N \mathbf{1}\{s_i = \kappa_i(a_i)\} \Pi^\nu(ds \mid t, \theta) \\ &= \Pi^\nu\left(\prod_{i=1}^N \{\kappa_i(a_i)\} \mid t, \theta\right). \end{aligned}$$

Using the product structure

$$\Pi^\nu(\cdot \mid t, \theta) = \bigotimes_{i=1}^N \eta_i^\nu(\cdot \mid t, \theta),$$

we obtain

$$\Pi^\nu\left(\prod_{i=1}^N \{\kappa_i(a_i)\} \mid t, \theta\right) = \prod_{i=1}^N \eta_i^\nu(\{\kappa_i(a_i)\} \mid t, \theta).$$

By the definition of η_i^ν in Lemma 1,

$$\eta_i^\nu(\{\kappa_i(a_i)\} \mid t, \theta) = \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta).$$

Hence

$$\int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta) = \prod_{i=1}^N \nu(a_i \mid t_i, n_{-i}(t_{-i}), \theta) = \nu(a \mid t, \theta).$$

Multiplying both sides by $\pi^1(t, \theta)$ yields

$$\pi^1(t, \theta) \int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta) = \pi^1(t, \theta) \nu(a \mid t, \theta) = \mu_{\nu, \mathcal{E}}(a, t, \theta).$$

By definition, under $(\pi^1, \Pi^\nu, \sigma^\nu)$ the induced joint distribution at (a, t, θ) is

$$\pi^1(t, \theta) \int_{s \in \prod_{i=1}^N S_i^*} \sigma^\nu(a \mid t, s, \theta) \Pi^\nu(ds \mid t, \theta).$$

Hence the joint distribution induced by $(\pi^1, \Pi^\nu, \sigma^\nu)$ is precisely $\mu_{\nu, \mathcal{E}}$.

□

Lemma 3 (Optimality of the canonical strategy). *Let ν be an SBCE recommendation rule for the fixed environment $\mathcal{E}$. Let $\{\Pi^\nu(\cdot \mid t, \theta)\}_{(t, \theta) \in T^N \times \Theta}$ be the canonical additional information structure constructed in Lemma 1, and let σ^ν be the canonical strategy profile defined in Lemma 2. Then σ^ν is a symmetric Bayes Nash equilibrium under the admissible original information structure*

$$I^\nu \in \mathcal{I}(\mathcal{E}).$$

Proof. Fix a player i , a state $\theta \in \Theta$, a retained information value $t_i \in T$, and an additional information realization $s_i \in S_i^*$.

Under the canonical construction, player i 's additional-information law is supported on $\kappa_i(A) \subseteq S_i^*$. Hence, if $s_i \notin \kappa_i(A)$, then the information event (t_i, s_i, θ) occurs with probability zero under the construction, so the optimality requirement at this information set is vacuous.

It therefore suffices to consider the case in which $s_i = \kappa_i(a_i)$ for some $a_i \in A$. In that case, the realized signal s_i encodes the unique action recommendation $a_i = \kappa_i^{-1}(s_i)$. Let $a'_i \in A$ be any alternative action.

Accordingly, under the canonical construction, the information event

$$(t_i, s_i, \theta) = (t_i, \kappa_i(a_i), \theta)$$

corresponds exactly to the recommendation event (t_i, a_i, θ) . Its probability under the induced (a, t, θ) -distribution μ_{σ^ν, I^ν} is therefore

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma^\nu, I^\nu}(a_i, a_{-i}, t_i, t_{-i}, \theta).$$

If

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma^\nu, I^\nu}(a_i, a_{-i}, t_i, t_{-i}, \theta) = 0,$$

then the information event (t_i, s_i, θ) occurs with probability zero under the construction, so the optimality requirement at this information set is vacuous. Hence it suffices to consider the case in which this quantity is strictly positive.

Consider first the conditional payoff to player i from following the encoded action a_i given the realized information triple (t_i, s_i, θ) . Since, under the canonical encoding, the realized signal $s_i = \kappa_i(a_i)$ identifies the unique recommended action a_i , conditioning on

(t_i, s_i, θ) is equivalent to conditioning on (t_i, a_i, θ) . Therefore, this conditional payoff is obtained by dividing

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma^\nu, I^\nu}(a_i, a_{-i}, t_i, t_{-i}, \theta) u_i(a_i, a_{-i}, t_i, t_{-i}, \theta)$$

by the same strictly positive normalizing factor

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma^\nu, I^\nu}(a_i, a_{-i}, t_i, t_{-i}, \theta).$$

By Lemma 2,

$$\mu_{\sigma^\nu, I^\nu} = \mu_{\nu, \mathcal{E}},$$

so the numerator is equal to

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) u_i(a_i, a_{-i}, t_i, t_{-i}, \theta).$$

Similarly, the conditional payoff from deviating to a'_i is obtained by dividing

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma^\nu, I^\nu}(a_i, a_{-i}, t_i, t_{-i}, \theta) u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta)$$

by the same strictly positive normalizing factor, and hence, again by Lemma 2, its numerator is equal to

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta).$$

Therefore, following a_i is optimal conditional on observing (t_i, s_i, θ) if and only if

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) [u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta)] \geq 0.$$

But this is the obedience inequality for the distribution $\mu_{\nu, \mathcal{E}}$, conditional on the triple (t_i, a_i, θ) and evaluated against the alternative action a'_i . Since ν is an SBCE recommendation rule for the fixed environment $\mathcal{E}$, the induced distribution $\mu_{\nu, \mathcal{E}}$ satisfies obedience. Hence the inequality holds, so following the encoded action a_i is optimal.

Because i , θ , t_i , s_i , and a'_i were arbitrary, σ^ν is a Bayes Nash equilibrium under the admissible original information structure I^ν . It remains to verify symmetry. First, for each player i ,

$$\sigma_i^\nu(a_i \mid t_i, s_i, \theta) = \mathbf{1}\{s_i = \kappa_i(a_i)\}$$

is invariant under any relabeling of opponents, because $\kappa_i(a_i) = (a_i, \dots, a_i)$ is unchanged by permutations of opponent-indexed coordinates. Second, the same canonical rule is used for every player. Hence σ^ν is symmetric. Therefore σ^ν is a symmetric Bayes Nash equilibrium under $I^\nu \in \mathcal{I}(\mathcal{E})$. $\square$

Lemma 4 (Anonymous aggregation and induced recommendation rule). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, and let σ be a symmetric Bayes Nash equilibrium under I . For each player i , define*

$$\gamma_i(t_i, m, \theta) := \sum_{t_{-i} \in T^{N-1}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta), \quad (t_i, m, \theta) \in T \times \mathcal{H} \times \Theta.$$

Then there exists a common anonymous marginal distribution $\gamma \in \Delta(T \times \mathcal{H} \times \Theta)$ such that $\gamma(t_i, m, \theta) = \gamma_i(t_i, m, \theta)$ for every player i .

Fix an arbitrary reference distribution $q^\circ \in \Delta(A)$. Using the anonymous marginal distribution γ and the equilibrium strategy profile σ , define for each player i

$$\nu_i^\sigma(a_i \mid t_i, m, \theta) := \frac{\sum_{t_{-i} \in T^{N-1}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta) \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)}{\gamma(t_i, m, \theta)},$$

whenever $\gamma(t_i, m, \theta) > 0$, and define $\nu_i^\sigma(\cdot \mid t_i, m, \theta) := q^\circ$ whenever $\gamma(t_i, m, \theta) = 0$.

Then there exists a common rule $\nu^\sigma(\cdot \mid t_i, m, \theta) \in \Delta(A)$ such that

$$\nu_i^\sigma(\cdot \mid t_i, m, \theta) = \nu^\sigma(\cdot \mid t_i, m, \theta) \quad \text{for all } i,$$

and the induced conditional distribution over action profiles is given by

$$\nu^\sigma(a \mid t, \theta) := \prod_{i=1}^N \nu^\sigma(a_i \mid t_i, n_{-i}(t_{-i}), \theta), \quad a \in \mathcal{A}.$$

Thus ν^σ defines a symmetric independent recommendation rule in the sense of Section 3.1.

Proof. We proceed in three steps.

Step 1: Construction of the anonymous marginal distribution. For each player i , define

$$\gamma_i(t_i, m, \theta) := \sum_{t_{-i} \in T^{N-1}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta), \quad (t_i, m, \theta) \in T \times \mathcal{H} \times \Theta.$$

Since π^1 is a probability measure on $T^N \times \Theta$, one has $\gamma_i(t_i, m, \theta) \geq 0$ for all (t_i, m, θ) , and

$$\sum_{\theta \in \Theta} \sum_{t_i \in T} \sum_{m \in \mathcal{H}} \gamma_i(t_i, m, \theta) = \sum_{\theta \in \Theta} \sum_{t \in T^N} \pi^1(t, \theta) = 1.$$

Thus $\gamma_i \in \Delta(T \times \mathcal{H} \times \Theta)$ for every player i .

We next show that γ_i does not depend on i . Fix players i and j , and fix $(\tau, m, \theta) \in T \times \mathcal{H} \times \Theta$. Let

$$E_i(\tau, m, \theta) := \{(t, \theta) : t_i = \tau, n_{-i}(t_{-i}) = m\},$$

and define $E_j(\tau, m, \theta)$ analogously. Choose any permutation ϕ of $\{1, \dots, N\}$ such that $\phi(i) = j$.

Then ϕ maps $E_i(\tau, m, \theta)$ bijectively onto $E_j(\tau, m, \theta)$. Indeed, if $(t, \theta) \in E_i(\tau, m, \theta)$, then $(\phi(t))_j = t_i = \tau$, and because ϕ relabels the opponents of player i bijectively into the opponents of player j , the count vector of opponents' retained signals is preserved: $n_{-j}((\phi(t))_{-j}) = n_{-i}(t_{-i}) = m$. Thus $(\phi(t), \theta) \in E_j(\tau, m, \theta)$, and the converse follows by applying the same argument to ϕ^{-1} . By exchangeability of π^1 ,

$$\pi^1(\phi(t), \theta) = \pi^1(t, \theta) \quad \text{for all } (t, \theta) \in T^N \times \Theta.$$

Therefore

$$\gamma_i(\tau, m, \theta) = \sum_{(t, \theta) \in E_i(\tau, m, \theta)} \pi^1(t, \theta) = \sum_{(t, \theta) \in E_j(\tau, m, \theta)} \pi^1(t, \theta) = \gamma_j(\tau, m, \theta).$$

Hence there exists a common anonymous marginal distribution $\gamma \in \Delta(T \times \mathcal{H} \times \Theta)$ such that $\gamma(t_i, m, \theta) = \gamma_i(t_i, m, \theta)$ for every player i .

Step 2: One-player formula and anonymity of the induced rule. We first verify the formula for ν_i^σ using the admissible factorization of the original information structure.

Fix player i , $(t_i, m, \theta) \in T \times \mathcal{H} \times \Theta$, and $t_{-i} \in T^{N-1}$ such that $n_{-i}(t_{-i}) = m$. Since $I \in I(\mathcal{E})$, conditional on (t, θ) the additional information profile factorizes according to the common kernel ρ . Hence, after integrating out s_{-i} from the measure $\pi(ds, t_i, t_{-i}, \theta)$ on $S = \prod_{k=1}^N S_k$, the induced measure on S_i is

$$\pi^1(t_i, t_{-i}, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta).$$

Because the integrand $\sigma_i(a_i \mid t_i, s_i, \theta)$ depends only on s_i , it follows that

$$\int_S \sigma_i(a_i \mid t_i, s_i, \theta) \pi(ds, t_i, t_{-i}, \theta) = \pi^1(t_i, t_{-i}, \theta) \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta).$$

Hence, whenever $\gamma(t_i, m, \theta) > 0$,

$$\nu_i^\sigma(a_i \mid t_i, m, \theta) = \frac{\sum_{t_{-i}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta) \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)}{\gamma(t_i, m, \theta)}.$$

If $\gamma(t_i, m, \theta) = 0$, then by definition $\nu_i^\sigma(\cdot \mid t_i, m, \theta) = q^\circ \in \Delta(A)$, so the rule is defined at such points. Now suppose $\gamma(t_i, m, \theta) > 0$. Summing over $a_i \in A$ and using that $\sigma_i(\cdot \mid t_i, s_i, \theta) \in \Delta(A)$, we obtain

$$\sum_{a_i \in A} \nu_i^\sigma(a_i \mid t_i, m, \theta) = \frac{\sum_{t_{-i}: n_{-i}(t_{-i})=m} \int_S \sum_{a_i \in A} \sigma_i(a_i \mid t_i, s_i, \theta) \pi(ds, t_i, t_{-i}, \theta)}{\gamma(t_i, m, \theta)} = 1.$$

Hence $\nu_i^\sigma(\cdot \mid t_i, m, \theta) \in \Delta(A)$.

We now show that, when $\gamma(t_i, m, \theta) > 0$, the conditional expectation term

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)$$

depends on t_{-i} only through the anonymous count vector $m = n_{-i}(t_{-i})$. Because $I \in \mathcal{I}(\mathcal{E})$ is admissible, the additional information structure is homogeneous conditionally independent. Therefore, for any fixed (t_i, t_{-i}, θ) ,

$$\pi(ds \mid t_i, t_{-i}, \theta) = \bigotimes_{k=1}^N \rho(ds_k \mid t_k, t_{-k}, \theta),$$

so in particular the marginal law of s_i conditional on (t_i, t_{-i}, θ) is $\rho(\cdot \mid t_i, t_{-i}, \theta)$. Hence

$$\int_S \sigma_i(a_i \mid t_i, s_i, \theta) \pi(ds, t_i, t_{-i}, \theta) = \pi^1(t_i, t_{-i}, \theta) \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta).$$

Fix now two opponents' retained information profiles $t_{-i}, \tilde{t}_{-i} \in T^{N-1}$ such that $n_{-i}(t_{-i}) = n_{-i}(\tilde{t}_{-i})$. Then there exists a permutation φ_{-i} of player i 's opponents such that $\tilde{t}_{-i} = \varphi_{-i}(t_{-i})$. By symmetry of the admissible original information structure,

$$\rho(B_i \mid t_i, t_{-i}, \theta) = \rho(\varphi_{-i}(B_i) \mid t_i, \tilde{t}_{-i}, \theta)$$

for every measurable set $B_i \subseteq S_i$. By condition (i) in the definition of a symmetric strategy profile,

$$\sigma_i(a_i \mid t_i, s_i, \theta) = \sigma_i(a_i \mid t_i, \varphi_{-i}(s_i), \theta) \quad \text{for all } s_i \in S_i.$$

Hence, by the change of variables induced by φ_{-i} ,

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta) = \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, \tilde{t}_{-i}, \theta).$$

Thus the integral

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)$$

depends on t_{-i} only through the anonymous count vector $m = n_{-i}(t_{-i})$.

Accordingly, for every player i and every (t_i, m, θ) such that $\gamma(t_i, m, \theta) > 0$, we may define

$$\bar{\sigma}_i(a_i \mid t_i, m, \theta) := \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta),$$

for any t_{-i} satisfying $n_{-i}(t_{-i}) = m$. This is well defined by the invariance argument just established. With this notation,

$$\nu_i^\sigma(a_i \mid t_i, m, \theta) = \frac{\sum_{t_{-i}: n_{-i}(t_{-i})=m} \pi^1(t_i, t_{-i}, \theta) \bar{\sigma}_i(a_i \mid t_i, m, \theta)}{\gamma(t_i, m, \theta)} = \bar{\sigma}_i(a_i \mid t_i, m, \theta).$$

Step 3: Equality across players and existence of a common rule. It remains to show that the one-player rule obtained above is the same for every player. Fix $(t^*, m, \theta) \in T \times \mathcal{H} \times \Theta$. If $\gamma(t^*, m, \theta) = 0$, then by definition $\nu_i^\sigma(\cdot \mid t^*, m, \theta) = q^\circ$ for every player i , so the rule is already common across players at (t^*, m, θ) . It therefore remains to consider the case $\gamma(t^*, m, \theta) > 0$.

Fix players i and j . Since $\gamma = \gamma_i = \gamma_j$ and $\gamma(t^*, m, \theta) > 0$, there exists $t_{-i} \in T^{N-1}$ such that $n_{-i}(t_{-i}) = m$ and $\pi^1(t^*, t_{-i}, \theta) > 0$. Likewise, there exists $t_{-j} \in T^{N-1}$ such that $n_{-j}(t_{-j}) = m$ and $\pi^1(t^*, t_{-j}, \theta) > 0$. Choose such t_{-i} and t_{-j} .

Because the two opponent profiles have the same count vector m , for each retained type $\tau \in T$ the number of opponents of type τ is the same in t_{-i} and t_{-j} . Hence there exists a bijection between the opponent set of player i and the opponent set of player j that matches coordinates with the same retained type. This bijection induces a relabeling map from S_i to S_j ; denote this induced relabeling by $\psi_{ij} : S_i \rightarrow S_j$. By construction, this is exactly the natural relabeling of opponent-indexed coordinates appearing in condition (ii) of the definition of a symmetric strategy profile.

By the definition of $\bar{\sigma}_i$ and $\bar{\sigma}_j$,

$$\begin{aligned} \bar{\sigma}_i(a \mid t^*, m, \theta) &= \int_{S_i} \sigma_i(a \mid t^*, s_i, \theta) \rho(ds_i \mid t^*, t_{-i}, \theta), \\ \bar{\sigma}_j(a \mid t^*, m, \theta) &= \int_{S_j} \sigma_j(a \mid t^*, s_j, \theta) \rho(ds_j \mid t^*, t_{-j}, \theta). \end{aligned}$$

By symmetry of the common kernel ρ , the relabeling map ψ_{ij} satisfies

$$(\psi_{ij})_{\#}\rho(\cdot \mid t^*, t_{-i}, \theta) = \rho(\cdot \mid t^*, t_{-j}, \theta),$$

that is, for every measurable set $C \subseteq S_j$,

$$\rho(C \mid t^*, t_{-j}, \theta) = \rho(\psi_{ij}^{-1}(C) \mid t^*, t_{-i}, \theta).$$

Indeed, this is exactly the symmetry requirement on the common kernel ρ applied to the relabeling of opponents that maps the profile t_{-i} into t_{-j} .

Moreover, by symmetry of the strategy profile, condition (ii) implies that for every $s_i \in S_i$,

$$\sigma_i(a \mid t^*, s_i, \theta) = \sigma_j(a \mid t^*, \psi_{ij}(s_i), \theta).$$

Therefore, using the pushforward identity above and a change of variables,

$$\begin{aligned} \bar{\sigma}_j(a \mid t^*, m, \theta) &= \int_{S_j} \sigma_j(a \mid t^*, s_j, \theta) \rho(ds_j \mid t^*, t_{-j}, \theta) \\ &= \int_{S_i} \sigma_j(a \mid t^*, \psi_{ij}(s_i), \theta) \rho(ds_i \mid t^*, t_{-i}, \theta) \\ &= \int_{S_i} \sigma_i(a \mid t^*, s_i, \theta) \rho(ds_i \mid t^*, t_{-i}, \theta) \\ &= \bar{\sigma}_i(a \mid t^*, m, \theta). \end{aligned}$$

Since i and j were arbitrary, the rule does not depend on the player index. Hence there exists a common stochastic kernel

$$\bar{\sigma}(\cdot \mid t^*, m, \theta) \in \Delta(A)$$

such that, for every player i ,

$$\bar{\sigma}_i(\cdot \mid t^*, m, \theta) = \bar{\sigma}(\cdot \mid t^*, m, \theta).$$

Therefore, for the fixed (t^*, m, θ) and for every player i ,

$$\nu_i^\sigma(\cdot \mid t^*, m, \theta) = \begin{cases} q^\circ, & \gamma(t^*, m, \theta) = 0, \\ \bar{\sigma}(\cdot \mid t^*, m, \theta), & \gamma(t^*, m, \theta) > 0. \end{cases}$$

Thus there exists a common rule

$$\nu^\sigma(\cdot \mid t^*, m, \theta) \in \Delta(A)$$

such that

$$\nu_i^\sigma(\cdot \mid t^*, m, \theta) = \nu^\sigma(\cdot \mid t^*, m, \theta) \quad \text{for all } i.$$

Thus ν^σ defines a symmetric independent recommendation rule in the sense of Section 3.1. □

Lemma 5 (Reproduction of the induced distribution). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, let σ be a symmetric Bayes Nash equilibrium under I , and let ν^σ be the symmetric independent recommendation rule constructed in Lemma 4. Then, for every $(t, \theta) \in T^N \times \Theta$ such that $\pi^1(t, \theta) > 0$,*

$$\nu^\sigma(a \mid t, \theta) = \frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} \quad \text{for all } a \in \mathcal{A}.$$

Consequently, the joint distribution on $\mathcal{A} \times T^N \times \Theta$ induced by the reduced game construction (π^1, ν^σ) is precisely $\mu_{\sigma, I}$.

Proof. Fix $(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta$. If $\pi^1(t, \theta) = 0$, then by definition

$$\mu_{\nu^\sigma, \mathcal{E}}(a, t, \theta) = \pi^1(t, \theta) \nu^\sigma(a \mid t, \theta) = 0.$$

Since $I \in \mathcal{I}(\mathcal{E})$ has marginal π^1 on $T^N \times \Theta$, one also has $\mu_{\sigma, I}(a, t, \theta) = 0$. Thus the stated equality holds trivially in this case.

Now consider (t, θ) such that $\pi^1(t, \theta) > 0$. By definition of $\mu_{\sigma, I}$,

$$\mu_{\sigma, I}(a, t, \theta) = \int_{s \in S} \sigma(a \mid t, s, \theta) \pi(ds, t, \theta).$$

Dividing by $\pi^1(t, \theta)$ gives

$$\frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} = \int_{s \in S} \sigma(a \mid t, s, \theta) \pi(ds \mid t, \theta).$$

Because $I \in \mathcal{I}(\mathcal{E})$ is admissible, its additional information layer is conditionally independent across players given (t, θ) . Equivalently, the rectangle factorization in Section 2.3 implies that the conditional law on $S = \prod_{i=1}^N S_i$ is the product measure

$$\pi(ds \mid t, \theta) = \bigotimes_{i=1}^N \rho(ds_i \mid t_i, t_{-i}, \theta),$$

where $\rho(\cdot \mid t_i, t_{-i}, \theta)$ denotes the one-player conditional law from Section 2.3. Since the equilibrium strategy profile is playerwise mixed,

$$\sigma(a \mid t, s, \theta) = \prod_{i=1}^N \sigma_i(a_i \mid t_i, s_i, \theta).$$

Therefore,

$$\frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} = \prod_{i=1}^N \left(\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta) \right).$$

Fix a player i , and write $m = n_{-i}(t_{-i})$.

We first show that

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)$$

depends on t_{-i} only through the anonymous count vector $m = n_{-i}(t_{-i})$. To verify this, fix two opponents' retained information profiles t_{-i} and $\tilde{t}_{-i}$ such that

$$n_{-i}(t_{-i}) = n_{-i}(\tilde{t}_{-i}).$$

Then there exists a permutation ϕ_{-i} of player i 's opponents such that

$$\tilde{t}_{-i} = \phi_{-i}(t_{-i}).$$

By symmetry of the admissible original information structure,

$$\rho(B_i \mid t_i, t_{-i}, \theta) = \rho(\phi_{-i}(B_i) \mid t_i, \tilde{t}_{-i}, \theta)$$

for every measurable set $B_i \subseteq S_i$. By condition (i) in the definition of a symmetric strategy profile,

$$\sigma_i(a_i \mid t_i, s_i, \theta) = \sigma_i(a_i \mid t_i, \phi_{-i}(s_i), \theta) \quad \text{for all } s_i \in S_i.$$

Hence, using the change of variables induced by the relabeling action ϕ_{-i} on S_i ,

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta) = \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, \tilde{t}_{-i}, \theta).$$

This proves the desired invariance.

It follows that for every realized profile t_{-i} with $m = n_{-i}(t_{-i})$, the integral

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta)$$

is constant over all t_{-i} with the same anonymous count vector m .

Since $\pi^1(t, \theta) > 0$ and $m = n_{-i}(t_{-i})$, we have

$$\gamma(t_i, m, \theta) \geq \pi^1(t_i, t_{-i}, \theta) = \pi^1(t, \theta) > 0.$$

Therefore, the case $\gamma(t_i, m, \theta) > 0$ applies in the definition of $\nu^\sigma(\cdot \mid t_i, m, \theta)$. Moreover, by Lemma 4, the integral

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, \tilde{t}_{-i}, \theta)$$

is the same for every $\tilde{t}_{-i} \in T^{N-1}$ such that $n_{-i}(\tilde{t}_{-i}) = m$. Hence, by Lemma 4, the numerator in the definition of $\nu^\sigma(a_i \mid t_i, m, \theta)$ is this common value times

$$\sum_{\tilde{t}_{-i} \in T^{N-1}: n_{-i}(\tilde{t}_{-i}) = m} \pi^1(t_i, \tilde{t}_{-i}, \theta) = \gamma(t_i, m, \theta).$$

Dividing by $\gamma(t_i, m, \theta)$ yields

$$\nu^\sigma(a_i \mid t_i, m, \theta) = \int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta).$$

Since $m = n_{-i}(t_{-i})$, this gives

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \rho(ds_i \mid t_i, t_{-i}, \theta) = \nu^\sigma(a_i \mid t_i, n_{-i}(t_{-i}), \theta).$$

Substituting this identity into the previous product representation yields

$$\frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} = \prod_{i=1}^N \nu^\sigma(a_i \mid t_i, n_{-i}(t_{-i}), \theta) = \nu^\sigma(a \mid t, \theta).$$

Thus, whenever $\pi^1(t, \theta) > 0$,

$$\nu^\sigma(a \mid t, \theta) = \frac{\mu_{\sigma, I}(a, t, \theta)}{\pi^1(t, \theta)} \quad \text{for all } a \in \mathcal{A}.$$

Finally, multiplying both sides by $\pi^1(t, \theta)$ gives

$$\mu_{\nu^\sigma, \mathcal{E}}(a, t, \theta) = \pi^1(t, \theta) \nu^\sigma(a \mid t, \theta) = \mu_{\sigma, I}(a, t, \theta)$$

for every $(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta$. □

Lemma 6 (Obedience of the induced recommendation rule). *Let $I \in \mathcal{I}(\mathcal{E})$ be an admissible original information structure, and let σ be a symmetric Bayes Nash equilibrium under I . Let $\nu^\sigma(\cdot \mid t_i, m, \theta) \in \Delta(A)$ be the induced recommendation rule constructed in Lemma 4. Then the induced joint distribution $\mu_{\nu^\sigma, \mathcal{E}}$ satisfies obedience.*

Proof. Fix a player i , a retained information value $t_i \in T$, a recommended action $a_i \in A$, a state $\theta \in \Theta$, and an alternative action $a'_i \in A$. To prove obedience, it is enough to verify the required inequality whenever

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu^\sigma, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) > 0.$$

If this quantity is zero, then the conditioning event (t_i, a_i, θ) has zero probability under $\mu_{\nu^\sigma, \mathcal{E}}$, and the obedience condition is vacuous.

Consider therefore the case in which

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu^\sigma, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) > 0.$$

By Lemma 5, $\mu_{\nu^\sigma, \mathcal{E}} = \mu_{\sigma, I}$. Hence

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma, I}(a_i, a_{-i}, t_i, t_{-i}, \theta) > 0.$$

By the definition of $\mu_{\sigma, I}$, this is equivalent to

$$\int_{S_i} \sigma_i(a_i \mid t_i, s_i, \theta) \pi(ds_i, t_i, \theta) > 0.$$

Therefore $\pi(t_i, \theta) > 0$, and for the fixed pair (t_i, θ) under consideration, the set

$$\{s_i \in S_i : \sigma_i(a_i \mid t_i, s_i, \theta) > 0\}$$

has strictly positive $\pi(\cdot \mid t_i, \theta)$ -probability.

For every $s_i \in S_i$, define

$$\begin{aligned} \Delta_i(a_i, a'_i; t_i, s_i, \theta) := & \sum_{t_{-i} \in T^{N-1}} \int_{s_{-i} \in \prod_{k \neq i} S_k} \sum_{a_{-i} \in A^{N-1}} \sigma_{-i}(a_{-i} \mid t_{-i}, s_{-i}, \theta) \\ & \times \left[u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta) \right] \pi(ds_{-i}, t_{-i} \mid t_i, s_i, \theta). \end{aligned}$$

We first show that, for almost every (t_i, s_i, θ) such that

$$\sigma_i(a_i \mid t_i, s_i, \theta) > 0,$$

one has, for every $a'_i \in A$

$$\Delta_i(a_i, a'_i; t_i, s_i, \theta) \geq 0.$$

This follows from the Bayes Nash equilibrium property by a measurable deviation argument. Indeed, fix $a_i \in A$. Suppose, toward a contradiction, that there exist $a'_i \in A$ and a measurable set

$$B \subseteq \{(\tau_i, s_i, \vartheta) \in T \times S_i \times \Theta : \sigma_i(a_i \mid \tau_i, s_i, \vartheta) > 0\}$$

of positive probability such that, for every $(\tau_i, s_i, \vartheta) \in B$,

$$\Delta_i(a_i, a'_i; \tau_i, s_i, \vartheta) < 0.$$

Define a measurable deviation $\tilde{\sigma}_i$ as follows. Outside B , let $\tilde{\sigma}_i = \sigma_i$. For each $(\tau_i, s_i, \vartheta) \in B$, let $\tilde{\sigma}_i(\cdot \mid \tau_i, s_i, \vartheta)$ be obtained from $\sigma_i(\cdot \mid \tau_i, s_i, \vartheta)$ by transferring all probability mass assigned to a_i to a'_i , that is,

$$\tilde{\sigma}_i(a \mid \tau_i, s_i, \vartheta) = \begin{cases} 0, & a = a_i, \\ \sigma_i(a'_i \mid \tau_i, s_i, \vartheta) + \sigma_i(a_i \mid \tau_i, s_i, \vartheta), & a = a'_i, \\ \sigma_i(a \mid \tau_i, s_i, \vartheta), & a \notin \{a_i, a'_i\}. \end{cases}$$

This defines a measurable deviation because B is measurable and the modification is coordinatewise measurable.

For every $(\tau_i, s_i, \vartheta) \in B$, the change in player i 's conditional expected payoff induced by this deviation is

$$-\sigma_i(a_i \mid \tau_i, s_i, \vartheta) \Delta_i(a_i, a'_i; \tau_i, s_i, \vartheta),$$

which is strictly positive because $\sigma_i(a_i \mid \tau_i, s_i, \vartheta) > 0$ and $\Delta_i(a_i, a'_i; \tau_i, s_i, \vartheta) < 0$ on B . Hence $\tilde{\sigma}_i$ yields a strictly higher conditional expected payoff than σ_i at every information realization in B , and therefore a strictly higher ex ante expected payoff, contradicting that σ is a Bayes Nash equilibrium.

Thus, for every a_i with

$$\sigma_i(a_i \mid t_i, s_i, \theta) > 0,$$

one has

$$\Delta_i(a_i, a'_i; t_i, s_i, \theta) \geq 0 \quad \text{for every } a'_i \in A,$$

for almost every (t_i, s_i, θ) .

Since the inequality

$$\Delta_i(a_i, a'_i; t_i, s_i, \theta) \geq 0$$

holds for almost every (t_i, s_i, θ) , multiplying by the nonnegative weight

$$\sigma_i(a_i \mid t_i, s_i, \theta)$$

preserves the inequality almost surely. Therefore, integrating over $s_i \in S_i$ with respect to the conditional distribution of s_i given (t_i, θ) gives

$$\int_{s_i \in S_i} \Delta_i(a_i, a'_i; t_i, s_i, \theta) \sigma_i(a_i \mid t_i, s_i, \theta) \pi(ds_i \mid t_i, \theta) \geq 0.$$

Multiplying both sides by the nonnegative factor $\pi^1(t_i, \theta)$ and using

$$\pi(ds, t, \theta) = \pi(ds_{-i}, t_{-i} \mid t_i, s_i, \theta) \pi(ds_i \mid t_i, \theta) \pi^1(t_i, \theta),$$

we obtain

$$\begin{aligned} & \sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \int_{s \in S} \pi(ds, t, \theta) \sigma_i(a_i \mid t_i, s_i, \theta) \sigma_{-i}(a_{-i} \mid t_{-i}, s_{-i}, \theta) \\ & \quad \times \left[u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta) \right] \geq 0. \end{aligned}$$

By the definition of $\mu_{\sigma, I}$, the preceding inequality can be rewritten as

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\sigma, I}(a_i, a_{-i}, t_i, t_{-i}, \theta) \left[u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta) \right] \geq 0.$$

By Lemma 5, $\mu_{\nu^\sigma, \mathcal{E}}(a, t, \theta) = \mu_{\sigma, I}(a, t, \theta)$ for every $(a, t, \theta) \in \mathcal{A} \times T^N \times \Theta$.

Substituting this equality into the preceding inequality yields

$$\sum_{t_{-i} \in T^{N-1}} \sum_{a_{-i} \in A^{N-1}} \mu_{\nu^\sigma, \mathcal{E}}(a_i, a_{-i}, t_i, t_{-i}, \theta) \left[u_i(a_i, a_{-i}, t_i, t_{-i}, \theta) - u_i(a'_i, a_{-i}, t_i, t_{-i}, \theta) \right] \geq 0.$$

This is the obedience inequality for $\mu_{\nu^\sigma, \mathcal{E}}$ conditional on (t_i, a_i, θ) .

Since i , t_i , a_i , θ , and a'_i were arbitrary, the obedience inequalities hold for all players, retained information values, recommended actions, states, and alternative actions whenever the conditioning event has positive probability. Hence the induced joint distribution $\mu_{\nu^\sigma, \mathcal{E}}$ satisfies obedience. $\square$

References

- Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics* 1, 67–96.
- Aumann, R.J., 1987. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society* , 1–18.
- Bergemann, D., Brooks, B., Morris, S., 2022. Counterfactuals with latent information. *American Economic Review* 112, 343–368.
- Bergemann, D., Morris, S., 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11, 487–522.
- Forges, F., 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision* 35, 277–310.
- Forges, F., 2006. Correlated equilibrium in games with incomplete information revisited. *Theory and Decision* 61, 329–344.
- Gualdani, C., Sinha, S., 2024. Identification in discrete choice models with imperfect information. *Journal of Econometrics* 244, 105854.
- Hara, K., Ito, Y., Koh, P.S., 2025. Identification and estimation of dynamic games with unknown information structure. Working paper.
- Liu, Q., 2015. Correlation and common priors in games with incomplete information. *Journal of Economic Theory* 157, 49–75.
- Magnolfi, L., Roncoroni, C., 2023. Estimation of discrete games with weak assumptions on information. *The Review of Economic Studies* 90, 2006–2041.
- Menzel, K., 2016. Inference for games with many players. *The Review of Economic Studies* 83, 306–337.
- Syrngkanis, V., Tamer, E., Ziani, J., 2021. Inference on auctions with weak assumptions on information. Working paper.