Generic Games Have Elaborations with Nearly Unique Equilibrium Outcomes

Tetsuya Hoshino*

September 10, 2026

Abstract

We show that generic finite games admit arbitrarily small incomplete-information elaborations whose equilibrium outcomes are nearly unique. For any finite number of players, we identify an explicit Lebesgue-null exceptional set of payoff vectors. Outside this set, for every $\varepsilon, \rho > 0$, there exists a finite common-prior elaboration in which, with probability at least $1 - \varepsilon$, every player knows that their payoffs coincide exactly with those of the base game. The ex ante action distributions induced by all Bayesian Nash equilibria of this elaboration have total-variation diameter less than ρ. The proof combines a finite common-noise construction and parabolic estimates with an auxiliary-type implementation that preserves exact normality. As a corollary, every game outside the exceptional set has at most one Kajii–Morris robust action profile.

1 Model and Main Results

We study finite elaborations whose Bayesian Nash equilibria induce nearly the same ex ante distribution of actions. The perturbation is small in the Kajii–Morris sense: with high probability all players know that their payoffs are exactly those of the base game. The construction implies that at most one pure action profile can be KM robust.

1.1 Model

Fix a finite set of players I and a finite set of actions A_i for each player i. Throughout the paper, we assume that $|A_i| \geq 2$ for every $i \in I$. As usual, let $A = \prod_i A_i$ and $A_{-i} = \prod_{j \neq i} A_j$. We write $\mathcal{U} = \prod_i \mathbb{R}^A$ for the full payoff space. A *base game* is $g = (g_i)_{i \in I} \in \mathcal{U}$, where $g_i(a)$ denotes player i's payoff at action profile $a \in A$.

*Kyoto University. Email: `tetsuyahoshino137@gmail.com`.

Disclosure of AI use and formal verification: This paper builds on the author's prior research on the problem. The mathematical argument was developed with substantial assistance from large language models, with OpenAI's GPT-5.6 Sol and GPT-6 Astra providing proof ideas. The author takes responsibility for all content. The main theorem and its corollary have been formalized in Lean 4. The Lean source code is available in the author's GitHub repository.

Let $\Delta(X)$ denote the set of probability distributions on a finite or countable set X. For two probability vectors μ, μ' on a finite set, we define the total variation (TV) distance by

$$\|\mu - \mu'\|_{\mathrm{TV}} = \max_B |\mu(B) - \mu'(B)| = \frac{1}{2}\sum_x |\mu(x) - \mu'(x)|.$$

The difference of two probability vectors at a single point is at most their TV distance.

Definition 1. An *elaboration* of base game g consists of countable type sets $(T_i)_i$, a common prior $P \in \Delta(T)$ on the set of type profiles $T = \prod_i T_i$, under which every type has positive marginal probability, and bounded payoff functions $u_i : A \times T \to \mathbb{R}$. The elaboration is *finite* if every T_i is finite.

We call a type profile $t \in T$ a *state*, and say that a state is *supported* if $P(t) > 0$. Boundedness of u_i is required within each elaboration; no bound common to all elaborations is imposed. In particular, the additional types used below may have payoff magnitudes that grow as the implementation is refined, while each individual elaboration has bounded payoffs.

Definition 2. Given a base game g and an elaboration, a type $t_i \in T_i$ is *normal* if

$$u_i(a, t_i, t_{-i}) = g_i(a) \qquad \text{for all } a \in A \text{ and } t_{-i} \text{ such that } P(t_i, t_{-i}) > 0. \tag{1}$$

For $\varepsilon \geq 0$, an elaboration is an ε-*elaboration* if the event that all players' types are normal has probability at least $1 - \varepsilon$.

Normality is exact: a normal type has the base-game payoff table at every supported state in which it occurs and for every action profile. Small ε bounds the probability that some player is not normal; it does not mean a uniform small change in payoff tables.

A strategy of player i is a mapping $\sigma_i : T_i \to \Delta(A_i)$, where $\sigma_i(a_i \mid t_i)$ denotes the probability that type t_i plays action a_i. A *Bayesian Nash equilibrium* (BNE) is a profile σ at which every type's mixed action is supported on maximizers of its interim expected payoff. Write $\mathrm{BNE}(E)$ for the set of BNE of an elaboration E. The ex ante action distribution induced by a strategy profile is

$$\mu_\sigma(a) = \sum_t P(t) \prod_i \sigma_i(a_i \mid t_i). \tag{2}$$

When σ is a BNE, we call μ_σ its *BNE outcome*.

Every finite elaboration has a BNE. Indeed, treating each player-type as an agent gives a finite game whose typewise best responses are exactly the BNE conditions; equilibrium existence follows from the usual finite-dimensional fixed-point argument. Thus the BNE outcome sets in the main result are nonempty.

1.2 Genericity and the Main Result

For the fixed player and action sets, let $Z \subset \mathcal{U}$ be the exceptional set defined in Appendix A.1. It is an explicit finite union of proper algebraic zero sets and hence is Lebesgue null. Excluding Z is a sufficient condition for the local belief correction used to realize the auxiliary feedback games below. The set also includes the pure unilateral payoff-tie hyperplanes, which settle the one-player case. The determinant definition and its elementary justification are given in that appendix.

Theorem 1. *For every $g \in \mathcal{U} \setminus Z$, and for every $\varepsilon > 0$ and $\rho > 0$, there exists a finite ε-elaboration E of base game g such that*

$$\sup_{\sigma,\sigma' \in \mathrm{BNE}(E)} \|\mu_\sigma - \mu_{\sigma'}\|_{\mathrm{TV}} < \rho.$$

1.3 KM Robustness as a Corollary

Definition 3. A pure action profile $a \in A$ is *KM robust*, or simply *robust*, if, for every $\delta > 0$, there exists $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}]$ and every ε-elaboration of base game g, that elaboration has a Bayesian Nash equilibrium σ such that $\mu_\sigma(a) \geq 1 - \delta$.

We write $\bar{\varepsilon}_a(\delta)$ for a threshold for the tolerance δ in Definition 3. The threshold may depend on δ and on a, but not on the elaboration; the equilibrium may depend on the elaboration.

Corollary 1. *For every $g \in \mathcal{U} \setminus Z$, there is at most one KM robust pure action profile.*

Proof. Suppose that distinct pure profiles a and b are both robust. Set $\delta = 1/4$ and let $\varepsilon = \min\{\bar{\varepsilon}_a(\delta), \bar{\varepsilon}_b(\delta)\} > 0$. Apply Theorem 1 with this ε and $\rho = 1/2$. Robustness supplies BNE σ, σ' of the resulting elaboration with $\mu_\sigma(a) \geq 3/4$ and $\mu_{\sigma'}(b) \geq 3/4$. Since $a \neq b$, we have $\mu_{\sigma'}(a) \leq 1/4$, and consequently $\|\mu_\sigma - \mu_{\sigma'}\|_{\mathrm{TV}} \geq \mu_\sigma(a) - \mu_{\sigma'}(a) \geq \frac{1}{2}$, which contradicts Theorem 1. ■

Remark 1. Kajii and Morris (1997, Definitions 2.2–2.5) use a countable state space, a common prior, information partitions with positive-probability cells, bounded state-dependent payoffs, and exact knowledge of one's own base payoffs. Morris and Ui (2005, Definitions 1–2) use countable product type spaces and impose the equality on positive-probability opponent-type profiles, as in Definition 2.

These formulations agree for our purpose. From a product-type elaboration, take only its positive-probability type profiles as states and let each player's information cells be the fibers of its type. Conversely, delete null states in a state/partition model, use information cells as types, push forward the prior, and average payoffs over states sharing the same tuple of cells. Strategies, interim incentives and action distributions are preserved, and any originally normal cell remains normal. Thus the probability of normality can only increase in this last reduction. The sources label an elaboration by its exact abnormal probability, whereas Definition 2 uses an upper bound; quantifying over all sufficiently small positive errors makes these conventions equivalent. Finally, for a pure point mass, $\|\mu - \delta_a\|_\infty = 1 - \mu(a)$, so that the closeness requirement in Definition 3 is the same as closeness of action distributions in the sources.

2 Feedback Games and Finite Realization

Sections 2–5 develop the construction for two players. Section 5 also settles the one-player case. When player i is fixed, j denotes the opponent.

2.1 Proof Strategy

The construction separates the design of nearly determinate action distributions from the requirement of exact normality. We first specify an auxiliary feedback game, in which each type compares a continuous evaluation vector at the candidate strategy profile. A common noise tree and a regularized response make the induced action distributions of all feedback equilibria close to a single comparison distribution when the tree is sufficiently fine.

Proposition 1 realizes this auxiliary problem by ordinary finite Bayesian games while keeping designated types exactly normal. At a fixed tree size, its limit-point conclusion and compactness make all BNE restrictions uniformly close to feedback equilibria. The proof then bounds the contribution of added states to the full BNE outcome. Small perturbation parameters control normality and added probability, the tree size controls feedback outcomes, and the implementation precision is chosen last.

2.2 Feedback Games and Feedback Equilibria

Throughout this section, $I = \{1, 2\}$ and both players have at least two actions. Fix a finite type space $T = T_1 \times T_2$ and a common prior P with positive type marginals $P_i(t_i)$. The space of behavioral strategy profiles is

$$\Sigma = \prod_{i=1}^{2} \prod_{t_i \in T_i} \Delta(A_i).$$

Each $\sigma \in \Sigma$ assigns a mixed action $\sigma_i(\cdot \mid t_i)$ to every player and type. As a finite product of probability simplexes, Σ is a compact convex subset of a Euclidean space.

For a strategy profile σ, type t_i's belief about the opponent's action is

$$q_i(t_i, \sigma)(a_j) = \sum_{t_j} P(t_j \mid t_i)\sigma_j(a_j \mid t_j). \tag{3}$$

This averages the opponent's mixed action over the conditional distribution of the opponent's type. Let G_i be the payoff matrix with own actions a_i as rows, opponent's actions a_j as columns, and entries $g_i(a_i, a_j)$, where each action is placed in its player's coordinate. Then $G_i q_i(t_i, \sigma)$ is the vector of expected payoffs that type t_i uses to compare its actions under the base payoff table.

For the auxiliary problem, we replace this base payoff vector by a prescribed continuous evaluation vector $w_i(t_i, \sigma)$. The following definition specifies the action comparisons that the implementation result will deliver.

Definition 4 (Feedback equilibrium)**.** Given continuous maps $w_i(t_i, \cdot) : \Sigma \to \mathbb{R}^{A_i}$, one for each player and type, a profile $\sigma \in \Sigma$ is a *feedback equilibrium of w* if

$$\mathrm{supp}\sigma_i(\cdot \mid t_i) \subseteq \operatorname*{argmax}_{a_i} w_i(t_i, \sigma)(a_i) \qquad \text{for every } i, t_i. \tag{4}$$

Remark 2 (Feedback games)**.** We call a finite information structure (T_1, T_2, P) with positive type marginals, together with the action sets A_i and continuous evaluation maps $w_i(t_i, \cdot) : \Sigma \to \mathbb{R}^{A_i}$, a *feedback game*. Its equilibrium notion is feedback equilibrium as defined in Definition 4.

For the auxiliary problem studied below, Section 3.1 specifies the information structure, and (15) in Section 3.4 specifies the evaluation maps. Together they define the feedback game that we analyze.

Remark 3 (Feedback equilibrium and Bayesian Nash equilibrium)**.** In an ordinary Bayesian game with fixed payoff functions, a type's interim pure-action payoff vector depends on the opponent's strategy but not on the player's own strategy. When $w_i(t_i, \sigma)$ is such an interim payoff vector, (4) is precisely the Bayesian Nash equilibrium condition. In particular, choosing $w_i(t_i, \sigma) = G_i q_i(t_i, \sigma)$ recovers Bayesian Nash equilibrium for the game with prior P and base payoffs g. Note that $q_i(t_i, \sigma)$ is independent of player i's own σ_i, as in (3).

In the auxiliary problem, however, $w_i(t_i, \sigma)$ may depend on the mixed action $\sigma_i(\cdot \mid t_i)$ prescribed at the comparing type itself. The feedback equilibrium condition evaluates this vector at the candidate profile σ and holds it fixed when comparing actions, allowing mixing among actions that maximize the vector at that profile. It therefore need not describe the optimal choice of a mixed action by a player who takes into account how changing that choice changes the evaluation vector.

The auxiliary problem and its feedback equilibria are proof devices. The finite elaborations constructed in Proposition 1 are ordinary Bayesian games with fixed payoff functions, and their equilibrium concept is Bayesian Nash equilibrium. The proposition connects these objects: for any choice of a Bayesian Nash equilibrium from each elaboration in the sequence, every limit point of the restrictions to the original types is a feedback equilibrium of w.

2.3 Finite Feedback Realization

Proposition 1 (Finite Feedback Realization)**.** *Fix a two-player base game $g \notin Z$, where Z is the exceptional set defined in Appendix A.1, and suppose that both players have at least two actions. Put*

$$m = \max\{|A_1|, |A_2|\}, \qquad K = 1 + 2(m-1).$$

Let $T = T_1 \times T_2$ be any finite type space, and let P be a common prior with positive type marginals. Let Σ be the corresponding space of behavioral strategy profiles. For each player i and type t_i, let $w_i(t_i, \cdot) : \Sigma \to \mathbb{R}^{A_i}$ be a continuous auxiliary interim payoff map, as in Section 2.2.

There exist constants $\kappa_0 > 0$ and $D < \infty$, depending only on g and valid uniformly over all *such choices of T, P, and w, with the following property.*

Choose any subsets $S_i \subseteq T_i$ and numbers $\varepsilon \geq 0$ and $0 \leq \kappa < \kappa_0$ such that

$$P(S_1 \times S_2) \geq 1 - \varepsilon$$

and, for every i, every $t_i \in S_i$, and every $\sigma \in \Sigma$,

$$\|w_i(t_i, \sigma) - G_i q_i(t_i, \sigma)\|_\infty \leq \kappa. \tag{5}$$

Then the base game g admits a sequence of finite elaborations E_l, obtained by enlarging the original type sets, in which every original type $t_i \in S_i$ is normal.

Write $T_l = T_{l,1} \times T_{l,2}$ for the type space of E_l, with $T_i \subseteq T_{l,i}$ for each player i, and let P_l be its common prior. The original states have positive total probability, and the conditional distribution on these states is P:

$$P_l(T) > 0, \qquad \frac{P_l(t)}{P_l(T)} = P(t) \quad \text{for every } t \in T.$$

Define the ratio of the probability of the added states to that of the original states by

$$\Lambda_l = \frac{P_l(T_l \setminus T)}{P_l(T)}.$$

Thus each original state $t \in T$ has probability

$$P_l(t) = \frac{P(t)}{1 + \Lambda_l}.$$

For some nonnegative sequence $\tau_l \to 0$, the following conclusions hold.

(a) The probability ratio Λ_l satisfies

$$\Lambda_l \leq D\kappa + (K-1)\varepsilon + \tau_l, \qquad \tau_l \to 0. \tag{6}$$

(b) The probability in E_l that at least one player's type is not normal is at most

$$K\varepsilon + D\kappa + \tau_l. \tag{7}$$

(c) For each l, choose any Bayesian Nash equilibrium of E_l, and let σ^l be its restriction to the original types. These restrictions all belong to the same fixed strategy space Σ. Every limit point of (σ^l) as $l \to \infty$ is a feedback equilibrium of w.

(d) Fix a pure action profile a and a tolerance $\delta > 0$. Suppose that a is robust, and let $\bar{\varepsilon}_a(\delta)$ be a

robustness threshold for this tolerance, as in Definition 3. If

$$K\varepsilon + D\kappa < \bar{\varepsilon}_a(\delta),$$

then there exists a feedback equilibrium σ of w such that

$$\mu_\sigma(a) \geq 1 - (1 + D\kappa + (K-1)\varepsilon)\delta. \tag{8}$$

Here μ_σ is the action distribution induced by σ under the original prior P, as in (2).

The constants κ_0 and D are independent of the prior P, the number of original types, and the magnitude of the auxiliary payoffs $w_i(t_i, \cdot)$ at types $t_i \notin S_i$.

2.4 Why Belief Correction and Additional Types Are Needed

At a designated normal type, the evaluation $w_i(t_i, \sigma)$ may differ slightly from the base interim payoff $G_i q_i(t_i, \sigma)$, but its fixed payoff table must remain exactly g_i. Appendix A.2 constructs a continuous correction distribution over opponent actions that reproduces the evaluation's best responses by changing beliefs, including all ties. The relative mass of the correction block is chosen uniformly over the allowed payoff errors.

That correction distribution depends on the unknown strategy profile. Appendix A.3 uses additional types of the existing players to approximate any such continuous distribution at every BNE. The construction gives the original normal types their base payoffs on all states where they occur; the added types' incentives are generated by fixed payoff tables. Small state masses are offset by large, finite helper payoffs, which is why no common payoff bound across elaborations is imposed.

Appendix A.4 combines these ingredients and proves all four parts of Proposition 1. Its construction is for an arbitrary fixed finite feedback game and is independent of the main theorem. The main proof uses parts (a)–(c); part (d) records the additional robustness-transfer consequence.

3 The Noise-Tree Feedback Game

We first specify the finite information structure and the state process encoded by its strategies. The terminal state includes a noise displacement and can leave the mixed-action simplexes, which motivates the payoff extension in Section 3.2. Put $d = |A_1| + |A_2|$, identify $\mathbb{R}^d$ with the two action blocks, and let $\nu > 0$ be the noise coefficient. The parameters are fixed in the order described in Section 5.1.

3.1 The Noise, the Dates, and the Types

Fix an integer $N \geq 1$, and let $h = 1/N$. We generate random objects in the following order.

Symbol	Meaning	Introduced
g, G_i	base game; player i's payoff matrix	Sections 1.1, 2.2
Z	exceptional set of payoff vectors	(31)
$R_{\kappa,\lambda}$	correction distribution	Lemma 7
P, T_i, Σ	original prior, type sets, original strategy space	Section 2.2
$q_i(t_i, \sigma)$	interim belief about the opponent's action	(3)
$w_i(t_i, \sigma)$	auxiliary (feedback) interim payoff vector	(4)
κ	uniform gap between w and the base interim payoff	(5)
S_i	set of good types of player i	Proposition 1
$\lambda = 2C\kappa$	output-block weight for a good type	Appendix A.4
Λ_l	total added unnormalized mass	(6)
$N, h = 1/N$	number of dates on the noise tree; step size	Section 3
$\xi_k, W_k, \mathcal{F}_k$	noise increments, noise path, its filtration	Section 3.1
K_i	player i's randomly sampled date	Section 3.1
Z_i^N, X_k^N	date-averaged action; accumulated state path	(9), (10)
$\mathcal{G}, V$	base payoff vector map and its smooth extension	(11)
η, f	regularization parameter; regularized best response	(12)
ν	diffusion coefficient	Section 3.1
v	solution of the parabolic system	Lemma 2
s	good-type threshold	Section 3.5
ε_*	upper bound on the probability of a bad original type	(19)
$\mu_{N,x}, \bar{\mu}_N$	action law of a profile; canonical action law	(14), (23)

Table 1: Standing notation.

1. Let $\xi_1, \dots, \xi_N$ be independent random vectors, with values in $\{-1, +1\}^d$, with each coordinate being -1 and 1 independently with equal probability, and then put

$$W_0 = 0, \qquad W_k = \frac{1}{\sqrt{N}} \sum_{\ell=1}^{k} \xi_\ell, \qquad \mathcal{F}_k = \sigma(\xi_1, \dots, \xi_k).$$

Thus $\mathbb{E}\xi_k = 0$, $\mathbb{E}\xi_k\xi_k^\top = I_d$, $|\xi_k| = \sqrt{d}$, and $\mathbb{E}\|W_N\|_2^2 = d$. The sequence $(\mathcal{F}_k)_{k \le N}$ is the natural filtration of a finite tree with 2^d branches at each node.

2. Independently of the above noise path, draw K_1 and K_2 independently and uniformly from the set $\{0, \dots, N-1\}$.

3. Define player i's type t_i as

$$t_i = (K_i, \xi_1, \dots, \xi_{K_i}).$$

Player i observes neither player j's date K_j nor the noise $\xi_{K_i+1}, \xi_{K_i+2}, \dots$ after date K_i.

4. Define the common prior P_N as the pushforward of the above experiment to the type space $T_1 \times T_2$. Every type has positive marginal probability: $P_N(t_i) = N^{-1} 2^{-dk}$ for a type with date

k. Pairs of inconsistent prefixes have zero probability, which is permitted under Definition 1. Conditionally on a type $t_i = (k, \xi_1, \ldots, \xi_k)$, the law of the noise path is exactly its $\mathcal{F}_k$-conditional law, because the date is independent of the path and the opponent's date remains uniform and independent.

This construction defines a finite information structure for the feedback game: the two players move once, simultaneously, and the only thing a player knows is its own type. It is *not* a dynamic game. The purpose of the dates is that a profile of behavioral strategies is then the same object as a sequence of adapted controls, so that this static information structure encodes a controlled recursion. The purpose of the common noise is that it correlates the two players' types.

5. We write x_i for player i's behavioral strategy on this information structure. Given date k and a noise path ω, we write $x_i(k, \omega) \equiv x_i(k, \xi_1(\omega), \ldots, \xi_k(\omega)) \in \Delta(A_i)$ for player i's mixed action at type $t_i = (k, \xi_1(\omega), \ldots, \xi_k(\omega))$. Note that each $x_i(k, \cdot)$ is $\mathcal{F}_k$-measurable by construction.

   Let $x(k, \omega) = (x_1(k, \omega), x_2(k, \omega)) \in \mathbb{R}^d$ be the mixed action profile for each (k, ω). Define the random vectors Z_i^N and the stochastic process X_k^N by

$$Z_i^N(x, \omega) = \frac{1}{N} \sum_{k=0}^{N-1} x_i(k, \omega), \tag{9}$$

$$X_k^N(x, \omega) = \frac{1}{N} \sum_{\ell < k} x(\ell, \omega) + \nu W_k(\omega). \tag{10}$$

   The vector Z_i^N is player i's date-averaged mixed action along the full path, and takes a value in $\Delta(A_i)$. The process X_k^N is $\mathcal{F}_k$-measurable, starting at $X_0^N = 0$, and satisfies the recursion

$$X_{k+1}^N = X_k^N + hx(k) + \nu\sqrt{h}\,\xi_{k+1}, \qquad X_N^N = Z^N + \nu W_N,$$

   where we write $Z^N = (Z_1^N, Z_2^N)$. Note that Z^N takes a value in a mixed-action set $\Delta(A_1) \times \Delta(A_2)$, but the noise-perturbed vector $X_N^N = Z^N + \nu W_N$ need not. As we will see, we still want to define a payoff for this X_N^N, which requires us to define the extension V, with domain $\mathbb{R}^d$, in Section 3.2.

   We suppress ω to mean random vectors, letting $x_i(k)(a)$ be the a-th coordinate of $x_i(k, \omega)$.

3.2 The Payoff Vector Map and the Regularized Best Response

On the set $\Delta(A_1) \times \Delta(A_2) \subset \mathbb{R}^d = \mathbb{R}^{|A_1|} \times \mathbb{R}^{|A_2|}$, we define the payoff vector by

$$\mathcal{G}_i(\sigma_j) = \left(\sum_{a_j} g_i(a_i, a_j)\sigma_j(a_j) \right)_{a_i}. \tag{11}$$

Thus, $\mathcal{G}(\sigma) = (\mathcal{G}_1(\sigma_2), \mathcal{G}_2(\sigma_1)) \in \mathbb{R}^d$ collects the players' expected payoff vectors given the opponents' mixed actions. Using the same formula, we extend the domain of this mapping $\mathcal{G} : \Delta(A_1) \times \Delta(A_2) \to \mathbb{R}^d$ to the set $\mathbb{R}^d$. We denote the extended mapping still by $\mathcal{G} : \mathbb{R}^d \to \mathbb{R}^d$.

We need an extension whose values and derivatives are bounded on all of $\mathbb{R}^d$. To construct one while preserving the expected payoffs near the product of simplexes, choose a smooth scalar function $\chi : \mathbb{R}^d \to [0, 1]$ that equals 1 on a neighborhood of the set $\Delta(A_1) \times \Delta(A_2)$ and equals 0 outside a sufficiently large ball. Such a function is called a *cutoff function*. We then define the function $V : \mathbb{R}^d \to \mathbb{R}^d$ by

$$V(z) = \chi(z)\mathcal{G}(z), \qquad \text{i.e.,} \qquad (V_1(z), V_2(z)) = (\chi(z)\mathcal{G}_1(z_2), \chi(z)\mathcal{G}_2(z_1)).$$

Consequently, V is smooth and compactly supported, so its values and all its derivatives are bounded and uniformly continuous (as required for Lemma 2 below). Fix a global Lipschitz constant $L > 0$, so that $\|V(z) - V(z')\|_2 \leq L\|z - z'\|_2$ for all $z, z' \in \mathbb{R}^d$. The extended payoff $\mathcal{G}_i$ depends only on player j's (extended) mixed action z_j, but not on player i's mixed action z_i, but since the cutoff χ may depend on both, this smooth extension $V_i(z)$ may depend on z_i away from the neighborhood at which $\chi(z) = 1$. No such independence is used below.

Fix $\eta > 0$ and define the regularized best response

$$f_i(y_i) = \underset{\sigma_i \in \Delta(A_i)}{\operatorname{argmax}} \{\sigma_i \cdot y_i - \eta\|\sigma_i\|_2^2\} = \operatorname{Proj}_{\Delta(A_i)}\left(\frac{y_i}{2\eta}\right), \tag{12}$$

where we have the second equality because $\sigma_i \cdot y_i - \eta\|\sigma_i\|_2^2 = -\eta\|\sigma_i - y_i/(2\eta)\|_2^2 + \|y_i\|_2^2/(4\eta)$. The two-block map $f = (f_1, f_2)$ takes values in the set $\Delta(A_1) \times \Delta(A_2)$, is bounded, and is globally $1/(2\eta)$-Lipschitz, since a Euclidean projection onto a convex set is 1-Lipschitz.[1] The unregularized best-response correspondence is bounded as well, its values being subsets of the simplexes; what it fails to provide is the single-valued globally Lipschitz map the lemma needs, since it is set-valued and its value jumps at ties. This is why η is introduced.

3.3 What a Type Believes and What the Game Induces

Lemma 1. *Fix N and a strategy profile x. For a type $t_i = (k, \xi_1, \ldots, \xi_k)$, let $q_i(t_i, x) \in \Delta(A_j)$ denote its conditional distribution of the opponent's action under P_N. Its base-game interim payoff vector satisfies*

$$G_i q_i(t_i, x) = \mathbb{E}[\mathcal{G}_i(Z_j^N(x)) \mid \mathcal{F}_k], \tag{13}$$

[1] To see the $1/(2\eta)$-Lipschitz continuity, we note that

$$\|f_i(y_i) - f_i(y_i')\|_2 = \left\|\operatorname{Proj}_{\Delta(A_i)}(y_i/(2\eta)) - \operatorname{Proj}_{\Delta(A_i)}(y_i'/(2\eta))\right\|_2 \leq \left\|(y_i - y_i')/(2\eta)\right\|_2 = \|y_i - y_i'\|_2/(2\eta).$$

Boundedness and Lipschitz continuity are needed for Lemma 2 below.

where the conditional expectation is evaluated at the observed prefix $(\xi_1, \dots, \xi_k)$*. The induced joint distribution of actions is*

$$\mu_{N,x}(a_1, a_2) = \mathbb{E}\big[Z_1^N(x)(a_1) Z_2^N(x)(a_2)\big]. \tag{14}$$

Proof. Fix $t_i = (k, \xi_1, \dots, \xi_k)$. Conditional on the complete noise path ω and on t_i, the opponent's date remains uniform because K_j is independent of (K_i, ω). The opponent's action distribution is therefore $N^{-1} \sum_{\ell=0}^{N-1} x_j(\ell, \omega) = Z_j^N(x, \omega)$. Averaging over the unobserved part of the path gives $q_i(t_i, x) = \mathbb{E}[Z_j^N(x) \mid \mathcal{F}_k]$, evaluated at the observed prefix, since K_i is independent of the noise path. Hence, for each action a_i,

$$\sum_{a_j} g_i(a_i, a_j) q_i(t_i, x)(a_j) = \mathbb{E}\left[\sum_{a_j} g_i(a_i, a_j) Z_j^N(x)(a_j) \,\middle|\, \mathcal{F}_k\right].$$

The sum inside the expectation is the a_i-th coordinate of $\mathcal{G}_i(Z_j^N(x))$, which proves (13).

For the joint action law, condition only on the complete noise path. The two dates and the players' behavioral randomizations are independent, so the conditional probability of (a_1, a_2) is $Z_1^N(x, \omega)(a_1) Z_2^N(x, \omega)(a_2)$. Averaging over the path gives (14). The product must remain inside the expectation: averaging over the shared noise can correlate the two players' actions. ■

3.4 The Auxiliary Payoffs and the Forward–Backward System

For each strategy profile $x = (x_1, x_2)$, with $x_i : T_i \to \Delta(A_i)$, define the following $\mathbb{R}^{|A_i|}$-valued random vector:

$$w_{i,k}(x) = \mathbb{E}[V_i(X_N^N(x)) \mid \mathcal{F}_k] - 2\eta x_i(k), \qquad k = 0, \dots, N-1. \tag{15}$$

This random vector is $\mathcal{F}_k$-measurable. *For a type* $t_i = (k, \xi_1, \dots, \xi_k)$*, its auxiliary interim payoff vector* $w_i(t_i, x)$ *is the value of* $w_{i,k}(x)$ *at that prefix.* For each fixed type, the map $x \mapsto w_i(t_i, x)$ is continuous, indeed smooth: $X_N^N(x)$ is affine in the finitely many strategy coordinates, V is smooth, and the conditional expectation is a finite weighted sum whose weights are determined by the fixed prior and independent of x.

For a given strategy profile x, write

$$Y_{i,k} = \mathbb{E}[V_i(X_N^N(x)) \mid \mathcal{F}_k] \in \mathbb{R}^{|A_i|}, \qquad Y_k = (Y_{1,k}, Y_{2,k}) \qquad \forall k = 0, \dots, N.$$

We suppress the dependence of Y_k on x in the notation. Then (15) becomes

$$w_{i,k}(x) = Y_{i,k} - 2\eta x_i(k).$$

Fix a type $t_i = (k, \xi_1, \dots, \xi_k)$. In the following comparisons, all random vectors are evaluated at this prefix. By Definition 4, the feedback equilibrium condition at this type is

$$\operatorname{supp} x_i(k) \subseteq \operatorname*{argmax}_{a_i} \{Y_{i,k}(a_i) - 2\eta x_i(k)(a_i)\}.$$

The vector $Y_{i,k} - 2\eta x_i(k)$ is evaluated at the candidate profile x and held fixed while its action coordinates are compared.

This support condition says that the mixed action $x_i(k)$ assigns positive probability only to actions with the highest value in this fixed vector. Equivalently, it gives at least as high an average value as every mixed action $\sigma_i \in \Delta(A_i)$:

$$\sigma_i \cdot \big(Y_{i,k} - 2\eta x_i(k)\big) \leq x_i(k) \cdot \big(Y_{i,k} - 2\eta x_i(k)\big).$$

Indeed, the support condition makes the right-hand side equal to the largest coordinate of the vector, which no weighted average can exceed. Conversely, if $x_i(k)$ assigns positive probability to an action with a strictly lower value, its average is strictly below that largest coordinate. Choosing σ_i to put all probability on a maximizing action would then violate the inequality.

Rearranging gives the equivalent *variational inequality*

$$(\sigma_i - x_i(k)) \cdot (Y_{i,k} - 2\eta x_i(k)) \leq 0 \qquad \forall \sigma_i \in \Delta(A_i). \tag{16}$$

To connect this condition to the regularized best response in (12), hold $Y_{i,k}$ fixed and consider

$$\max_{\sigma_i \in \Delta(A_i)} \Big\{ \sigma_i \cdot Y_{i,k} - \eta \|\sigma_i\|_2^2 \Big\}.$$

The gradient of this objective at $\sigma_i = x_i(k)$ is $Y_{i,k} - 2\eta x_i(k)$. Thus (16) says that moving from $x_i(k)$ toward any other feasible mixed action cannot increase the objective to first order.

Here the sufficiency of this condition can also be seen directly:

$$[\sigma_i \cdot Y_{i,k} - \eta\|\sigma_i\|_2^2] - [x_i(k) \cdot Y_{i,k} - \eta\|x_i(k)\|_2^2] = (\sigma_i - x_i(k)) \cdot (Y_{i,k} - 2\eta x_i(k)) - \eta\|\sigma_i - x_i(k)\|_2^2.$$

Under (16), the right-hand side is nonpositive and is strictly negative whenever $\sigma_i \neq x_i(k)$. Hence $x_i(k)$ is the unique maximizer. Conversely, if $x_i(k)$ maximizes the objective, the right derivative at $\epsilon = 0$ along $(1-\epsilon)x_i(k) + \epsilon\sigma_i$ must be nonpositive for every $\sigma_i \in \Delta(A_i)$. This gives (16).

By the definition of f_i in (12), the feedback equilibrium condition at this type is therefore equivalent to

$$x_i(k) = f_i(Y_{i,k}).$$

Throughout this argument, $Y_{i,k}$ is fixed at its value for the candidate profile x. Differentiation with respect to the comparison mixed action σ_i does not differentiate the dependence of $Y_{i,k}$ on x. The coefficient 2η in the auxiliary payoff was chosen precisely because the gradient of the quadratic penalty $-\eta\|\sigma_i\|_2^2$ is $-2\eta\sigma_i$. It makes the feedback condition coincide with the optimality condition defining the single-valued Lipschitz map f_i.

Combining the two players' conditions gives $x(k) = f(Y_k)$. Substituting this into the state recursion of Section 3.1, and suppressing the dependence of X^N on x, shows that every feedback equilibrium satisfies

$$\begin{aligned} X_0^N &= 0, \\ X_{k+1}^N &= X_k^N + hf(Y_k) + \nu\sqrt{h}\,\xi_{k+1}, \quad & k &= 0,\dots,N-1, \\ Y_k &= \mathbb{E}[V(X_N^N) \mid \mathcal{F}_k], & k &= 0,\dots,N. \end{aligned} \tag{17}$$

This is a coupled forward–backward system on the fixed finite noise tree.[2]

The *forward* equation starts from $X_0^N = 0$ and updates the state by adding the regularized response $hf(Y_k)$ and the next noise increment. The *backward* equation takes the terminal payoff vector $V(X_N^N)$ and averages it using the information available at date k. In particular,

$$Y_N = V(X_N^N), \qquad Y_k = \mathbb{E}[Y_{k+1} \mid \mathcal{F}_k], \qquad k = 0,\dots,N-1,$$

where the second equality follows from iterated conditional expectations. This explains the two directions: X^N is linked to its initial value, while Y is linked to its terminal value.

The equations are coupled because updating X^N requires Y_k, and Y_k depends on the terminal state X_N^N. They must therefore be satisfied together. The strategies may depend on the entire observed prefix; no assumption that they depend only on X_k^N has been imposed.

The implication needed below is that every feedback equilibrium produces an adapted solution of (17). The comparison in Section 4 will give estimates that hold uniformly over all such solutions, and hence apply to every feedback equilibrium.

3.5 Good Types and the Uniform Payoff Gap

We verify the probability and payoff bounds needed to apply Proposition 1 to the auxiliary problem of Section 3. Fix N and let $x \in \Sigma$ be arbitrary. Since $Z^N(x)$ belongs to the product of simplexes, the cutoff extension satisfies

$$V_i(Z^N(x)) = \mathcal{G}_i(Z_j^N(x)).$$

For a type $t_i = (k, \xi_1, \dots, \xi_k)$, equations (13) and (15) give

$$w_i(t_i, x) - G_i q_i(t_i, x) = \mathbb{E}\Big[V_i(Z^N(x) + \nu W_N) - V_i(Z^N(x)) \mid \mathcal{F}_k\Big] - 2\eta x_i(k),$$

[2] The type space, prior P_N, and information $\mathcal{F}_k$ are fixed before any strategy is chosen. For each x, the process $X^N(x)$ is computed on this same noise tree. Different feedback equilibria may therefore produce different state and action distributions, but they do not change the information structure.

where the right-hand side is evaluated at the observed prefix $(\xi_1, \dots, \xi_k)$. The two contributions are the displacement caused by the noise and the regularization term. Since $x_i(k)$ is a mixed action, the latter satisfies $\|2\eta x_i(k)\|_\infty \le 2\eta$.

The noise contribution can be bounded using the global L-Lipschitz property of V. For every behavioral profile x, every noise path, and every action a_i,

$$\begin{aligned}&\Big|V_i(Z^N(x) + \nu W_N)(a_i) - \mathcal{G}_i(Z_j^N(x))(a_i)\Big| \\ &\qquad \le \Big\|V(Z^N(x) + \nu W_N) - V(Z^N(x))\Big\|_2 \le L\nu \|W_N\|_2.\end{aligned} \tag{18}$$

Although $Z^N(x)$ depends on the strategy profile, the final bound does not. Taking conditional expectations coordinate by coordinate and adding the regularization bound yields

$$\|w_i(t_i, x) - G_i q_i(t_i, x)\|_\infty \le \mathbb{E}[L\nu\|W_N\|_2 \mid \mathcal{F}_k] + 2\eta,$$

with the right-hand side again evaluated at the type's prefix. This bound holds for every $x \in \Sigma$.

Fix $s > 0$ as a tolerance for the noise contribution. Call an original type $t_i = (k, \text{prefix})$ *good* if

$$\mathbb{E}[L\nu\|W_N\|_2 \mid \mathcal{F}_k] \le s$$

at that prefix, and *bad* otherwise. Let S_i be the set of good types of player i. This criterion depends only on the date and the observed noise prefix. Hence the sets S_i are fixed independently of x, as required by Proposition 1.

We next bound the probability of bad types. Since $\mathbb{E}\|W_N\|_2^2 = d$,

$$\mathbb{E}\|W_N\|_2 \le \big(\mathbb{E}\|W_N\|_2^2\big)^{1/2} = \sqrt{d}.$$

Thus the unconditional expectation of the noise bound is at most $L\nu\sqrt{d}$, independently of N. This upper bound is controlled by ν and does not decrease as the tree is refined.

For a fixed date k, the conditional expectation defining good types is a nonnegative random variable. Markov's inequality and the tower property give

$$\begin{aligned}\mathbb{P}(\mathbb{E}[L\nu\|W_N\|_2 \mid \mathcal{F}_k] > s) &\le \frac{\mathbb{E}[\mathbb{E}[L\nu\|W_N\|_2 \mid \mathcal{F}_k]]}{s} \\ &= \frac{L\nu\,\mathbb{E}\|W_N\|_2}{s} \le \frac{L\nu\sqrt{d}}{s}.\end{aligned}$$

This bound holds at every date. Since each player's date is uniform and independent of the noise path, that player's probability of being a bad type is the average of the corresponding probabilities over the N dates and obeys the same bound. Taking the union bound over the two players,

$$P_N\big((S_1 \times S_2)^c\big) \le \varepsilon_* := \frac{2L\nu\sqrt{d}}{s}. \tag{19}$$

Equivalently, $P_N(S_1 \times S_2) \geq 1 - \varepsilon_*$. The bound ε_* is independent of N and becomes small when ν/s is small.

At every good type, the conditional noise bound is at most s. Consequently,

$$\|w_i(t_i, x) - G_i q_i(t_i, x)\|_\infty \leq \kappa \equiv s + 2\eta \qquad \text{for every } t_i \in S_i \text{ and every } x \in \Sigma. \tag{20}$$

Thus the same designated sets satisfy both the probability bound and the uniform payoff bound required by Proposition 1, with $\varepsilon = \varepsilon_*$ and $\kappa = s + 2\eta$. At bad types, the auxiliary payoff maps remain bounded and continuous for each fixed N; the proposition requires no smallness bound there. The parameters will be chosen in Section 5 so that $\kappa < \kappa_0$.

4 Asymptotic Determinacy of Feedback Outcomes

Every feedback equilibrium yields an adapted solution of (17). We compare all these solutions to one recursion obtained from a parabolic equation. The analysis is valid in arbitrary finite dimension and for any finite partition into player blocks. Its constants may depend on the fixed data η, ν, V, d and a Hölder exponent β, but not on N or on the candidate solution.

4.1 The Comparison Equation and Canonical Recursion

Consider a solution of (17) and ask what would have to be true if the backward variable were a smooth function of time and of the forward state, say $Y_k = v(t_k, X_k)$ with $t_k = kh$. The definition $Y_k = \mathbb{E}[V(X_N) \mid \mathcal{F}_k]$ makes (Y_k) a martingale, so $Y_k = \mathbb{E}[Y_{k+1} \mid \mathcal{F}_k]$ and hence

$$\mathbb{E}\big[v(t_{k+1}, X_{k+1}) - v(t_k, X_k) \mid \mathcal{F}_k\big] = 0.$$

Insert $X_{k+1} = X_k + hf(Y_k) + \nu\sqrt{h}\,\xi_{k+1}$ and expand. The increment has conditional mean $hf(Y_k)$ and conditional covariance $\nu^2 h I_d$ up to $O(h^2)$, so a second-order Taylor expansion in space and a first-order expansion in time give

$$h\left(\partial_t v + Dv\, f(v) + \frac{\nu^2}{2}\Delta v\right)(t_k, X_k) + o(h) = 0.$$

Dividing by h produces the equation. The terminal condition is forced by $Y_N = V(X_N)$.

So the object to study is

$$\partial_t v + Dv\, f(v) + \frac{\nu^2}{2}\Delta v = 0, \qquad v(1, x) = V(x), \tag{21}$$

a system of d equations for the $\mathbb{R}^d$-valued unknown v, semilinear in the sense that the second-order term is the constant-coefficient Laplacian and the nonlinearity $Dv\, f(v)$ involves only v and its first derivatives. Note the direction of time: the data are given at $t = 1$ and the equation is solved backwards, matching the backward variable it is meant to describe.

The logic of the section is the reverse of the heuristic above. We do not assume that Y_k is a smooth function of X_k — it is not, in general, since the controls may depend on the whole prefix. We solve (21) first, and then use its solution as a *comparison object* against which every adapted solution of the finite system is measured.

Lemma 2 (A Bounded-Drift Parabolic System)**.** *Let $f : \mathbb{R}^d \to \mathbb{R}^d$ be bounded and globally Lipschitz, let $\nu > 0$, and let $V : \mathbb{R}^d \to \mathbb{R}^d$ be smooth with bounded uniformly continuous derivatives of every order. Then (21) has a bounded classical solution with bounded spatial gradient on $[0,1] \times \mathbb{R}^d$. It is unique among mild solutions that are continuous in time with values in the space of bounded uniformly continuously differentiable functions and have bounded norm in that space. For every $0 < \beta < 1$, its Hessian and time derivative are uniformly β-Hölder in space and uniformly $\beta/2$-Hölder in time. In particular, their bounds and these Hölder constants are global in x.*

The proof is given in Appendix B.1. It is a self-contained argument by heat-semigroup fixed point, an a priori gradient bound, and a bootstrap; nothing in it interacts with the game-theoretic part of the paper. Three of its features are worth recording here, because they are used in ways that a weaker statement would not support.

Global existence on the whole interval. The system is quasilinear, and for such systems solutions may cease to exist in finite time. Here they do not, and the reason is structural: the nonlinearity is $Dv\, f(v)$ with f *bounded*, so the gradient satisfies a *linear* Volterra inequality of the form (45), whose Gronwall bound (46) holds on all of $[0,1]$ regardless of the Lipschitz constant of f. This matters because the Lipschitz constant of f is $1/(2\eta)$, which is large precisely in the regime we need.

Global-in-space bounds. All bounds are uniform in $x \in \mathbb{R}^d$, not merely on compact sets. This is essential: the discrete state X_k^N of Section 3 is a random walk displacement added to a bounded drift, so it can be far from the origin, and the comparison must hold on every branch of the noise tree, including improbable ones.

Hölder regularity of the Hessian and of $\partial_t v$. These are exactly the ingredients of the one-step expansion (50): the spatial expansion is carried to second order with a β-Hölder remainder on the Hessian, and the temporal expansion to first order with a $\beta/2$-Hölder remainder on $\partial_t v$. Since the noise increments scale like $\sqrt{h}$, a β-Hölder Hessian converts into a one-step error $O(h^{1+\beta/2})$, which summed over $N = 1/h$ steps leaves $O(h^{\beta/2})$. The global Hölder estimates are what supply this explicit power rate. A uniform modulus of continuity for the Hessian would give a uniform $o(h)$ per step, which still sums to $o(1)$, but with no rate; and mere pointwise continuity on an unbounded domain would not deliver the uniformity of the one-step estimate in the first place.

Let v be the solution of Lemma 2. For $h = 1/N$, $t_k = kh$, a deterministic initial state x_0, and a given noise sequence, define the canonical recursion by

$$
\begin{aligned}
\widehat{X}_0 &= x_0, \qquad \widehat{s}_k = f(v(t_k, \widehat{X}_k)), \\
\widehat{X}_{k+1} &= \widehat{X}_k + h\widehat{s}_k + \nu\sqrt{h}\,\xi_{k+1}.
\end{aligned}
\tag{22}
$$

For the noise tree of Section 3.1, use $x_0 = 0$. Its controls lie in the product of action simplexes, so independent uniform date sampling defines an action distribution $\bar{\mu}_N$, equivalently

$$\widehat{Z}_i^N = h \sum_{k=0}^{N-1} \widehat{s}_{i,k}, \qquad \bar{\mu}_N(a) = \mathbb{E} \prod_i \widehat{Z}_i^N(a_i).$$

Thus $\bar{\mu}_N$ is defined from the PDE solution and the fixed noise tree before it is used to compare feedback outcomes.

4.2 Asymptotic Determinacy

Proposition 2 (Asymptotic Determinacy of Feedback Action Laws)**.** *Fix $\eta, \nu > 0$ and V as in Section 3.2, and let $\bar{\mu}_N$ be the action law induced by the canonical recursion (22) with $x_0 = 0$ under the date sampling of Section 3.1. Then*

$$\sup_{x \text{ a feedback equilibrium of } (15)} \|\mu_{N,x} - \bar{\mu}_N\|_{\mathrm{TV}} \to 0 \qquad (N \to \infty). \tag{23}$$

The conclusion is uniform over all feedback equilibria. It gives approximate agreement of their induced action distributions, without asserting that the feedback equilibrium itself is unique.

4.3 From Adapted Solutions to Action Distributions

The comparison applies to history-dependent controls, not just controls that are functions of the current state. Its one-step consistency estimate gives a backward error bound; subtracting the canonical recursion then cancels the common noise pathwise and bounds the states and controls. The following precise statement also allows an error ρ in the response condition, although our application uses $\rho = 0$.

Lemma 3 (Comparison of Every Adapted Finite-Noise Solution)**.** *Under the assumptions of Lemma 2, fix $0 < \beta < 1$ and a bound $M < \infty$ for noise increments, and let v be the solution of (21). Put $h = 1/N$ and $t_k = kh$. On any filtered probability space, suppose ξ_{k+1} is $\mathcal{F}_{k+1}$-measurable and independent of $\mathcal{F}_k$, has mean zero and covariance I_d, and satisfies $|\xi_{k+1}| \le M$. Let $X_0 = x_0$ be deterministic, and let s_k be $\mathcal{F}_k$-measurable controls satisfying, for some $0 \le \rho \le 1$,*

$$X_{k+1} = X_k + h s_k + \nu\sqrt{h}\,\xi_{k+1}, \tag{24}$$

$$Y_k = \mathbb{E}[V(X_N) \mid \mathcal{F}_k], \qquad |s_k - f(Y_k)| \le \rho. \tag{25}$$

Then, for all sufficiently large N,

$$\max_{0 \le k \le N} \|Y_k - v(t_k, X_k)\|_\infty \le C(\rho + h^{\beta/2}), \tag{26}$$

where $\|\cdot\|_\infty$ for a random vector means its essential supremum. For the canonical recursion (22) on the same noise path, almost surely,

$$\max_{0\le k\le N} |X_k - \widehat{X}_k| + \max_{0\le k<N} |s_k - \widehat{s}_k| \le C'(\rho + h^{\beta/2}). \tag{27}$$

The constants C, C' and the lower bound on N depend only on the fixed data, including M and β. They are uniform over candidate controls, over filtrations, and over N.

The detailed consistency estimate and the backward and forward error calculations are in Appendix B.2. One set of constants applies to every adapted solution, which is essential when the controls are generated by an arbitrary feedback equilibrium.

Corollary 2 (Vanishing Diameter of Time-Agent Action Laws)**.** *Suppose the filtration is the natural filtration of a fixed finite noise tree, the coordinates of $\mathbb{R}^d$ are partitioned into player blocks, the controls satisfy $s_{i,k} \in \Delta(A_i)$, and f maps into $\prod_i \Delta(A_i)$. Conditionally on the complete noise path, sample each player's date independently and uniformly from $\{0, \ldots, N-1\}$ and let that player's action be drawn from the indicated behavioral control at that date. Then, over all exact solutions of (24)–(25) with $\rho = 0$, the induced action distributions on $\prod_i A_i$ have total variation diameter at most $CN^{-\beta/2}$.*

Proof. Fix a noise path and put $Z_i = h\sum_{k=0}^{N-1} s_{i,k}$ and $\widehat{Z}_i = h\sum_{k=0}^{N-1} \widehat{s}_{i,k}$, both in $\Delta(A_i)$. Conditionally on that path, the dates are independent and uniform and the behavioral randomizations are independent, so the conditional action distribution is the product $\bigotimes_i Z_i$, and likewise $\bigotimes_i \widehat{Z}_i$ for the canonical recursion. For probability vectors, the hybrid argument — replace one coordinate at a time and use the triangle inequality — gives

$$\Big\| \bigotimes_i Z_i - \bigotimes_i \widehat{Z}_i \Big\|_{\mathrm{TV}} \le \sum_i \|Z_i - \widehat{Z}_i\|_{\mathrm{TV}}.$$

Each summand is bounded by a constant multiple of $\max_k |s_{i,k} - \widehat{s}_{i,k}|$, since $Z_i - \widehat{Z}_i$ is an average of the $s_{i,k} - \widehat{s}_{i,k}$ and all norms on $\mathbb{R}^{A_i}$ are equivalent, hence by $CN^{-\beta/2}$ using (27) with $\rho = 0$. The bound holds on every noise path and for every candidate exact solution, so it survives averaging over the path: every candidate action law is within $CN^{-\beta/2}$ of the single law induced by the canonical recursion (22). The triangle inequality gives the stated diameter bound. No Brownian limit and no convergence assertion for filtrations is involved. ■

We now prove Proposition 2. Recall from Section 3.2 that f is bounded and $1/(2\eta)$-Lipschitz and V is smooth with bounded derivatives of every order, so Lemmas 2 and 3 apply with these data, with $M = \sqrt{d}$ and with any fixed $\beta \in (0, 1)$.

Proof of Proposition 2. By Section 3.4, every feedback equilibrium x of (15) yields controls $s_k = x(k) = f(Y_k)$ that satisfy (24)–(25) with $\rho = 0$, on the natural filtration of the noise tree, with $X_k = X_k^N(x)$ and $X_0 = 0$. By (14) the induced action law $\mu_{N,x}$ is exactly the law described in Corollary 2, since $Z_i^N(x)$ is the date average of the controls. That corollary bounds $\|\mu_{N,x} - \bar{\mu}_N\|_{\mathrm{TV}}$

by $CN^{-\beta/2}$, with C independent of x and of N. Taking the supremum over feedback equilibria and letting $N \to \infty$ gives (23). ■

5 Proof of the Main Result for Two Players

Proof of Theorem 1 for two players. If $|I| = 1$ and $g \notin Z$, exclusion of the pure payoff ties gives the sole player a unique maximizing action. The complete-information game itself, with one type and payoff function g_1, is therefore a finite 0-elaboration with a unique BNE outcome, proving the theorem in this case for every $\varepsilon, \rho > 0$. We now consider two players, fix $g \notin Z$, and fix $\varepsilon, \rho > 0$.

5.1 Choosing the Parameters for Both Tolerances

Take κ_0, D and $K = 1+2(m-1)$ from Proposition 1, where $m = \max\{|A_1|, |A_2|\}$. These constants depend only on the base game and its action sets. Choose $s, \eta > 0$ so that

$$\kappa = s + 2\eta < \kappa_0, \qquad D\kappa < \min\{\varepsilon/4, \rho/16\}.$$

Next choose $\nu > 0$ so that the bound in (19) satisfies

$$K\varepsilon_* = \frac{2KL\nu\sqrt{d}}{s} < \min\{\varepsilon/4, \rho/16\}.$$

Here V and its global Lipschitz constant L are those of Section 3.2. Fix any $\beta \in (0,1)$. All the data of the parabolic comparison are now fixed; its constants may depend on η and ν, but are finite and will not change as N varies.

By Proposition 2, choose N sufficiently large that

$$\sup_{y \text{ a feedback equilibrium of (15)}} \|\mu_{N,y} - \bar{\mu}_N\|_{\mathrm{TV}} < \rho/8.$$

Keep this N fixed. The good-type sets of Section 3.5 are independent of the strategy profile and satisfy (19) and (20). Apply Proposition 1 to the resulting finite feedback game with ε_* and κ in place of its two error parameters. This gives a sequence of finite elaborations $E_{N,l}$ with normal original good types. For this fixed N, write Λ_l for their added-to-original probability ratio and $\tau_l \to 0$ for the residual mass bound in that proposition.

5.2 Uniform Approximation of All Equilibrium Restrictions

Let Σ be the original strategy space for the fixed level N. It is a finite product of simplexes and hence compact. The feedback equilibrium set is closed, since its defining conditions are the finitely many continuous equalities

$$x_i(a_i \mid t_i)\left(\max_{b_i \in A_i} w_i(t_i, x)(b_i) - w_i(t_i, x)(a_i)\right) = 0 \qquad \text{for every } i, t_i, a_i.$$

It is also nonempty: choose a BNE of each finite elaboration $E_{N,l}$, take a convergent subsequence of its restrictions in Σ, and apply Proposition 1(c).

In fact, every open neighborhood of this feedback equilibrium set contains all BNE restrictions of $E_{N,l}$ for every sufficiently large l. To prove the uniform assertion, suppose it fails for some open neighborhood U. Choose increasing indices l_j and BNE restrictions $x^{l_j} \notin U$. Compactness supplies a convergent subsequence of these restrictions. At the indices not chosen, select arbitrary BNE of the corresponding elaborations, so that Proposition 1(c) applies to a sequence with one equilibrium at every index. It places the subsequential limit in the feedback equilibrium set. But that limit lies in the closed complement of U, a contradiction. Thus part (c), although phrased in terms of limit points of arbitrarily selected equilibrium sequences, gives uniform proximity over all equilibria at this fixed N.

The map $x \mapsto \mu_{N,x}$ is polynomial and therefore continuous. The open set

$$U = \{x \in \Sigma : \|\mu_{N,x} - \bar{\mu}_N\|_{\mathrm{TV}} < \rho/4\}$$

contains every feedback equilibrium by the choice of N. Consequently, for every sufficiently large l, every BNE σ of $E_{N,l}$ has an original-type restriction x satisfying

$$\|\mu_{N,x} - \bar{\mu}_N\|_{\mathrm{TV}} < \rho/4.$$

No strategy or equilibrium is selected in this conclusion.

5.3 The Diameter of the Full BNE Outcomes

Choose one l sufficiently large for the preceding uniform conclusion and for

$$\tau_l < \min\{\varepsilon/4, \rho/16\}.$$

By Proposition 1(b), the probability that some type is not normal in $E_{N,l}$ is at most

$$K\varepsilon_* + D\kappa + \tau_l < 3\varepsilon/4 < \varepsilon.$$

Thus $E_{N,l}$ is a finite ε-elaboration. Part (a) of the same proposition also gives

$$\Lambda_l \leq D\kappa + (K-1)\varepsilon_* + \tau_l < 3\rho/16.$$

For a BNE σ of this elaboration, let x be its original-type restriction and let χ_σ be the unnormalized action measure generated by the added states. The original states retain total unnormalized mass one and conditional prior P_N, so

$$\mu_\sigma = \frac{\mu_{N,x} + \chi_\sigma}{1 + \Lambda_l}, \qquad \chi_\sigma(a) \geq 0, \qquad \sum_{a \in A} \chi_\sigma(a) = \Lambda_l. \tag{28}$$

The added states therefore have probability $\Lambda_l/(1+\Lambda_l)$. When $\Lambda_l > 0$, equation (28) is a mixture of the two probability laws $\mu_{N,x}$ and χ_σ/Λ_l; when $\Lambda_l = 0$, its added term is zero. In both cases,

$$\|\mu_\sigma - \mu_{N,x}\|_{\mathrm{TV}} \le \frac{\Lambda_l}{1+\Lambda_l} \le \Lambda_l < 3\rho/16.$$

Combining this with the uniform bound on the original restrictions yields, for every BNE σ,

$$\|\mu_\sigma - \bar{\mu}_N\|_{\mathrm{TV}} < \rho/4 + 3\rho/16 = 7\rho/16.$$

The triangle inequality now gives

$$\sup_{\sigma,\sigma' \in \mathrm{BNE}(E_{N,l})} \|\mu_\sigma - \mu_{\sigma'}\|_{\mathrm{TV}} \le 7\rho/8 < \rho.$$

This proves the theorem for two players, including all equilibria and the entire prior distribution on original and added states. ■

Remark 4 (The order of the limits)**.** The construction first fixes the two tolerances ε, ρ, then s, η, and then ν. Only after these choices have fixed all analytic constants do we choose a sufficiently large N. The compactness argument and the limit in Proposition 1(c) are used only at this fixed N to choose a sufficiently large implementation index l. No limit in N is interchanged with a limit in l, and no analytic estimate is required to be uniform in η or ν. Nor do we assume that $\Lambda_l \to 0$ for fixed perturbation parameters: the possibly nonvanishing bound $D\kappa + (K-1)\varepsilon_*$ is made small through the earlier parameter choices, and only the residual mass τ_l vanishes with l.

6 Extension to Any Finite Number of Players

We extend the construction of Theorem 1 to every finite player set. The implementation must produce joint distributions over opponents' action profiles, and the conditional-expectation identities must retain the correlation generated by the common noise. After verifying these two points, we apply the same fixed-N argument and estimate the complete BNE outcomes.

Put $n = |I|$. The case $n = 1$ was settled in Section 5, so assume $n \ge 2$. Whenever player i is a recipient, fix a host player $j \in I \setminus \{i\}$.

6.1 Beliefs Over Opponents' Profiles

For $i \in I$, the relevant belief is the joint distribution

$$q_i(t_i,\sigma)(c) = \sum_{t_{-i}} P(t_{-i} \mid t_i) \prod_{j \ne i} \sigma_j(c_j \mid t_j), \qquad c \in A_{-i}.$$

The payoff matrix G_i has $m_i = |A_i|$ rows and $k_i = |A_{-i}|$ columns, as in the definition of Z in Appendix A.1. Lemma 7 applies with these arbitrary finite dimensions. The interim payoff vector is still $G_i q_i(t_i, \sigma)$, since a player's payoff is linear in the joint distribution of the opponents' profiles.

For $n \geq 3$, the belief $q_i(t_i, \sigma) \in \Delta(A_{-i})$ obtained from a common prior need not factor across opponents, and neither need the correction $r_i(\sigma)$ produced by Lemma 7. The belief correction acts on the whole simplex $\Delta(A_{-i})$. Its implementation must therefore deliver an arbitrary correlated distribution over A_{-i} at the output states.

6.2 Implementing Correlated Opponent Actions

The implementation in Appendix A.5 extends the gates, probes, and output states of Lemma 8. Fresh committed types fill unused positions in each state. A controller reads a coordinate of the host player's original strategy directly, while probes transfer coordinates of any other player, including the recipient, to quantities carried by the host. This preserves the input estimate (39) and leaves every original type free to receive its exact base payoff at every added state where it appears.

To produce a correlated output, enumerate $A_{-i} = \{c_1, \ldots, c_{\bar{k}}\}$ and use the cumulative probabilities of the target $r_i(\sigma)$, as in the output-label construction of Lemma 8. For each label, place the recipient and one final output type of every opponent in a single output state. All these output types read the same flags. At a clean label, the telescoping score calculation forces every opponent onto its coordinate of the same profile c_p. Mixing over labels therefore implements the target joint distribution, with the correlation carried by the common label in the prior.

At an unclean label, the opponents may independently randomize in any way allowed by their incentives. The error estimate (41) uses only the total mass of these labels and therefore still bounds the error in the joint distribution on A_{-i} at every BNE. The bad-type construction uses a binary target for one opponent in $I \setminus \{i\}$ and commits the remaining opponents. Appendix A.5 gives the details, including the preservation of exact normality; the corresponding mass bounds are recorded in Section 6.4.

6.3 The Prior and the Conditional-Expectation Identities

Set $d = \sum_{i \in I} |A_i|$ and write $z \in \prod_{i \in I} \mathbb{R}^{A_i}$. Fix $\nu > 0$ and construct the d-dimensional noise and its prefixes as in Section 3.1. Draw independent uniform dates K_i for all players and give player i its own date and the shared noise prefix. Replace (11) by

$$\mathcal{G}_{i,a_i}(z) = \sum_{c \in A_{-i}} g_i(a_i, c) \prod_{j \neq i} z_j(c_j), \qquad \forall i \in I, \quad a_i \in A_i. \tag{29}$$

This is a polynomial of degree at most $n - 1$. Multiplying by the same smooth cutoff gives a compactly supported V with bounded derivatives of every order and a global Lipschitz constant L. These are the properties of V used in Sections 4 and 3.5; neither argument requires linearity.

For a type $t_i = (k, \text{prefix})$, the identities of Lemma 1 take the form

$$q_i(t_i, x)(c) = \mathbb{E}\left[\prod_{j \neq i} Z_j^N(x)(c_j) \,\middle|\, \mathcal{F}_k \right], \qquad c \in A_{-i},$$

$$\mu_{N,x}(a) = \mathbb{E}\prod_{i\in I} Z_i^N(x)(a_i), \qquad a \in A,$$

where the first conditional expectation is evaluated at the observed prefix. Conditionally on the complete noise path, the opponents' dates are independent of one another and their behavioral randomizations are independent, so the conditional joint distribution of the opponents' actions is the product measure $\bigotimes_{j\neq i} Z_j^N(x)$. Conditioning on the type then averages this over the unobserved part of the path, and the product stays inside the expectation. Multiplying by $g_i(a_i, c)$ and summing over $c \in A_{-i}$ gives

$$\big(G_i q_i(t_i, x)\big)_{a_i} = \mathbb{E}[\mathcal{G}_{i,a_i}(Z^N(x)) \mid \mathcal{F}_k],$$

which is (13) in unchanged form. Only linearity of conditional expectation over the finitely many coordinates $\prod_{j\neq i} Z_j^N(x)(c_j)$ has been used. Keeping the product inside the expectation retains the correlation created by the unobserved common noise; the interim payoff cannot in general be computed from the opponents' marginal conditional action distributions alone.

The auxiliary payoffs are (15) with i ranging over I. The variational characterization in Section 3.4 applies to each player's simplex, and the resulting system is again (17). Section 4 treats arbitrary finite dimension and arbitrary partitions of coordinates into player blocks, so Lemma 3 and Corollary 2 apply with n blocks. In the latter proof, the total variation distance between product laws is bounded by the sum of the marginal distances over these blocks. This changes only the constant and gives Proposition 2 with the canonical law $\bar{\mu}_N$ defined from the corresponding d-dimensional recursion.

6.4 Constants and Completion

Put $m = \max_i |A_i|$ and let C be the maximum over $i \in I$ of the constants C_{G_i} from Lemma 7. Choose $\kappa_0 > 0$ small enough that every player's correction map is available for $0 < \kappa < \kappa_0$ with the common mass $\lambda = 2C\kappa \le 1$. The constants of Proposition 1 and the good-type estimate become

$$K = 1 + n(m-1), \qquad D = 2nC, \qquad \varepsilon_* = nL\nu\sqrt{d}/s.$$

The constants κ_0 and D depend only on g and remain uniform over finite type spaces, priors, and continuous auxiliary payoff maps.

For good sets satisfying $P(\prod_i S_i) \ge 1 - \varepsilon$, the output-mass bound in the realization proof in Appendix A becomes

$$\lambda \sum_i P_i(S_i) + (m-1)\sum_i P_i(S_i^c) \le 2nC\kappa + n(m-1)\varepsilon.$$

Here $\sum_i P_i(S_i) \le n$ and $P_i(S_i^c) \le \varepsilon$ for each i, because $(\prod_i S_i)^c$ contains $S_i^c \times \prod_{j\neq i} T_j$. The normality estimate (19) uses the union bound over the n players. The good-type payoff estimate (20) remains valid with $\kappa = s + 2\eta$. Together with the implementation above, these observations give parts (a)–(c) of Proposition 1 with the displayed constants. The proof of part (d), based on the same decomposition into original and added states, also carries over unchanged.

Fix the desired tolerances $\varepsilon, \rho > 0$. As in Section 5.1, choose $s, \eta > 0$ and then $\nu > 0$ so that

$$\kappa = s + 2\eta < \kappa_0, \qquad D\kappa < \min\{\varepsilon/4, \rho/16\}, \qquad K\varepsilon_* < \min\{\varepsilon/4, \rho/16\}.$$

Fix these parameters before choosing N. The finite-player version of Proposition 2 just proved permits us to fix N sufficiently large that all feedback outcomes are within $\rho/8$ of $\bar{\mu}_N$. Apply the realization construction to this fixed noise tree and its fixed good-type sets. The compactness argument of Section 5.2 applies to the finite product of strategy simplexes for these players: part (c) rules out a sequence of BNE restrictions staying away from the feedback-equilibrium set. Hence, for all sufficiently large l, every restricted BNE action law is within $\rho/4$ of $\bar{\mu}_N$. Increase l if necessary so that $\tau_l < \min\{\varepsilon/4, \rho/16\}$, and keep this one finite elaboration $E_{N,l}$.

Its abnormal probability is less than $3\varepsilon/4$, so it is an ε-elaboration. Its added-state ratio satisfies $\Lambda_l < 3\rho/16$. For every BNE σ with original restriction x, decomposition (28) gives

$$\|\mu_\sigma - \mu_{N,x}\|_{\mathrm{TV}} \le \frac{\Lambda_l}{1 + \Lambda_l} \le \Lambda_l.$$

Combining this bound with the uniform estimate for restricted BNE action laws, exactly as in Section 5.3, yields

$$\sup_{\sigma, \sigma' \in \mathrm{BNE}(E_{N,l})} \|\mu_\sigma - \mu_{\sigma'}\|_{\mathrm{TV}} \le \frac{\rho}{2} + 2\Lambda_l \le \frac{7\rho}{8} < \rho.$$

Thus all complete BNE outcomes of this finite ε-elaboration have the required diameter. This proves Theorem 1 for every finite player set. The implementation accuracy is chosen after N is fixed, and the added-state probability is controlled by the chosen parameters; no claim that $\Lambda_l \to 0$ is used.

Remark 5 (The two extensions used above)**.** The realization argument requires a joint correction distribution on A_{-i}, implemented by common output labels and an error bound that tolerates arbitrary mixing at unclean labels. The noise-tree argument requires the product of date averages inside the conditional expectation and permits the resulting payoff map to be a polynomial of degree at most $n - 1$. These facts preserve the incentive and comparison arguments for arbitrary finite numbers of players.

A Implementation

A.1 Genericity and Elementary Facts

For each player i, we consider her payoffs g_i as an m_i-by-k_i matrix G_i, where $m_i = |A_i|$ and $k_i = |A_{-i}|$, and denote the (a_i, c)-th entry by $g_i(a_i, c)$. Letting $q \in \Delta(A_{-i})$ denote a belief about the opponents' actions, we write $G_i q \in \mathbb{R}^{A_i}$ for the interim expected payoff vector. Form the row-labeled matrix

$$\mathcal{N}_i(G_i) = \begin{bmatrix} e_1 \\ \vdots \\ e_{k_i} \\ G_{i,a} - G_{i,b} \quad \forall a < b \\ \mathbf{1}^\top \end{bmatrix} \in \mathbb{R}^{(k_i + \binom{m_i}{2} + 1) \times k_i}, \tag{30}$$

whose rows are the k_i coordinate rows, the $\binom{m_i}{2}$ payoff-difference rows, and the all-ones row.

Let Z_i be the union of the zero sets of all k_i-by-k_i determinant polynomials obtained by selecting distinct row labels from (30), *omitting those polynomials that are identically zero.* Also let H be the union of the pure unilateral payoff-tie hyperplanes

$$g_i(a_i, c) = g_i(b_i, c), \qquad a_i \neq b_i, \quad c \in A_{-i}.$$

Regard each Z_i as a subset of the full payoff space by allowing the other players' payoffs to vary freely, and put

$$Z = H \cup \bigcup_i Z_i. \tag{31}$$

This is an explicit finite union of proper algebraic zero sets; in particular it is Lebesgue null, by Lemma 4 below.

The role of Z_i is described in Appendix A.2: excluding it is a *sufficient* condition for the arrangement of payoff-comparison hyperplanes inside $\Delta(A_{-i})$ to be combinatorially stable under small perturbations of G_i, which is what makes the belief transformation available. Only sufficiency is used, and only sufficiency is true: a payoff matrix in Z_i may have a vanishing determinant whose comparison hyperplane misses the simplex altogether, in which case the cell structure is stable anyway. The role of H is confined to the one-player case. In fact H is redundant: selecting one payoff-difference row and $k_i - 1$ coordinate rows produces the determinant $\pm(G_{i,a} - G_{i,b})_c$, which is not identically zero, so $g \notin \bigcup_i Z_i$ already forces $g_i(a_i, c) \neq g_i(b_i, c)$. We keep H in (31) because it makes the treatment of the one-player case self-contained.

The three facts collected here are standard. They are proved because the proof of Lemma 7 uses them in the specific forms stated, and because the paper is otherwise self-contained.

Lemma 4. *The zero set of a polynomial on $\mathbb{R}^n$ that is not identically zero is Lebesgue null. Hence so is any finite union of such zero sets.*

Proof. Induct on n. For $n = 1$ a nonzero polynomial has finitely many roots. For $n > 1$, write the polynomial as a polynomial in the last variable with coefficients that are polynomials in the first $n - 1$ variables. Since it is not identically zero, at least one coefficient is not identically zero, so by induction the set $\mathcal{B}$ of points of $\mathbb{R}^{n-1}$ at which all coefficients vanish is null. For each point outside $\mathcal{B}$ the corresponding univariate polynomial is nonzero and has finitely many roots, hence a null section. By Fubini's theorem applied on each bounded box, the zero set is null. A nonzero constant has empty zero set. A finite union of null sets is null. ■

Lemma 5. *A bounded polyhedron is the convex hull of its finitely many vertices.*

Proof. Let Q be a bounded polyhedron and $x \in Q$. If x is not a vertex, there is a direction $u \neq 0$ along which Q contains a segment through x; since Q is bounded and closed, the maximal such segment has two endpoints, each lying in a face of Q of strictly smaller dimension, and x is a convex combination of them. Induction on the dimension of the smallest face containing the point terminates at vertices, of which there are finitely many since each is determined by an active set of constraints of full rank. ■

Lemma 6. *Let $\mathcal{K}$ be a finite polytopal complex whose union is a compact convex set, and for each nonempty face D let b_D be a point of the relative interior of D. The collection of simplexes with vertex sets $\{b_{D_0}, \dots, b_{D_\ell}\}$, over strictly increasing chains $D_0 \subsetneq \cdots \subsetneq D_\ell$ of faces of $\mathcal{K}$, is a triangulation of that union refining $\mathcal{K}$.*

Proof. Induct on dimension. The points $b_{D_0}, \dots, b_{D_\ell}$ are affinely independent, because b_{D_r} lies in the relative interior of D_r and the affine hulls of the D_r are strictly increasing, so each b_{D_r} lies outside the affine hull of its predecessors. For a face D of dimension p, assume the construction triangulates the boundary of D; coning each boundary simplex to b_D triangulates D, since every point of D other than b_D lies on a unique ray from b_D meeting the boundary exactly once. Two cones intersect in the cone over the intersection of their bases, which is a common face by the inductive hypothesis, and the construction on a face of D is the one induced from $\mathcal{K}$, so the pieces agree on common faces. Since every simplex of the subdivision is contained in the maximal element D_ℓ of its chain, the triangulation refines $\mathcal{K}$. ■

A.2 Direct Belief Correction

What we have. A player with payoff matrix G and belief q over the opponents' actions ranks its own actions by the vector Gq.

What we want. We want a designated type to rank its actions by a perturbed vector $Gq + e$, while keeping its payoff table G, because it has to be a normal type in the sense of Definition 2. The only object we may change is the belief.

Why the following construction solves it. We construct a probability distribution $R_{\kappa,\lambda}(q, e)$ for the opponent's actions on extra states. Here κ bounds the payoff error and λ is the mass of those states relative to the original ones, fixed for the whole error ball. The resulting belief is $(q + \lambda R_{\kappa,\lambda}(q, e))/(1 + \lambda)$, and under the unchanged payoff table G it has exactly the best responses prescribed by $Gq + e$. The main work is to make the correction a genuine probability distribution that depends continuously on q and e.

The base matrix G has arbitrary finite dimensions, since the construction is applied to each player separately and, in Section 6, with columns indexed by opponents' profiles. The perturbation always enters as a payoff vector e, so only the family $H_e = G + e\mathbf{1}^\top$ is needed: on the simplex, $H_e q = Gq + e$ because $\mathbf{1}^\top q = 1$. We therefore state the lemma for this family alone, with $\|H_e - G\|_\infty = \|e\|_\infty$.

Let Δ_k denote the probability simplex of dimension $k-1$:

$$\Delta_k = \{q \in \mathbb{R}^k : q_c \geq 0 \ (1 \leq c \leq k),\ \textstyle\sum_c q_c = 1\}.$$

The norm of a matrix in this section is the maximum absolute value of its entries. If $G \in \mathbb{R}^{m\times k}$, its row indexed by a is G_a, and

$$\mathrm{BR}_G(q) = \operatorname*{argmax}_{1\leq a\leq m} G_a q.$$

A.2.1 Example

Example 1. Let $k = 2$, so that Δ_2 is the segment from e_1 to e_2 and may be parametrized by $q_1 \in [0,1]$, with $q = (q_1, 1-q_1)$. Take $m = 3$ and

$$G = \begin{pmatrix} 4 & 0 \\ 0 & 3 \\ 2 & 1 \end{pmatrix}, \qquad \text{so that} \qquad G_1 q = 4q_1, \quad G_2 q = 3 - 3q_1, \quad G_3 q = 1 + q_1.$$

Each row of G gives an affine function of q_1: a straight line over the interval. The whole of the lemma is a statement about the configuration of these lines.

A payoff-difference row $G_a - G_b$ vanishes at the single point

$$q_1 = \frac{(G_b - G_a)_2}{(G_a - G_b)_1 - (G_a - G_b)_2},$$

if this number lies in $[0,1]$. We call it the *switch point* s_{ab} of the pair, since it is where the ranking of a against b reverses. For the matrix above, $G_1 - G_3 = (2,-1)$ gives $s_{13} = 1/(2-(-1)) = 1/3$, and the same computation gives $s_{12} = 3/7$ and $s_{23} = 1/2$. Note that action 3 is never a best response here and yet s_{13} and s_{23} still matter below. The object that has to be stable under perturbation is the full list of signs, $\mathrm{sign}((G_a - G_b)q)$, not only the identity of the maximizer.

The *cells* of the configuration are the maximal sets on which that list of signs is constant:

$$\{0\}, \quad (0, \tfrac{1}{3}), \quad \{\tfrac{1}{3}\}, \quad (\tfrac{1}{3}, \tfrac{3}{7}), \quad \{\tfrac{3}{7}\}, \quad (\tfrac{3}{7}, \tfrac{1}{2}), \quad \{\tfrac{1}{2}\}, \quad (\tfrac{1}{2}, 1), \quad \{1\}.$$

The two endpoints $\{0\}$ and $\{1\}$ are the coordinate faces of Δ_2.

We consider the perturbation in the form in which it will be used: $H_e = G + e\mathbf{1}^\top$ for a vector $e \in \mathbb{R}^3$. For a belief q, we have $H_e q = Gq + e$ since $\mathbf{1}^\top q = 1$. Adding e_a to every entry of row a shifts the ath line vertically by e_a and does nothing else. Take $e = (3/10, 0, 0)$, so that only the

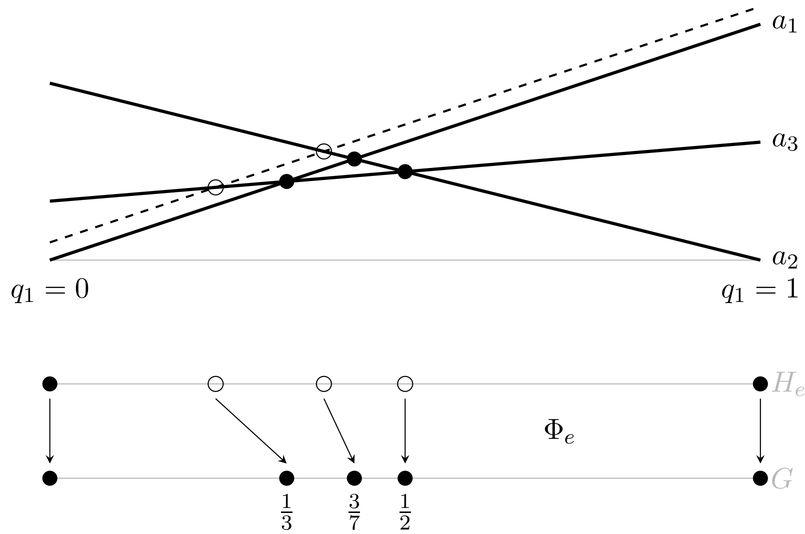

Figure 1: The case $k = 2$ of Example 1. Above: the three payoff lines $q_1 \mapsto G_a q$ over Δ_2, with the line of action 1 shifted by $e_1 = 3/10$ shown dashed. Filled dots are the switch points of G, hollow dots those of H_e. Below: the same points on two copies of the interval, and the piecewise affine map Φ_e that carries each switch point of H_e to the corresponding switch point of G and fixes both endpoints.

line of action 1 moves, upwards. The switch points become

$$s_{12}^{H_e} = \frac{3 - e_1}{7} = \frac{27}{70}, \qquad s_{13}^{H_e} = \frac{1 - e_1}{3} = \frac{7}{30}, \qquad s_{23}^{H_e} = \frac{1}{2}.$$

That is, $s_{12}^{H_e}$ and $s_{13}^{H_e}$ move left, while $s_{23}^{H_e}$ does not move at all. They move continuously in e, and for $\|e\|_\infty$ small enough they remain distinct, remain in the open interval, and keep their original order. Figure 1 shows the configuration.

In this example, write $\Phi_e(q)$ for the desired corrected belief:

$$\Phi_e(q)_1 = \begin{cases} \frac{10}{7} q_1 & \text{if } 0 \le q_1 < \frac{7}{30} \\ \frac{5}{8} q_1 + \frac{3}{16} & \text{if } \frac{7}{30} \le q_1 < \frac{1}{2} \\ q_1 & \text{if } \frac{1}{2} \le q_1 \le 1, \end{cases}$$

$$\Phi_e(q)_2 = 1 - \Phi_e(q)_1.$$

Here, $\Phi_e(\cdot)_1$ is an increasing piecewise affine bijection of $[0, 1]$ that fixes 0 and 1, sends $s_{ab}^{H_e}$ to s_{ab}^G for each (a, b), and interpolates affinely in between.

Why all comparison signs are preserved. Fix a pair (a, b) and let q_1 run from 0 to 1. The function $q_1 \mapsto (G_a - G_b)q$ is affine and vanishes only at s_{ab}^G, so its sign is determined by exactly two pieces of data: its value at the left endpoint, which is $(G_a - G_b)_2$, and whether q_1 has passed s_{ab}^G. The same is true for H_e. The two pieces of data agree: $(H_{e,a} - H_{e,b})_2 = (G_a - G_b)_2 + e_a - e_b$, and $(G_a - G_b)_2 \neq 0$ by genericity, so $|e_a - e_b| \leq 2\|e\|_\infty$ ensures the same endpoint sign for $\|e\|_\infty$ small; and Φ_e is an increasing bijection carrying $s_{ab}^{H_e}$ to s_{ab}^{G}, so q_1 lies to the right of $s_{ab}^{H_e}$ exactly when $\Phi_e(q)_1$ lies to the right of s_{ab}^{G}. Hence the two signs agree at every q, ties included:

$$\operatorname{sign}\big((G_a - G_b)\Phi_e(q)\big) = \operatorname{sign}\big((Gq+e)_a - (Gq+e)_b\big). \tag{32}$$

In particular, $\mathrm{BR}_G(\Phi_e(q)) = \mathrm{BR}_{H_e}(q)$. The correction distribution below will produce exactly this belief.

It is worth seeing why an absolute bound would not do. In the middle of the interval the displacement is about 0.056 at $q_1 = 0.35$. A bound of this size does not ensure that a coordinate close to zero remains large enough to be produced by adding states. For a proposed corrected belief q' and $\lambda > 0$, the required condition is

$$\frac{(1+\lambda)q'_c - q_c}{\lambda} \geq 0 \quad \iff \quad q'_c \geq \frac{q_c}{1+\lambda}. \tag{33}$$

Thus a coordinate-dependent bound matters. We will enforce this condition at the vertices and obtain a probability distribution everywhere else by convex interpolation.

Constructing the distribution on the extra states. The calculation can be seen explicitly in this example. Choose $\lambda = 6/7$ for illustration. List the perturbed endpoints and switch points as s_j, and their original counterparts as t_j. At $q = (s_j, 1-s_j)$ assign the correction distribution $(p_j, 1-p_j)$, where $p_j = ((1+\lambda)t_j - s_j)/\lambda$:

j	0	1	2	3	4
s_j	0	7/30	27/70	1/2	1
t_j	0	1/3	3/7	1/2	1
p_j	0	9/20	67/140	1/2	1

Every p_j is between zero and one, so these are genuine probability distributions. At either endpoint the correction is the same pure belief. For a point between two consecutive s_j, write $q_1 = (1-\alpha)s_j + \alpha s_{j+1}$, with $0 \leq \alpha \leq 1$, and assign the same convex combination of the two correction distributions. The resulting $R(q,e)$ is still a probability distribution. Its mixture with q has first coordinate $(1-\alpha)t_j + \alpha t_{j+1}$, which is exactly $\Phi_e(q)_1$. Adjacent intervals assign the same value at their shared endpoint, so this interpolation is continuous in q.

For example, at $q = (0.35, 0.65)$ this gives

$$R(q,e) = (151/320, 169/320), \qquad \frac{q + \lambda R(q,e)}{1+\lambda} = (13/32, 19/32) = \Phi_e(q).$$

The target payoffs are $Gq + e = (1.7, 1.95, 1.35)$; under the corrected belief the unchanged matrix gives $(1.625, 1.78125, 1.40625)$, with the same ranking. The chosen perturbation is large enough to make the figure legible, and the table verifies feasibility directly. The lemma below gives uniform bounds for every sufficiently small perturbation, rather than relying on these particular numbers.

What genericity excludes. A switch point at an endpoint means $(G_a - G_b)_c = 0$ for a coordinate c. An arbitrarily small vertical perturbation can move it into or out of the interval, changing the list of cells. If the switch points of different action pairs coincide in the interior, a perturbation can separate them, and their order can depend on its direction. We exclude these configurations so that each perturbed endpoint, switch point and intervening cell has a fixed counterpart for interpolation. Each degeneracy is the vanishing of a nonzero polynomial in the entries of G; the exceptional set is a finite union of such zero sets, hence Lebesgue null. These exclusions are sufficient for the construction, without claiming that every possible correction fails at an excluded matrix.

A.2.2 The Direct Belief-Correction Lemma

The example suggests the output we need: a continuous probability distribution for the extra states, together with one mass bound that works for all beliefs and errors. The next lemma states this directly. Its constants use exactly the determinant exclusions specified in Appendix A.1.

Lemma 7 (Direct Belief Correction)**.** *For every pair of positive integers m, k, let $\mathcal{Z}_{m,k}$ be the finite union of determinant zero sets defined in Step 1 of the proof. For every $G \notin \mathcal{Z}_{m,k}$ there are constants $r_G > 0$ and $C_G > 0$ with the following property. For every $0 < \kappa < r_G$ with $C_G\kappa \leq 1/2$ and every fixed $\lambda \geq 2C_G\kappa$, there is a jointly continuous map*

$$R_{\kappa,\lambda} : \Delta_k \times \{e \in \mathbb{R}^m : \|e\|_\infty \leq \kappa\} \longrightarrow \Delta_k$$

such that, for all (q, e) in this domain,

$$\mathrm{BR}_G\left(\frac{q + \lambda R_{\kappa,\lambda}(q,e)}{1+\lambda}\right) = \operatorname*{argmax}_a((Gq)_a + e_a). \tag{34}$$

The equality includes all ties. The map can be chosen with $R_{\kappa,\lambda}(q, 0) = q$ and $\mathrm{supp} R_{\kappa,\lambda}(q, e) \subseteq \mathrm{supp} q$. The set $\mathcal{Z}_{m,k}$ is Lebesgue null.

The parameters κ and λ are fixed before the input (q, e) varies. This is what allows one prescribed mass of extra states to handle every strategy profile. The option to take λ larger than $2C_G\kappa$ lets us use a common constant $C = \max_i C_{G_i}$ and $\lambda = 2C\kappa$ for all players. If $q = q(x)$ and $e = e(x)$ are continuous and obey the uniform error bound, then $R_{\kappa,\lambda}(q(x), e(x))$ is continuous as well, even when x includes the recipient's own mixed strategy.

For $k = 2$, the table in Example 1 illustrates the construction at vertices, convex interpolation makes the output a probability distribution, and the transported switch points explain the equality of best responses. In higher dimensions we must supply a subdivision whose pieces fit together on common faces, and verify continuity when the subdivision moves with e. The proof does this with barycentric subdivision and continuous barycentric coordinates. It constructs the correction distributions directly, without requiring a global inverse belief map.

Proof of Lemma 7. We first identify the vertices and cells that persist as e varies. We then put a correction distribution at each vertex of a triangulation and interpolate those distributions. This order separates the geometric question of matching cells from the probability question of keeping every coordinate nonnegative.

Step 1: the exceptional set, and what it fixes. Let $\mathcal{M}(H)$ be the finite labelled list of row vectors consisting of the coordinate rows $e_1, \ldots, e_k$ and the rows $H_a - H_b$ for $1 \le a < b \le m$. Adjoin the row $\mathbf{1} = (1, \ldots, 1)$ to obtain $\mathcal{N}(H)$. For each choice of k distinct row labels from $\mathcal{N}(H)$, take the determinant of those rows in a fixed order; each is a polynomial in the entries of H. Retain precisely those determinant polynomials which are not identically zero, and let $\mathcal{Z}_{m,k}$ be the union of their zero sets. Identically zero determinants — for instance those arising from the linear relation among three payoff-difference rows $H_a - H_b$, $H_b - H_c$, $H_a - H_c$ — impose no exclusion, which is why the definition must omit them: otherwise $\mathcal{Z}_{m,k}$ would be everything.

There are finitely many retained polynomials and each is not identically zero, so their zero sets are Lebesgue null by Lemma 4, and so is their finite union. Fix $G \notin \mathcal{Z}_{m,k}$. We retain this entire exceptional set and restrict only the perturbations to $H_e = G + e\mathbf{1}^\top$.

Step 2: affine candidate vertices and locally constant signs. Write $r(H_e) = r(G) + \delta_r(e)\mathbf{1}$, where $\delta_r(e) = 0$ for a coordinate row and $\delta_r(e) = e_a - e_b$ for the row labelled (a, b). For a list B of $k-1$ distinct row labels, put $A_B = [B(G); \mathbf{1}]$ and $d_B = \det A_B$. Subtracting $\delta_r(e)\mathbf{1}$ from each row of $B(H_e)$ shows that $\det[B(H_e); \mathbf{1}] = d_B$ for every e. Thus a singular list never becomes nonsingular. For $d_B \neq 0$ its candidate is

$$v_B(e) = A_B^{-1} \begin{pmatrix} -\delta_B(e) \\ 1 \end{pmatrix}, \qquad \delta_B(e) = (\delta_r(e))_{r \in B}.$$

To see why this is the required candidate, use $\mathbf{1}v = 1$ to rewrite $B(H_e)v = 0$ as $B(G)v = -\delta_B(e)$. Thus the coefficient matrix is fixed, only the right-hand side moves, and $v_B(e)$ is affine in e. The empty list is allowed when $k = 1$. For every row label r, its evaluation is also affine:

$$s_{rB}(e) := r(H_e)v_B(e) = r(G)v_B(e) + \delta_r(e). \tag{35}$$

If $s_{rB}(0) = 0$, then $\det[B(G); r(G)] = 0$. This determinant polynomial is identically zero by the choice of G (or because a label repeats). Hence $\det[B(H_e); r(H_e)] = 0$ for every e; since the rows of $B(H_e)$ are independent, $r(H_e)$ lies in their span and $s_{rB}(e) = 0$ for every e. Otherwise $s_{rB}(0) \neq 0$, and continuity preserves its nonzero sign on a small ball. Finitely many evaluations are involved, so choose one radius $r_G > 0$ on which all their signs are fixed.

These signs determine feasibility through the coordinate rows, and determine coordinate support as well. Two candidates coincide exactly when one satisfies all the other's defining equations, so their coincidence relations are fixed too. We may therefore label the distinct candidates in the simplex by one finite index set $\mathcal{V}$, writing their affine positions as $v(e)$, $v \in \mathcal{V}$.

Step 3: cells, and the local constancy of the face structure. For an assignment ς of signs to the rows of $\mathcal{M}$, let $C_\varsigma(e)$ be the set of points of Δ_k at which each row takes exactly the prescribed sign; coordinate rows are allowed only the signs zero and positive. The sets $C_\varsigma(e)$ partition Δ_k, since every point has a well-defined sign vector. Let $P_\varsigma(e)$ be the polytope obtained by replacing all required strict inequalities by weak inequalities, retaining the required equalities and $\mathbf{1}q = 1$.

Every vertex of a nonempty $P_\varsigma(e)$ is one of the candidates of Step 2. Indeed, let q be a vertex and consider the homogeneous rows active at q, that is, those with $r(H_e)q = 0$. Their rank is at most $k-1$, because they all annihilate the nonzero vector q. If their rank were less than $k-1$, there would be a direction $u \neq 0$ annihilated by all of them and by $\mathbf{1}$; moving from q a small distance along $\pm u$ would preserve the active equalities and the inactive strict inequalities, contradicting extremality. So the rank is exactly $k-1$, and selecting $k-1$ independent active rows produces a list B with $d_B \neq 0$ and $v_B(e) = q$. Conversely, every candidate satisfying the weak constraints of $P_\varsigma(e)$ is a vertex of that polytope, since its $k-1$ independent active equations together with the normalization determine it uniquely. Thus the vertex set of $P_\varsigma(e)$ is read off from the fixed sign data of (35).

A bounded polytope is the convex hull of its vertices (Lemma 5). Consequently $C_\varsigma(e)$ is nonempty if and only if $P_\varsigma(e)$ has a vertex and each required strict inequality is strict at at least one vertex: necessity holds because a linear functional that is nonzero somewhere on a polytope is nonzero at some vertex, and sufficiency because averaging finitely many witnessing vertices makes all the required inequalities strict simultaneously. Whenever $C_\varsigma(e)$ is nonempty its closure equals $P_\varsigma(e)$: mix any point of $P_\varsigma(e)$ with a point of $C_\varsigma(e)$, putting positive weight on the latter, and let that weight decrease to zero.

It follows that the same collection of sign patterns has nonempty cells, with the same vertices and the same vertex–face incidences, throughout the neighborhood. These cell closures form a finite polytopal complex. To see the face property directly, intersect two of them: on the intersection, every row whose prescribed signs differ must vanish, so the intersection is obtained from each by turning supporting inequalities into equalities, hence is a face of each and is itself a cell closure. Every coordinate face of Δ_k is a union of cells, since the coordinate rows are among $\mathcal{M}$. Cell dimensions are fixed as well. Because all remaining inequalities are strict, the directions in the affine hull of a nonempty cell are exactly those satisfying $\mathbf{1}u = 0$ and $r(H_e)u = 0$ for its prescribed zero rows. Since $r(H_e)u = r(G)u$ when $\mathbf{1}u = 0$, this direction space is independent of e.

Step 4: probability distributions at the vertices of a triangulation. For each nonempty face D of the complex, let $b_D(e)$ be the average of its vertices. It lies in the relative interior of D. By Lemma 6, the simplexes with vertices $b_{D_0}(e), \dots, b_{D_\ell}(e)$ for chains $D_0 \subsetneq \cdots \subsetneq D_\ell$ triangulate Δ_k. The same chains work for every e, because Step 3 fixed the face incidences.

To mix $b_D(e)$ back to $b_D(0)$ by adding a nonnegative distribution, we need $b_D(e)_c \leq (1 + \lambda)b_D(0)_c$. We therefore first bound coordinate increases. Let $E_B e = (\delta_B(e), 0)^\top$. The affine formula in Step 2 gives a common coordinatewise Lipschitz constant

$$L_\mathcal{V} = \max_{B: d_B \neq 0} \|A_B^{-1} E_B\|_{\infty \to \infty}, \qquad \varrho = \min_{\substack{v \in \mathcal{V},\ 1 \leq c \leq k \\ v(0)_c > 0}} v(0)_c > 0, \qquad \theta = \max\{1, L_\mathcal{V}/\varrho\},$$

where $\|M\|_{\infty\to\infty} = \max_i \sum_j |M_{ij}|$. If $v(0)_c = 0$, Step 2 gives $v(e)_c = 0$; otherwise $v(0)_c \geq \varrho$ and $|v(e)_c - v(0)_c| \leq L_{\mathcal{V}}\|e\|_\infty \leq \theta\|e\|_\infty v(0)_c$. Averaging this estimate over the vertices of D yields the one-sided bound

$$b_D(e)_c \leq (1+\theta\kappa) b_D(0)_c \qquad (\|e\|_\infty \leq \kappa). \tag{36}$$

Take $C_G = 2\theta$, reduce r_G if necessary so that $\theta r_G \leq 1/2$, and fix κ, λ as in the statement. At the vertex $b_D(e)$ prescribe

$$R_D(e) = \frac{(1+\lambda) b_D(0) - b_D(e)}{\lambda}.$$

Its coordinates sum to one, and (36) gives $R_D(e)_c \geq (\lambda - \theta\kappa) b_D(0)_c/\lambda \geq 0$, since $\lambda \geq 2C_G\kappa = 4\theta\kappa$. Thus every prescribed value is already a probability distribution. Mixing it with $b_D(e)$ in the proportions $1 : \lambda$ returns exactly $b_D(0)$.

Step 5: interpolation and joint continuity. If q belongs to the simplex for a chain $D_0 \subsetneq \cdots \subsetneq D_\ell$, write it in barycentric coordinates and define

$$q = \sum_{j=0}^{\ell} t_j b_{D_j}(e), \qquad R_{\kappa,\lambda}(q,e) = \sum_{j=0}^{\ell} t_j R_{D_j}(e), \qquad t_j \geq 0, \quad \sum_j t_j = 1.$$

The second expression is a convex combination of probability distributions, so it belongs to Δ_k. On the intersection of two simplexes, only vertices of their common face have nonzero weights; both formulas therefore give the same value. This makes $R_{\kappa,\lambda}$ well defined.

Continuity must hold as e changes, not just within one fixed simplex. For each maximal simplex, put its k vertices in the columns of a matrix $A_\tau(e)$. They are affinely independent and have coordinate sum one, so this matrix is invertible: a linear dependence would have coefficients summing to zero and would contradict affine independence. Its barycentric coordinates are $t = A_\tau(e)^{-1} q$, hence continuous in (e, q). The set on which these weights are nonnegative is closed in the parameter domain, and these finitely many sets cover it. The interpolation formula is continuous on each such set and agrees on overlaps. The finite pasting property for continuous functions on closed sets gives joint continuity on the whole domain.

Step 6: the corrected best responses. By the vertex prescription and the interpolation formula,

$$\widehat{q} := \frac{q + \lambda R_{\kappa,\lambda}(q,e)}{1+\lambda} = \sum_j t_j b_{D_j}(0).$$

Let D_* be the largest face in the chain whose weight is positive. All terms with positive weight lie in D_*, and the term $b_{D_*}(e)$ lies in its relative interior with positive weight. Every inequality defining D_* that is not an equality is strict at that term, so it remains strict in their convex combination. Consequently q lies in the relative interior of $D_*(e)$, and the same argument places $\widehat{q}$ in the relative interior of $D_*(0)$. Step 3 identifies these interiors by the same sign pattern. In particular, $(G_a - G_b)\widehat{q}$

and $(Gq+e)_a - (Gq+e)_b$ have the same sign for every pair. An action is a best response exactly when all its comparisons are nonnegative, which proves (34), including ties.

At $e = 0$ each $R_D(0) = b_D(0)$, so interpolation gives $R_{\kappa,\lambda}(q, 0) = q$. If $q_c = 0$, nonnegativity of its barycentric representation forces $b_{D_j}(e)_c = 0$ whenever $t_j > 0$. The fixed coordinate supports from Step 2 then give $b_{D_j}(0)_c = 0$ and $R_{D_j}(e)_c = 0$, proving the support inclusion. Throughout the construction $\kappa > 0$ and $\lambda > 0$ were fixed; no continuity as λ tends to zero is required. ■

A.3 Implementation by Additional Types

In this subsection, we return to the two-player setting of Section 2.2. Fix an original finite type space $T = T_1 \times T_2$ and a common prior P with positive type marginals, and let Σ be the corresponding strategy space. When player i is fixed, write $j = 3 - i$.

What we have. A continuous map $r : \Sigma \to \Delta(A_j)$ that we would like a designated type t_i to face as its conditional distribution over the opponent's action on a designated block of states.

What we still need. The value $r(\sigma)$ depends on the equilibrium strategies, which the designer of the elaboration does not know. The elaboration must therefore contain machinery that produces $r(\sigma)$, to any prescribed accuracy, at every equilibrium, whatever σ turns out to be, and that machinery must not disturb the payoffs of the original types.

Why the following construction solves it. We append extra types to the two existing players. Their payoffs are chosen so that their own best-response conditions force their behavior to reproduce a prescribed arithmetic circuit evaluated at σ, and so that the last stage of the circuit converts the computed distribution into an actual distribution of opponent actions on the designated block. All the added states other than that block can be given arbitrarily small probability, by scaling the helper payoffs accordingly; the block itself has the mass we prescribe.

Before the formal statement, three structural points.

Adding types is not adding players. A helper type of player j is an element of the type set T_j of the *existing* player j, with the same action set A_j. When we say that a helper type "uses two designated actions", we mean that its payoff table makes the other elements of A_j strictly worse. Consequently every state is still a pair (a type of player 1, a type of player 2), the base payoff $g_i(a_1, a_2)$ is meaningful at every state, and the constructed game is a two-player Bayesian game. This is what allows an original type to be given its exact base payoff table at states where it meets a helper.

Unnormalized masses. We assign the new states unnormalized masses while leaving the original states with masses $P(t_1, t_2)$, and divide everything by the final total at the end. Dividing by a positive constant changes no incentive. Moreover, to check a type's incentives it suffices to compare its *unnormalized expected payoff numerators*, that is, the sums over its states of mass times payoff: the conditional denominator is positive and identical across that type's actions.

Small masses and large payoffs. A state of mass ε' whose payoff coefficients have been divided by ε' contributes to a helper's payoff numerator exactly what a state of mass 1 with the undivided coefficients would contribute, since $\varepsilon' \cdot (u/\varepsilon') = u$. This is why the added states can be made

Role	Owner	Appears in states with	Read by	Engineered payoff difference
controller R	i	input types (read states); its own output C (internal state)	nobody	$\ell(z) - z_{\mathsf{C}}$
output C	j	its controller; later readers	later controllers and flags	$y_{\mathsf{R}} - \frac{1}{2}$
probe B_n	j	one original type s_i of player i	controllers	$\sigma_i(a \mid s_i) - u_n$
flag	i	the output it reads; final outputs	final output types	$\widehat{F}_l - u_m$
final output	j	the flags; the recipient t_i (output state)	nobody	the score (40)

Table 2: Helper roles in Lemma 8. “Read by” means: appears in a state that gives another helper an incentive depending on this type’s action.

negligible in probability while retaining their incentive effect, and it is also why the helper payoffs are unbounded across the sequence of constructed elaborations. Definition 1 permits this, since each individual elaboration is finite and has bounded payoffs.

Table 2 summarizes the roles.

Lemma 8 (Approximate Actuators with Arbitrary Inputs). *Fix $i \in \{1, 2\}$, put $j = 3 - i$, and fix an original type t_i, a continuous map $r : \Sigma \to \Delta(A_j)$, and an output mass $\varpi > 0$. For every $\alpha, \tau > 0$ one can append finitely many helper types and states with these properties.*

(i) Each output state contains t_i and a fresh player-j helper type. The output states have total unnormalized mass exactly ϖ. At every Bayesian equilibrium, if σ denotes the strategies at the original types, the opponent's action distribution conditional on these output states is within total variation α of $r(\sigma)$.

(ii) All other new states have total mass less than τ. Each contains at most one original type, and that original type can be assigned any prescribed bounded payoff table at the state, including its exact base-game payoff.

(iii) All helper payoffs are bounded within the constructed finite game. The bounds need not be uniform in α or τ.

These conclusions hold for every behavioral Bayesian equilibrium, including all possible mixing at helper indifferences. An actuator *means the collection of helper types and states producing the output distribution in part (i).*

Proof. *Step 1: a basic gate.* A binary helper type uses two designated actions of its own action set, called low and high. At its *active* states — those where its payoffs are used to create incentives

— all its other actions have payoff -1; at its *passive* states — those where it appears only to be read by somebody else — every action has payoff zero. Low has payoff zero at every state. Each helper has at least one active state of positive mass, so its unused actions are strictly worse than low and are never played. By scaling its high-action payoffs, any prescribed affine expression in the quantities it faces can be realized as its unnormalized high-minus-low payoff numerator; adding passive states changes only the positive conditional denominator, hence no incentive.

Suppose $z_1, \ldots, z_p$ are the high-action probabilities of p player-j types and $\ell(z)$ is an affine function of them. Introduce a player-i controller R and a player-j output C, and write y_{R} and z_{C} for their high-action probabilities. Read states, each pairing R with one input type, together with one internal state pairing R with C, give R and C respectively the payoff differences

$$\ell(z) - z_{\mathsf{C}}, \qquad y_{\mathsf{R}} - \frac{1}{2}. \tag{37}$$

Concretely: to give R the term $c\, z_n$, pair R with the nth input type in a read state of mass ε' and let R's high payoff there be $(c/\varepsilon')\mathbf{1}_{\{a_j=\text{high}\}}$ and its low payoff be zero; the contribution to R's numerator is then cz_n, of either sign according to the sign of c. A constant term is obtained the same way from a state whose player-j occupant is a fresh type committed to a fixed action. The term $-z_{\mathsf{C}}$ comes from the internal state. Both R and C receive zero own payoffs at every later state where they are read or used, so (37) is their complete incentive.

Every Bayesian equilibrium of this pair satisfies

$$z_{\mathsf{C}} = \operatorname{clip}(\ell(z), 0, 1), \qquad \operatorname{clip}(x, 0, 1) = \max\{0, \min\{x, 1\}\}. \tag{38}$$

Indeed, if $0 < z_{\mathsf{C}} < 1$ then C is indifferent, so $y_{\mathsf{R}} = 1/2$, so R is indifferent, so $\ell(z) = z_{\mathsf{C}} \in (0,1)$ and both sides of (38) agree. If $z_{\mathsf{C}} = 0$, optimality of low for C requires $y_{\mathsf{R}} \le 1/2$; were $\ell(z) > 0 = z_{\mathsf{C}}$, then R would strictly prefer high and $y_{\mathsf{R}} = 1$, a contradiction, so $\ell(z) \le 0$ and the clip is 0. If $z_{\mathsf{C}} = 1$, the symmetric argument gives $\ell(z) \ge 1$. The boundary values $\ell(z) \in \{0, 1\}$ are covered by these three cases. Note that (38) is an exact identity for the gate's *output* z_{C}, holding at *every* equilibrium and whatever tie-breaking occurs elsewhere. It does not pin down the controller: when $\ell(z) = z_{\mathsf{C}}$ the type R is indifferent and may mix, and nothing below uses its probability.

Every controller belongs to player i and every arithmetic output to player j. Since a controller reads only player-j quantities and produces a player-j quantity, gates compose into a feedforward circuit: order the outputs so that each controller reads only earlier ones, and (38) determines every output value by induction along that order, as an exact function of the circuit's actual inputs. Those inputs that are player-j original probabilities are the original strategies themselves. The ones supplied by the probes of Step 2 below are determined by the original strategies only up to a tie in a single probe, so the circuit output is not in general a function of σ alone; what will be shown is that it approximates its target uniformly over every permissible resolution of such ties.

Step 2: reading both players' original strategies. A player-i controller can read an original player-j type directly, since that type's action probability is a player-j quantity. To store the coordinate

$x = \sigma_j(a \mid t_j)$ as a gate output, use the gate with $\ell = 1/4 + x/2$, whose value lies in $[1/4, 3/4]$ and is therefore never clipped; its player-j output stores $z = 1/4 + x/2$, from which $x = 2z - 1/2$ is recovered by an affine expression in any later gate.

Reading a coordinate of player i requires an extra device, because a player-i controller's incentives cannot depend on player i's own behavior. Fix an original type s_i and put $x = \sigma_i(a \mid s_i)$. Introduce player-j binary probe types $B_1, \ldots, B_Q$ and pair each with s_i in a probe state. Give B_n the payoff difference

$$x - u_n, \qquad u_n = (n - \frac{1}{2})/Q,$$

by letting its high-action payoff at that state be proportional to $\mathbf{1}_{\{a_i = a\}} - u_n$; all later states give B_n zero own payoffs. Its equilibrium high probability z_n is one when $x > u_n$, zero when $x < u_n$, and arbitrary only in the tie case $x = u_n$. Hence, whatever happens at ties,

$$\left| Q^{-1} \sum_{n=1}^{Q} z_n - x \right| \leq Q^{-1}, \tag{39}$$

because $Q^{-1}\#\{n : u_n < x\}$ differs from x by at most Q^{-1} and at most one index can be tied. A controller reads the average $Q^{-1}\sum_n z_n$ through Q separate read states and can store it with another gate. The device applies verbatim when $s_i = t_i$, so the recipient's *own* strategy is among the readable inputs. Every original strategy coordinate is thus available to the circuits, exactly for player j and within Q^{-1} for player i. The probe states may be given arbitrarily small total mass, since the payoff coefficient at a probe state of mass ε' is scaled by $1/\varepsilon'$.

Step 3: approximating continuous functions. We claim that for every continuous $F : \Sigma \to [0, 1]$ and every $\gamma > 0$ there is a finite circuit of the above gates whose output differs from $F(\sigma)$ by at most γ at every equilibrium, uniformly in σ.

Extend F continuously to the ambient cube by composing with the Euclidean projection onto Σ. For a continuous F on a p-dimensional cube, the Bernstein polynomial of degree Q' in each coordinate is

$$B_{Q'}F(x) = \mathbb{E}\, F(U_1/Q', \ldots, U_p/Q'), \qquad U_n \sim \mathrm{Bin}(Q', x_n) \text{ independent.}$$

If ω_F is a modulus of continuity for the maximum norm, then splitting on the event that some $|U_n/Q' - x_n|$ exceeds ε', and using the union bound with $\mathrm{Var}(U_n/Q') \leq 1/(4Q')$, bounds the uniform approximation error by

$$\omega_F(\varepsilon') + 2\|F\|_\infty \frac{p}{4Q'\varepsilon'^2}.$$

Choosing ε' small and then Q' large makes this less than $\gamma/3$; so polynomials approximate F uniformly. A polynomial is built from its variables by additions, scalar multiplications and products, and a product is reduced to squares by $xy = \big((x+y)^2 - (x-y)^2\big)/4$. A square, or any continuous function of one variable on a bounded interval, is uniformly approximated by piecewise linear functions, and on a bounded interval a piecewise linear function is an affine function plus a finite linear combination of hinges $(x-c)_+$. Each hinge is a rescaled clipped affine expression, hence one gate: for x in a bounded range and a suitable $R \geq 1$, $(x-c)_+ = R\,\mathrm{clip}((x-c)/R, 0, 1)$. Every intermediate quantity is encoded affinely in $[0,1]$ before being stored in a gate output and decoded affinely when read. Composing these finitely many gates gives a circuit computing, exactly, a fixed piecewise linear function $\widehat{F}$ of its inputs with $\sup_\sigma |\widehat{F}(\sigma) - F(\sigma)| < 2\gamma/3$.

The order of the choices matters. First fix the circuit, which fixes a Lipschitz constant L_{cir} for $\widehat{F}$ as a function of its inputs; only then choose the probe resolution Q so large that $L_{\mathrm{cir}}/Q < \gamma/3$, which by (39) bounds the effect of the input error. If the target takes values in $[0,1]$, a final clipping gate keeps the output there without increasing the error. This proves the claim.

Step 4: output labels. Enumerate the opponent's actions as $A_j = \{c_1, \ldots, c_{\bar{k}}\}$ with $\bar{k} = |A_j|$. For $1 \leq l < \bar{k}$, use Step 3 to compute the cumulative probabilities

$$F_l(\sigma) = \sum_{n=1}^{l} r(\sigma)(c_n)$$

by player-j outputs $\widehat{F}_l$, uniformly within γ. Choose Θ output labels, each carrying state mass ϖ/Θ, and put $u_m = (m - 1/2)/\Theta$ for $1 \leq m \leq \Theta$. For each pair (m, l), a fresh player-i flag reads $\widehat{F}_l$ and has high-minus-low payoff difference $\widehat{F}_l - u_m$; write D_{ml} for its high probability, which is 1 if $\widehat{F}_l > u_m$, is 0 if $\widehat{F}_l < u_m$, and is arbitrary in the tie case. The flag receives zero own payoffs in the states where it is subsequently read.

For each label m, introduce a final player-j output type whose score for the action a_j is

$$\mathbf{1}_{\{a_j = c_{\bar{k}}\}} + \sum_{l=1}^{\bar{k}-1} D_{ml}\big(\mathbf{1}_{\{a_j = c_l\}} - \mathbf{1}_{\{a_j = c_{l+1}\}}\big). \tag{40}$$

Read states pairing this type with the player-i flags implement (40) as its unnormalized payoff numerator, exactly as in Step 1, the constant term coming from a state with a committed player-i type. The final output type has zero own payoffs in its output state, which pairs it with the original recipient t_i and has mass ϖ/Θ; its incentives therefore come entirely from the read states.

Call the label m *clean* if $|u_m - F_l(\sigma)| > \gamma$ for every l. At a clean label, every flag is pure and correct: $D_{ml} = \mathbf{1}_{\{F_l > u_m\}}$, since $|\widehat{F}_l - F_l| \leq \gamma$. Set $F_0 = 0$ and $F_{\bar{k}} = 1$ and let n be the unique index with $F_{n-1} < u_m < F_n$. Then $D_{ml} = \mathbf{1}_{\{l \geq n\}}$, and the coefficient of $\mathbf{1}_{\{a_j = c_p\}}$ in (40) is

$$\mathbf{1}_{\{p = \bar{k}\}} + D_{mp}\mathbf{1}_{\{p \leq \bar{k}-1\}} - D_{m,p-1}\mathbf{1}_{\{p \geq 2\}},$$

which equals 1 when $p = n$ and 0 otherwise, in each of the cases $p < n$, $p = n < \bar{k}$, $p = n = \bar{k}$, $n < p < \bar{k}$ and $n < p = \bar{k}$. So at a clean label the final output type has the unique maximizer c_n and plays it purely. The action distribution produced by the output states is therefore the mixture, over labels, of the actions forced at each label.

It remains to bound the error. Ignoring unclean labels for a moment, the label mixture is the midpoint quantile distribution of F, whose cumulative probabilities differ from F_l by at most $1/(2\Theta)$; telescoping the successive masses shows that it differs from $r(\sigma)$ in total variation by at most $(\bar{k}-1)/\Theta$. The unclean labels are those with $|u_m - F_l| \leq \gamma$ for some l; for each l there are at most $2\gamma\Theta + 1$ of them, so their fraction is at most $2(\bar{k}-1)\gamma + 2(\bar{k}-1)/\Theta$. At an unclean label the flags may be mixed and the final output type may randomize arbitrarily, but its contribution to the mixture is bounded by its mass. Hence at every equilibrium the total output error is at most

$$2(\bar{k}-1)\gamma + 3(\bar{k}-1)/\Theta. \tag{41}$$

Choosing γ small and Θ large makes this less than α. The estimate is uniform over the original strategies and over all helper responses at ties, which is what part (i) asserts.

Step 5: assembling the prior. Give all probe, read and internal states positive masses with sum less than τ, and scale their helper payoffs accordingly to implement the expressions above. Use disjoint fresh type labels for distinct controllers, arithmetic outputs, probes, flags and final outputs. A read pair is (reader, source), an internal pair is (controller, its fresh output), and an output pair is (original recipient, final output); these pairs do not collide, so the assignment of payoffs to states is consistent. Repeated reads of the same source by the same reader are combined by adding their coefficients in one state. A constant is placed in a designated read or internal state, or in a separate state with a fresh opponent type committed to one action, that commitment being enforced by giving that action a strictly higher payoff.

Every helper type has positive marginal probability, and there are finitely many states with finite payoff coefficients, so payoffs are bounded within the constructed game. Finally, at every state containing an original type — a direct-read state, a probe state, or an output state — the construction has specified only the *helper's* payoffs; the original type's payoff table at that state is therefore free, and may be set to its exact base-game payoff, or to zero, with no effect on any helper incentive. This is part (ii), and it is the property that will make (exact) normality available in Proposition 1. ■

Two features of the proof will be used later and are worth isolating. First, every conclusion holds at *every* Bayesian equilibrium of the constructed game, with no selection among equilibria and no convention about how indifferent helpers break ties: the gate identity (38) holds at all equilibria, and the two places where ties can genuinely bite — the probes of Step 2 and the flags of Step 4 — are absorbed into the explicit error bounds (39) and (41). Second, the map r may depend on the recipient's own strategy, because Step 2 transfers player-i coordinates to player-j quantities; this is used for both classes of types in the next subsection.

A.4 Proof of Finite Feedback Realization

Proof of Proposition 1. *Stage 1: the correction distributions at good types.* Let r_{G_i} and C_i be the constants of Lemma 7 for G_i. Put $C = \max\{C_1, C_2\} > 0$ and choose $\kappa_0 > 0$ so that $\kappa_0 < \min_i r_{G_i}$

and $C\kappa_0 \le 1/2$. Set $D = 4C$.

Assume first $\kappa > 0$ and put $\lambda = 2C\kappa$, so that $\lambda \le 1$ and $\lambda \ge 2C_i\kappa$ for both players. Lemma 7 supplies the continuous correction map $R^{(i)}_{\kappa,\lambda}$ for each G_i, with this same mass λ. Fix a good type $t_i \in S_i$ and suppress it in the notation. Define

$$\begin{aligned} e_i(\sigma) &= w_i(t_i,\sigma) - G_i q_i(t_i,\sigma),\\ r_i(\sigma) &= R^{(i)}_{\kappa,\lambda}(q_i(t_i,\sigma), e_i(\sigma)),\\ \widetilde{q}_i(\sigma) &= \frac{q_i(t_i,\sigma) + \lambda r_i(\sigma)}{1+\lambda}. \end{aligned}$$

By (5), $\|e_i(\sigma)\|_\infty \le \kappa$ for every σ. Thus the lemma applies on the whole strategy space: r_i is a continuous map from Σ into $\Delta(A_j)$, and $\operatorname{argmax} G_i\widetilde{q}_i(\sigma) = \operatorname{argmax} w_i(t_i,\sigma)$, including all ties. The fact that r_i is a probability distribution is part of the lemma's conclusion, so it can be used directly as an actuator output.

Stage 2: the construction at good types. Give every good original type its exact payoff g_i at *every* state where it appears. Retain the original masses and append, for each good type t_i, an actuator in the sense of Lemma 8 with output mass $\varpi = \lambda P_i(t_i)$ and output map r_i.

Consider the ideal case in which each actuator delivers its target exactly. Type t_i then meets the opponent's actions with distribution $q_i(t_i,\sigma)$ on its original states, of total mass $P_i(t_i)$, and with distribution $r_i(\sigma)$ on its output states, of total mass $\lambda P_i(t_i)$. Its unnormalized payoff numerator over these states is therefore

$$P_i(t_i)(G_i q_i + \lambda G_i r_i) = P_i(t_i)(1+\lambda) G_i \widetilde{q}_i(\sigma), \tag{42}$$

by the definition of $\widetilde{q}_i$. By (34), its set of maximizers is exactly $\operatorname{argmax}(G_i q_i + e_i) = \operatorname{argmax} w_i(t_i,\sigma)$, ties included. This is the whole point of Appendix A.2: the type's payoff table is untouched, and the auxiliary ranking has been produced by the belief alone.

Two perturbations of (42) are present in the actual construction and both vanish in the limit. The actuator output is only within α_l of $r_i(\sigma)$ in total variation, which contributes at most $2\lambda P_i(t_i)\alpha_l \|G_i\|_\infty$ to the numerator. And the type t_i also occurs in probe states where it is the source and in direct-read states where another actuator reads one of its coordinates, of total mass at most τ_l, where it again receives its exact payoff g_i; these contribute at most $\tau_l\|G_i\|_\infty$. If $\kappa = 0$, set $\lambda = 0$, omit all good-type output blocks and set $\widetilde{q}_i(\sigma) = q_i(t_i,\sigma)$: $w_i(t_i,\sigma) = G_i q_i(t_i,\sigma)$ and the original payoff vector already has the required maximizers.

Stage 3: the construction at bad types. Fix $t_i \notin S_i$ and a reference action $a_i^0 \in A_i$. For $a \ne a_i^0$ let

$$d_a(\sigma) = w_i(t_i,\sigma)_a - w_i(t_i,\sigma)_{a_i^0},$$

which is continuous on the compact set Σ; choose $M > 1 + \max_{a,\sigma} |d_a(\sigma)|$, a finite number depending on the fixed feedback problem but on nothing constructed later. Put $q_a(\sigma) = (d_a(\sigma) + M)/(2M) \in$

$(0, 1)$.

Give this type payoff zero, for every action profile, at every state other than its own output blocks constructed below: at its original states, at probe states where it is the source, at direct-read states where another actuator reads one of its coordinates, and at any other auxiliary state in which it appears. Part (ii) of Lemma 8 permits this, since the construction there prescribes only the helpers' payoffs at those states and leaves the original type's payoff table free; the assignment changes no helper incentive, no state mass and no estimate, and it makes the contribution of all those states to this type's payoff numerator exactly zero. For each $a \neq a_i^0$, append an output block of mass $P_i(t_i)$ whose target distribution puts probability $q_a(\sigma)$ on one designated opponent action, called high, and $1 - q_a(\sigma)$ on another, called low; Lemma 8 supplies it, and applies even though q_a depends on the recipient's own strategy, by Step 2 of its proof. In that block give the original type the payoff

$$2M\big(\mathbf{1}_{\{\text{opponent high}\}} - \frac{1}{2}\big) \quad \text{if it chooses } a, \qquad 0 \quad \text{for every other own action.}$$

In the ideal case the block contributes $P_i(t_i) \cdot 2M(q_a(\sigma) - \frac{1}{2}) = P_i(t_i)d_a(\sigma)$ to the numerator of action a and nothing to any other action. Summing the $|A_i| - 1$ blocks, the type's numerator vector is $P_i(t_i)\big(w_i(t_i, \sigma) - w_i(t_i, \sigma)_{a_i^0}\mathbf{1}\big)$, whose maximizers are those of $w_i(t_i, \sigma)$. A bad type is not normal, and is not claimed to be; the cost of it is the mass of its blocks, accounted for next.

Stage 4: mass and normality accounting. There are finitely many blocks in total. For the lth constructed game, choose the actuator errors $\alpha_l \downarrow 0$ and let $\tau_l \downarrow 0$ bound the sum of all probe, read and internal masses. Using $P_i(S_i^c) \leq \varepsilon$ and $P_i(S_i) \leq 1$, the total output mass is at most

$$\lambda\big(P_1(S_1) + P_2(S_2)\big) + (m-1)\big(P_1(S_1^c) + P_2(S_2^c)\big) \leq 4C\kappa + 2(m-1)\varepsilon = D\kappa + (K-1)\varepsilon,$$

which with the internal masses proves (6).

Every original state in $S_1 \times S_2$ consists of two exactly normal types. Indeed, a good type t_i occurs only in original states, in the output states of its own actuator, in probe states where it is the source, and in direct-read states where another actuator reads one of its coordinates; by Stage 2 it receives g_i at all of them, and no other state contains it, because output, probe and read states pair one original type with fresh helpers. Since the original states carry total unnormalized mass one, after normalization by $1 + \Lambda_l$ the event that all players are normal has probability at least

$$\frac{1 - \varepsilon}{1 + \Lambda_l} \geq 1 - \varepsilon - \Lambda_l,$$

which is (7). All types used have positive marginal probability, all payoffs are bounded within each finite construction, and each E_l has a behavioral Bayesian equilibrium.

Stage 5: passage to the limit. Let σ^l be the restriction to the original types of a Bayesian equilibrium of E_l, and suppose $\sigma^l \to \sigma^*$ along a subsequence; some subsequence converges because Σ is compact. Fix a type and consider its unnormalized numerator vector $\Pi^l \in \mathbb{R}^{A_i}$ in E_l.

At a good type, Stage 2 gives

$$\Pi^l = P_i(t_i)(1+\lambda)G_i\widetilde{q}_i(\sigma^l) + O(\alpha_l) + O(\tau_l),$$

where the errors are uniform in σ^l. The maps $q_i(t_i,\cdot)$ and r_i are continuous when $\kappa > 0$, so $\widetilde{q}_i$ is continuous; it is also continuous when $\kappa = 0$, by its definition above. Hence $\Pi^l \to P_i(t_i)(1+\lambda)G_i\widetilde{q}_i(\sigma^*)$, whose maximizers are exactly $\operatorname{argmax} w_i(t_i,\sigma^*)$. At a bad type, Stage 3 gives $\Pi^l \to P_i(t_i)\big(w_i(t_i,\sigma^*) - w_i(t_i,\sigma^*)_{a_i^0}\mathbf{1}\big)$, with the same set of maximizers.

Now let $a \in \operatorname{supp}\sigma_i^*(\cdot \mid t_i)$. Then $\sigma_i^l(a \mid t_i) > 0$ for all large l in the subsequence, so the equilibrium condition in E_l gives $\Pi_a^l \geq \Pi_{a'}^l$ for every $a' \in A_i$; here we use that comparing conditional payoffs is the same as comparing unnormalized numerators, the conditional denominator being positive and common to all actions of the type. Letting $l \to \infty$ preserves the weak inequalities, so a maximizes the limit vector, hence $a \in \operatorname{argmax} w_i(t_i,\sigma^*)$. This is (4), and it holds for every behavioral equilibrium of every E_l and every resolution of helper indifferences, proving (c).

Stage 6: transfer of robustness. Write χ_l for the unnormalized action measure generated by the new states, so $\chi_l \geq 0$ and $\sum_b \chi_l(b) = \Lambda_l$. If σ^l is the original restriction of an equilibrium of E_l, the action distribution of that equilibrium is

$$\mu_{E_l} = \frac{\mu_{\sigma^l} + \chi_l}{1+\Lambda_l},$$

since the original states carry mass one and the equilibrium's play there is σ^l. Hence $\mu_{E_l}(a) \geq 1-\delta$ implies

$$\mu_{\sigma^l}(a) \geq (1+\Lambda_l)(1-\delta) - \Lambda_l = 1 - (1+\Lambda_l)\delta.$$

Suppose now $K\varepsilon + D\kappa < \bar{\varepsilon}_a(\delta)$. Discarding finitely many terms, every remaining E_l has abnormal probability below $\bar{\varepsilon}_a(\delta)$ by (7), hence is an ε'-elaboration for every ε' between that probability and $\bar{\varepsilon}_a(\delta)$; no lower bound on the abnormal probability is required, and a construction in which every type happens to be normal is admissible as well. Definition 3 therefore supplies an equilibrium of E_l with $\mu_{E_l}(a) \geq 1-\delta$. Choose one for each l, pass to a convergent subsequence of original restrictions, and use (c): the limit σ^* is a feedback equilibrium. Since μ_σ is a polynomial in σ, hence continuous, and $\limsup_l \Lambda_l \leq D\kappa + (K-1)\varepsilon$ by (6), the limit satisfies (8).

All the approximation and payoff-scaling choices above are made after the finite prior P and the maps w have been fixed, and none of the constants κ_0, C, D, K depends on them. ■

A.5 Implementation with Several Players

Filling unused positions. A state must now specify a type for each of the n players, whereas the gates, probes and reads of Lemma 8 involve only two of them. Fill each unused position with a fresh type of that player, used at that state only, and committed to a fixed action by giving that action a strictly higher payoff there. Because a filler occurs in exactly one state, its best response is determined by that state alone, so the commitment is unconditional and requires no argument about the rest of the game. All payoffs of the active helpers are defined to be constant in the fillers' actions, so the incentive computations of Lemma 8 are unchanged.

Reading strategy coordinates. The gate of Step 1 still consists of a player-i controller and a player-j output, and all arithmetic is carried by player-j outputs, so that gates compose in a feedforward circuit exactly as before. A coordinate of an original player-j type is read directly by a controller. A coordinate of any other player, including player i itself and any third player, is transferred to a player-j quantity by the probes of Step 2: the probe state now contains the probe (player j), the original type being probed (its owner), and fillers for the remaining $n-2$ positions, and the probe's high-action payoff there is proportional to $\mathbf{1}_{\{a_h=a\}}$ minus the same threshold as in Step 2, where h is the owner of the probed type. Estimate (39) is unaffected.

Correlated outputs. Enumerate $A_{-i} = \{c_1, \ldots, c_{\bar{k}}\}$ with $\bar{k} = |A_{-i}|$, and use the same cumulative probabilities $F_l(\sigma) = \sum_{p \le l} r(\sigma)(c_p)$, the same circuits, the same Θ labels with midpoints u_m, and the same player-i flags with high probabilities D_{ml}. The single change is in the last stage. For each label m and each opponent $h \neq i$, create a final output type of player h whose score for the action a_h is

$$\mathbf{1}_{\{a_h=(c_{\bar{k}})_h\}} + \sum_{l=1}^{\bar{k}-1} D_{ml}\big(\mathbf{1}_{\{a_h=(c_l)_h\}} - \mathbf{1}_{\{a_h=(c_{l+1})_h\}}\big),$$

that is, (40) with $a_j = c_l$ replaced by $a_h = (c_l)_h$. Place the recipient t_i and the final output types of all $n-1$ opponents carrying the *same* label m in one output state of mass ϖ/Θ.

The verification is the one already carried out in Step 4 of Lemma 8, applied coordinate by coordinate. At a clean label m — one with $|u_m - F_l(\sigma)| > \gamma$ for every l — all flags are pure and correct, so with p the unique index satisfying $F_{p-1} < u_m < F_p$ the telescoping computation gives, for every opponent h simultaneously, score one at $(c_p)_h$ and score zero at every other action of h. Each opponent's output type therefore plays purely, and all of them play the coordinates of the *same* profile c_p, because they are driven by the same flags. The state thus produces the pure profile $c_p \in A_{-i}$, and the mixture over labels produces the midpoint quantile distribution of F, an arbitrary element of $\Delta(A_{-i})$ up to the quantile resolution. No product structure is imposed on the target joint distribution: the correlation is carried by the label, which is a feature of the prior and not of anybody's strategy.

At an unclean label the flags may be mixed, and then the opponents' output types may be indifferent and may randomize arbitrarily. At one fixed label, however, the tuple of opponent types is fixed and their behavioral randomizations are independent, so the conditional action law *at that label* is a product distribution on A_{-i} — an arbitrary one, but a product. The correlation in the output is not manufactured within a label. It comes from the mixture over labels: each clean label forces the coordinates of one common pure profile, and averaging those pure profiles over the label distribution, which is a feature of the prior, realizes an arbitrary correlated element of $\Delta(A_{-i})$. That the fixed-label laws at unclean labels are unknown is harmless, and it is worth being explicit about why: the bound (41) never used anything about the behavior at unclean labels except that their total mass is at most $2(\bar{k}-1)\gamma + 2(\bar{k}-1)/\Theta$, and any probability distribution supported on a set of labels of that mass can move the total variation distance by at most that mass. Hence estimate (41) bounds the error in the *joint* distribution on A_{-i}, at every Bayesian equilibrium, with arbitrary mixing at every tie.

The remaining ingredients of Appendix A.4 carry over unchanged. In Stage 3 of Proposition 1 the output block for a bad type uses a binary target for one opponent in $I \setminus \{i\}$ and commits the other opponents, so the recipient's payoff $2M(\mathbf{1}_{\{\text{that opponent plays high}\}} - \frac{1}{2})$ has the same effect as before. The good-type construction, the mass accounting and the limit argument of Stages 2, 4 and 5 do not refer to the number of players except through the sums recorded in Section 6.4.

For clarity, this construction gives the full many-player version of Proposition 1. For $\kappa > 0$, apply Lemma 7 to the players' matrices and use the same relative mass $\lambda = 2C\kappa$ at their good types. If $\kappa = 0$, omit the good-type output blocks as in Appendix A.4. Every state containing a good original type assigns it exactly g_i at all action profiles. Each new state contains at most one original type; fillers and output types are fresh, so these assignments do not conflict with helper incentives. At bad types, implement each payoff difference with a binary host output as above.

For arbitrary fixed good sets with $P(\prod_i S_i) \geq 1 - \varepsilon$, the total added mass is at most $D\kappa + (K-1)\varepsilon + \tau_l$, with $D = 2nC$ and $K = 1 + n(m-1)$; good original states retain probability at least $(1-\varepsilon)/(1+\Lambda_l)$. This proves parts (a) and (b). At every good or bad original type, the unnormalized payoff numerators converge to the vectors used in Stage 5 of Appendix A.4, since the actuator error is now controlled in joint TV on A_{-i}. The same support argument proves part (c) for every equilibrium sequence. Finally, the action-measure decomposition and robustness argument in Stage 6 use only the original-state mass and are independent of the number of players, giving part (d) with these constants. This proof uses no conclusion of Theorem 1.

B Analysis

B.1 The Parabolic Equation

This appendix proves Lemma 2. The argument is self-contained and uses only the Gaussian heat semigroup; it is placed here because it interacts with nothing else in the paper. The statement being proved is repeated for convenience.

Lemma 2. *Let $f : \mathbb{R}^d \to \mathbb{R}^d$ be bounded and globally Lipschitz, let $\nu > 0$, and let $V : \mathbb{R}^d \to \mathbb{R}^d$ be smooth with bounded uniformly continuous derivatives of every order. Then*

$$\partial_t v + Dv\, f(v) + \frac{\nu^2}{2}\Delta v = 0, \qquad v(1,x) = V(x)$$

has a bounded classical solution with bounded spatial gradient on $[0,1] \times \mathbb{R}^d$, unique among mild solutions that are continuous in time with values in the space of bounded uniformly continuously differentiable functions and have bounded norm in that space. For every $0 < \beta < 1$ its Hessian and *time derivative are uniformly β-Hölder in space and uniformly $\beta/2$-Hölder in time, with bounds and Hölder constants global in x.*

Two remarks on the structure of the proof before it begins. The equation is quasilinear, so the standard route is a fixed point in a space of functions with one bounded derivative, followed by a bootstrap for higher regularity; the bootstrap must be done by hand because f is only Lipschitz and cannot be differentiated. The reason the local solution extends to the whole interval, rather than blowing up, is that the nonlinearity is $Du\, f(u)$ with f bounded, so the gradient obeys the *linear* Volterra inequality (45); this is the step to watch.

Proof of Lemma 2. *Step 0: reversal of time and the mild formulation.* Write $B = \sup_y |f(y)|$ and let L_f be a global Lipschitz constant of f. Put $u(t,x) = v(1-t,x)$, so that u solves the forward problem with initial datum V, and let P_t be the Gaussian heat semigroup with covariance $\nu^2 t I_d$, that is, convolution with the centered Gaussian density of that covariance. The mild form of the problem is

$$u(t) = P_t V + \int_0^t P_{t-s} N(u(s))\, ds, \qquad N(u) = Du\, f(u). \tag{43}$$

Work in the Banach space $\mathcal{C}^1$ of bounded uniformly continuously differentiable maps $\mathbb{R}^d \to \mathbb{R}^d$ with the norm $\|u\|_\infty + \|Du\|_\infty$, and in $C([0,T];\mathcal{C}^1)$ for the time-dependent problem. Differentiating the Gaussian kernel gives the two estimates

$$\|P_t\phi\|_\infty \le \|\phi\|_\infty, \qquad \|DP_t\phi\|_\infty \le \frac{C_d}{\nu\sqrt{t}}\|\phi\|_\infty, \tag{44}$$

the second because the spatial derivative of the kernel has L^1 norm of order $(\nu\sqrt{t})^{-1}$.

Step 1: local existence and uniqueness. On a ball of radius R in $C([0,T];\mathcal{C}^1)$,

$$\|N(u) - N(w)\|_\infty \le B\|Du - Dw\|_\infty + L_f R\|u - w\|_\infty,$$

using $|Du\, f(u) - Dw\, f(w)| \le |Du - Dw||f(u)| + |Dw||f(u) - f(w)|$. Combining this with (44) and integrating the singularity $\int_0^T (T-s)^{-1/2} ds = 2\sqrt{T}$, the integral operator in (43) has Lipschitz constant at most

$$\left(T + \frac{2C_d\sqrt{T}}{\nu}\right)(B + L_f R)$$

in the norm of $C([0,T];\mathcal{C}^1)$. Choosing R larger than twice the norm of V and then T small makes the map $u \mapsto P_\cdot V + \int_0^\cdot P_{\cdot-s} N(u(s)) ds$ a contraction of that ball into itself. The heat semigroup is strongly continuous on $\mathcal{C}^1$, so the fixed point is a continuous mild solution on $[0,T]$. The local existence time depends only on R, B, L_f, ν and d. The same contraction, applied on successive intervals, gives uniqueness in the stated class.

Step 2: an a priori gradient bound, and global existence. Let $L(t) = \|Du(t)\|_\infty$, $L_0 = \|DV\|_\infty$, and $\varkappa = C_d B/\nu$. Since $\|N(u(t))\|_\infty \le B\,L(t)$, differentiating (43) in space and using (44) gives

$$L(t) \le L_0 + \varkappa \int_0^t (t-s)^{-1/2} L(s)\,ds. \tag{45}$$

This inequality is linear in L, which is the crucial point: the Lipschitz constant L_f does not appear. Substituting (45) into its own right-hand side and using

$$\int_r^t (t-s)^{-1/2}(s-r)^{-1/2} ds = \pi, \qquad \int_0^t (t-s)^{-1/2} ds = 2\sqrt{t} \le 2$$

for $t \le 1$, we get $L(t) \le L_0(1+2\varkappa) + \pi\varkappa^2 \int_0^t L(r)\,dr$, and Gronwall's inequality yields

$$L(t) \le L_0(1+2\varkappa)e^{\pi\varkappa^2 t} \qquad (0 \le t \le 1). \tag{46}$$

Also $\|u(t)\|_\infty \le \|V\|_\infty + B\int_0^t L(s)ds$ is bounded on $[0,1]$. Since the local existence time of Step 1 depends on the solution only through the radius R, and (46) bounds that radius uniformly on $[0,1]$, the solution cannot cease to exist before time 1: starting sufficiently close to a putative maximal time would extend it past that time. So the mild solution exists on all of $[0,1]$, with $\|u\|_\infty + \|Du\|_\infty$ bounded there.

Step 3: Hölder estimates for the heat semigroup. For a spatial function ϕ write $[\phi]_\alpha$ for its global α-Hölder seminorm. Gaussian scaling, the vanishing integral of each derivative of the kernel, and interpolation between consecutive derivative bounds give, for $0 < \alpha < 1$,

$$[DP_t\phi]_\alpha \le Ct^{-(1+\alpha)/2}\|\phi\|_\infty, \tag{47}$$

$$\|D^j P_t\phi\|_\infty \le Ct^{(\alpha-j)/2}[\phi]_\alpha \qquad (j \ge 1). \tag{48}$$

Estimate (48) follows by writing the differentiated convolution against $\phi(x-z) - \phi(x)$, which is legitimate because $\int D^j p_t(z)dz = 0$ for $j \ge 1$, and then using $|\phi(x-z) - \phi(x)| \le [\phi]_\alpha |z|^\alpha$ together with the Gaussian moment $\int |D^j p_t(z)||z|^\alpha dz \le Ct^{(\alpha-j)/2}$. Estimate (47) follows by interpolating $\|DP_t\phi\|_\infty \le Ct^{-1/2}\|\phi\|_\infty$ and $\|D^2P_t\phi\|_\infty \le Ct^{-1}\|\phi\|_\infty$. Constants here may depend on ν and d.

Step 4: spatial regularity. The forcing $N(u)$ is uniformly bounded by (46). Applying (47) to the spatially differentiated mild equation,

$$[Du(t)]_\alpha \le [DP_tV]_\alpha + \int_0^t C(t-s)^{-(1+\alpha)/2}\|N(u(s))\|_\infty ds,$$

whose time singularity is integrable for every $\alpha < 1$; the first term is bounded because V is smooth with bounded derivatives. So Du is uniformly α-Hölder in space, for every $\alpha < 1$. Consequently $N(u) = Du\,f(u)$ is uniformly α-Hölder in space as well: Du is bounded and α-Hölder, $f(u)$ is bounded and Lipschitz in x with constant $L_f\|Du\|_\infty$, hence α-Hölder, and a product of bounded α-Hölder functions is α-Hölder.

Fix $0 < \beta < \alpha < 1$. Interpolating the cases $j = 2$ and $j = 3$ of (48) gives

$$[D^2 P_t \phi]_\beta \le C t^{-1+(\alpha-\beta)/2} [\phi]_\alpha, \tag{49}$$

and both this singularity and the one in $\|D^2 P_t \phi\|_\infty \le C t^{-1+\alpha/2}[\phi]_\alpha$ are integrable. Differentiating the mild equation twice in space therefore yields a Hessian that is uniformly bounded and uniformly β-Hölder in space. The differentiation under the integral is justified first with the integral truncated away from $s = t$, where the integrand is smooth, and then in the limit using the integrable bounds just displayed.

Step 5: regularity in time. Restart the mild equation at time t: for $t, t+h \in [0, 1]$,

$$u(t+h) = P_h u(t) + \int_t^{t+h} P_{t+h-s} N(u(s))\, ds.$$

Its zeroth and first derivative versions give

$$\|u(t+h) - u(t)\|_\infty \le C\sqrt{h} + Ch, \qquad \|Du(t+h) - Du(t)\|_\infty \le Ch^{\alpha/2} + C\sqrt{h},$$

where the heat increment of $u(t)$ is estimated by its bounded gradient and that of $Du(t)$ by its spatial α-Hölder bound, in both cases through $\|P_h\phi - \phi\|_\infty \le [\phi]_\gamma (C\nu\sqrt{h})^\gamma$. Hence $N(u)$ is uniformly $\alpha/2$-Hölder in time. Similarly,

$$D^2 u(t+h) = P_h D^2 u(t) + \int_t^{t+h} D^2 P_{t+h-s} N(u(s))\, ds$$

shows that $D^2 u$ is uniformly $\beta/2$-Hölder in time: the first term changes by at most $Ch^{\beta/2}$ by the spatial β-Hölder bound on $D^2 u$, and the integral is at most $Ch^{\alpha/2}$ by the integrable singularity of $\|D^2 P_\cdot\|$ against an α-Hölder forcing.

Step 6: the classical equation. Dividing the restarted mild equation by h and letting $h \downarrow 0$,

$$u_t = \frac{\nu^2}{2} \Delta u + N(u) :$$

the heat increment divided by h converges uniformly to $(\nu^2/2)\Delta u$ because the Hessian is bounded and uniformly continuous, and the averaged forcing converges uniformly to $N(u(t))$ because $N(u)$ is continuous in time. The right-hand side is continuous in (t, x), so this right-derivative identity gives the classical time derivative, with one-sided interpretations at the endpoints. The estimates of Steps 4 and 5 then show that u_t is uniformly β-Hölder in space and $\beta/2$-Hölder in time, and the smooth bounded initial datum ensures that all estimates include $t = 0$. Reversing time, $v(t, x) = u(1 - t, x)$ has all the stated properties and solves (21). ■

B.2 Uniform Comparison on Finite Noise Trees

Proof of Lemma 3. *Step 1: the one-step consistency estimate.* Since f is bounded by $B = \sup_y |f(y)|$ and $\rho \le 1$, the controls satisfy $|s_k| \le B+1$. We claim that, uniformly in $t \in [0, 1-h]$, in $x \in \mathbb{R}^d$ and in $|b| \le B+1$, and for x, b that are $\mathcal{F}_k$-measurable,

$$\begin{aligned}&\mathbb{E}[v(t+h, x+hb+\nu\sqrt{h}\,\xi_{k+1}) \mid \mathcal{F}_k] \\ &\quad = v(t,x) + h\left(v_t + Dv\,b + \frac{\nu^2}{2}\Delta v\right)(t,x) + R, \qquad |R| \le Ch^{1+\beta/2}.\end{aligned} \tag{50}$$

To see this, write $\Delta_h = hb + \nu\sqrt{h}\,\xi_{k+1}$ and note $|\Delta_h| \le (B+1)h + \nu M\sqrt{h} \le C\sqrt{h}$. Expand in space at $(t+h, x)$ to second order with the integral form of the remainder; since the Hessian of $v(t+h,\cdot)$ is β-Hölder with a global constant, the remainder is bounded by $C|\Delta_h|^{2+\beta} \le Ch^{1+\beta/2}$. Taking the conditional expectation and using $\mathbb{E}[\xi_{k+1} \mid \mathcal{F}_k] = 0$ and $\mathbb{E}[\xi_{k+1}\xi_{k+1}^\top \mid \mathcal{F}_k] = I_d$ turns the first-order term into $hDv(t+h,x)b$ and the second-order term into $\frac{\nu^2 h}{2}\Delta v(t+h,x) + O(h^2)$, the $O(h^2)$ coming from the drift part of the quadratic term. It remains to replace $t+h$ by t in the three terms. For the value, $v(t+h,x) - v(t,x) = \int_t^{t+h} v_t(u,x)\,du = hv_t(t,x) + O(h \cdot h^{\beta/2})$ by the temporal $\beta/2$-Hölder continuity of v_t. For the other two terms, which already carry a factor h, the temporal Hölder continuity of Dv and of $D^2 v$ contributes $O(h^{1+\beta/2})$. Since $|\xi_{k+1}| \le M$, all these bounds are uniform over the noise tree, including on branches far from the origin, which is where global-in-x regularity is used.

Step 2: the backward estimate (26). Set $\mathcal{D}_k = Y_k - v(t_k, X_k)$. The tower property gives $Y_k = \mathbb{E}[Y_{k+1} \mid \mathcal{F}_k]$. Applying (50) with $x = X_k$ and $b = s_k$, both $\mathcal{F}_k$-measurable, and then using (21) at (t_k, X_k) to replace $v_t + \frac{\nu^2}{2}\Delta v$ by $-Dv\,f(v)$, we obtain

$$\mathcal{D}_k = \mathbb{E}[\mathcal{D}_{k+1} \mid \mathcal{F}_k] + hDv(t_k, X_k)\big(s_k - f(v(t_k, X_k))\big) + R_k, \qquad |R_k| \le Ch^{1+\beta/2}.$$

Let $L_v = \sup_{t,x} |Dv(t,x)|$, finite by Lemma 2, and $C_0 = L_v L_f$ with L_f a Lipschitz constant of f. By (25),

$$|s_k - f(v(t_k, X_k))| \le |s_k - f(Y_k)| + |f(Y_k) - f(v(t_k, X_k))| \le \rho + L_f|\mathcal{D}_k|,$$

so taking essential suprema and using $\|\mathbb{E}[\,\cdot \mid \mathcal{F}_k]\|_\infty \le \|\cdot\|_\infty$,

$$(1 - hC_0)\|\mathcal{D}_k\|_\infty \le \|\mathcal{D}_{k+1}\|_\infty + hL_v\rho + Ch^{1+\beta/2}.$$

Assume N large enough that $hC_0 \le 1/2$, so $(1-hC_0)^{-1} \le 1 + 2hC_0$. Iterating backwards from $\mathcal{D}_N = 0$, which holds because $Y_N = V(X_N) = v(1, X_N)$, gives

$$\|\mathcal{D}_k\|_\infty \le \sum_{n=k}^{N-1} (1+2hC_0)^{n-k+1}\big(hL_v\rho + Ch^{1+\beta/2}\big) \le e^{2C_0} N\big(hL_v\rho + Ch^{1+\beta/2}\big),$$

and $Nh = 1$ turns the right side into $e^{2C_0}(L_v\rho + Ch^{\beta/2})$. This is (26); the case $C_0 = 0$ follows by direct summation.

Step 3: the forward estimate (27). The drift field $b(t, x) = f(v(t, x))$ is bounded and globally Lipschitz in x with constant C_0. By Step 2,

$$|s_k - b(t_k, X_k)| \le \rho + L_f|\mathcal{D}_k| \le C(\rho + h^{\beta/2})$$

almost surely. Subtract (22) from (24): the noise increments are identical and cancel pathwise, leaving

$$X_{k+1} - \widehat{X}_{k+1} = X_k - \widehat{X}_k + h\big(s_k - b(t_k, \widehat{X}_k)\big),$$

and since $|b(t_k, X_k) - b(t_k, \widehat{X}_k)| \le C_0|X_k - \widehat{X}_k|$,

$$|X_{k+1} - \widehat{X}_{k+1}| \le (1 + hC_0)|X_k - \widehat{X}_k| + Ch(\rho + h^{\beta/2}).$$

The discrete Gronwall inequality, starting from $X_0 = \widehat{X}_0$ and using $Nh = 1$, gives the state bound. The control bound follows from $|s_k - \widehat{s}_k| \le |s_k - b(t_k, X_k)| + C_0|X_k - \widehat{X}_k|$. No Markov property of s_k has been used at any point: the controls enter only through their values on the realized path. ∎

References

A. Kajii and S. Morris. The robustness of equilibria to incomplete information. *Econometrica*, 65 (6):1283–1309, 1997.

S. Morris and T. Ui. Generalized potentials and robust sets of equilibria. *Journal of Economic Theory*, 124(1):45–78, 2005.