

Cross-Validation Equilibrium*

Ran Spiegler[†] and Stephan Waizmann[‡]

June 10, 2026

We study strategic interaction when players delegate belief formation to predictive machine learning (ML). In a static Bayesian game, each player's ML agent predicts a payoff-relevant outcome variable as a function of the player's type. The ML agent's training sample is endogenous: it is drawn from the outcome distribution generated by players' ML-guided behavior. In Cross-Validation Equilibrium (CVE), each player's ML agent selects a predictive model to minimize expected out-of-sample squared error, given its realized training sample; and each player best-responds to the belief generated by the model her ML agent selected. We analyze CVE and relate it to other equilibrium concepts. We apply CVE to jury voting, speculative betting, and games with linear-quadratic payoffs. E.g., in a team-effort game, endogenous model selection can give rise to multiple equilibria.

*We acknowledge financial support from UKRI Frontier Research Grant no. EP/Y033361/1. We thank participants at the 5th Southeast Theory Festival and the referees at EC2026, for helpful comments.

[†]Tel Aviv University and University College London, rani@tauex.tau.ac.il.

[‡]University College London, stephanwaizmann@gmail.com.

1. Introduction

The use of machine learning (ML) algorithms to form predictive beliefs in economic settings has become ubiquitous. Credit agencies employ ML to evaluate borrowers' creditworthiness. Judges make use of pretrial risk assessment algorithms to predict whether a person accused of a crime will fail to show up at court or engage in criminal activity if they are bailed out. In these and other cases, ML algorithms use a random training sample to generate conditional probabilistic predictions; the ultimate decision power rests with the human decision-maker who employs these tools.

This paper addresses the following question: What are the effects of predictive ML when its users are players in a static game with possibly incomplete information? In this context, an individual player tasks an "ML agent" with constructing a predictive model (selected on the basis of a random training sample) that outputs a belief over payoff-relevant variables (e.g., an opponent's action) as a function of the player's signal. The random training sample used by the ML agent is drawn from an endogenous distribution induced by all players' responses to their own ML-generated beliefs. We are interested in the implications of this two-way linkage between players' behavior and ML-generated predictive models.

We develop a game-theoretic equilibrium concept, dubbed *Cross-Validation Equilibrium (CVE)*, in which players' equilibrium beliefs are formed according to a stylized model of predictive ML. Our concept captures the idea that predictive ML has access to a noisy training sample — drawn from the equilibrium distribution — and uses it to construct a predictive model of a payoff-relevant variable. The predictive model is selected from some feasible set of models according to its *out-of-sample predictive performance*, measured by the *expected mean squared prediction error*. Thus, while predictive models are estimated against the training sample, their selection is based on the out-of-sample predictive success of each estimated model. This model-selection criterion captures in stylized form the notion of *cross validation* that characterizes many ML tools (e.g., see Chapter 7 in Hastie, Tibshirani, and Friedman (2009)).

Since players' endogenous training samples are random, so is the predictive model they end up adopting, and therefore so is their resulting behavior. Thus, CVE involves randomness in players' subjective models, the beliefs these models generate, and the behavior they induce. Nash equilibrium is a limit case of our concept, corresponding to CVE with perfectly noiseless samples, which give rise to correct models and therefore correct beliefs. Another limit case is Analogy Based Expectation Equilibrium Jehiel

(2005), which corresponds to a CVE in which each player’s model is exogenously given and training samples are noiseless.

For a simple example that illustrates our modeling approach, consider a standard symmetric two-firm Cournot competition with linear demand and a random, commonly observed demand parameter θ . Firm i ’s ML agent aims to predict firm j ’s expected quantity as a function of θ . Imagine that for each value of θ , firm i ’s ML agent observes firm j ’s expected quantity conditional on θ plus independent Gaussian noise. This noisy observation is *not* a standard game-theoretic signal; it is an observation of the *equilibrium* link between θ and firm j ’s behavior. The profile of noisy measurements constitutes the “training sample” at the disposal of firm i ’s ML agent.

Firm i ’s ML agent selects a *predictive model*, which (following Jehiel (2005)) is a *partition* of the set of parameter values. For each partition cell, the prediction is the weighted training-sample average of firm j ’s quantity conditional on the cell. The ML agent chooses a partition that minimizes the expected *out-of-sample squared prediction error*. This procedure naturally involves a bias-variance trade-off: For some training-sample noise realizations, the ML agent will adopt a coarse partition that bundles together different values of θ , even if they are associated with different behavior by firm j . In CVE, firm i best-responds to the prediction induced by the selected partition.

We establish that a CVE exists in finite environments. Moreover, a CVE converges to an ABEE with respect to the finest feasible partition as the training-sample noise vanishes. We also highlight basic properties of CVE. First, the expected training-sample noise conditional on any partition cell is zero. This unbiasedness property, which improves the concept’s tractability, is not obvious a priori, as the partition itself is endogenously selected. It arises from our criterion for out-of-sample predictive success, and from the assumption of symmetric noise. Second, the ML agent avoids over-fitting, in the sense that if the predicted variable is measurable with respect to some partition, the ML agent will (almost) never adopt a finer one.

We then apply CVE to various games, to explore its economic applications. The concept is especially tractable in games with linear-quadratic payoffs. For example, in the above-mentioned Cournot competition example (when the only feasible partitions are the fully fine and fully coarse ones), we show that the volatility of quantities and prices in the unique symmetric CVE rises with the variance of the training-sample noise. We also show that unlike Nash equilibrium, asymmetric CVE are possible in this game. Another example involves a team effort problem, where unlike the Cournot example, the game exhibits strategic complementarities. As a result, and despite the linear best-

replies that imply a unique Nash equilibrium, there can be *multiple* symmetric CVE. We also establish a sense in which CVE in the team effort game is discontinuously different from Nash equilibrium, even when the training sample noise is arbitrarily small. We also apply CVE to a two-player version of Feddersen and Pesendorfer (1998)’s jury model, as well as a simple two-player game of speculative trade, where we show that when one player is restricted to use the correct model, this player may make an expected loss to her opponent in CVE.

The rest of the paper is organized as follows. Section 2 presents the model and defines Cross-Validation Equilibrium. The general results are contained in Section 3. In Section 4, we apply CVE to examples. Proofs of the main results are in the Appendix.

2. The Model

Consider the following, slightly non-standard formulation of a Bayesian game. The set of players is $N = \{1, \dots, n\}$. Let Θ be a set of states of Nature. For each player $i \in N$, let A_i be the player’s action set, and let T_i be the set of types for this player. Denote $T = \times_{i \in N} T_i$ and $A = \times_{i \in N} A_i$. Let $\mu \in \Delta(\Theta \times T)$ be an objective prior distribution over the exogenous variables. A strategy for player i is a function $\sigma_i : T_i \rightarrow \Delta(A_i)$.

So far, the exposition is standard. We slightly deviate from the norm in our definition of payoffs. Let $X_i = \mathbb{R}$ be a set of *payoff-relevant outcomes* for player i . Let $f_i : \Theta \times A_{-i} \rightarrow X_i$ be a function that maps states and profiles of actions by player i ’s opponents to her payoff-relevant outcomes. Let $u_i : A_i \times T_i \times X_i \rightarrow \mathbb{R}$ be player i ’s *continuous* vNM utility function, which is in particular *linear* in x_i . The scalar variable x_i is a “sufficient statistic” that pins down player i ’s payoff given her action and type. This variable is what player i ’s ML-based belief-formation procedure will aim to predict.

The following examples illustrate this non-standard aspect of the formalism. In both examples, we leave the information structure given by μ unspecified, because it is immaterial for the illustration.

Example 2.1 (mixtures in binary-action games). Let $n = 2$ and $A_1 = A_2 = [0, 1]$, where a_i represents the probability that player i chooses a particular action in a binary-action game. In this case, let $f_i(\theta, a_j) = a_j$. When u_i is player i ’s expected utility in the binary-action game, linearity of u_i in x_i follows automatically.

Example 2.2 (a linear Cournot game). Let $n = 2$ and $A_1 = A_2 = \mathbb{R}$. Let θ be a demand parameter, such that $u_i = a_i[\theta - a_i - a_j]$. (Allow the price $\theta - a_i - a_j$ to take negative

values.) Consider two alternative specifications of x_i . First, let $f_i(\theta, a_j) = a_j$ — i.e., the sufficient statistic for u_i is player j 's quantity. Second, let $f_i(\theta, a_j) = \theta - a_j$ — i.e., the sufficient statistic for u_i is the residual demand that player i faces.

The combination of μ and a strategy profile $\sigma = (\sigma_1, \dots, \sigma_2)$ induces a joint distribution over $\Theta \times T \times A$,

$$p^\sigma(\theta, t, a) = \mu(\theta, t) \cdot \prod_{i \in N} \sigma_i(a_i | t_i) \quad (1)$$

For convenience, we also use the notation p^σ for marginal and conditional distributions derived from this joint distribution. Let

$$\bar{x}_i^\sigma(t_i) = \sum_{\theta, a_{-i}} p^\sigma(\theta, a_{-i} | t_i) f_i(\theta, a_{-i}) \quad (2)$$

be the expectation of x_i conditional on player i 's type, according to p^σ . Note that $p^\sigma(\theta, a_{-i} | t_i)$ — and therefore also \bar{x}_i^σ — only depend on the profile σ_{-i} ; they are invariant to player i 's own mixed strategy σ_i .

We now describe how player i forms a prediction of her payoff-relevant outcome given her type. Fix σ_{-i} . This means that $\bar{x}_i^\sigma(t_i)$ is also held fixed for every t_i . An *ML agent* at player i 's service has access to a *training sample*, described by

$$y^{t_i} = \bar{x}_i^\sigma(t_i) + \varepsilon_{t_i} \quad (3)$$

for every $t_i \in T$, where ε_{t_i} is *i.i.d* according to a distribution ν that is distributed symmetrically around zero. Since $\bar{x}_i^\sigma(t_i)$ is calculated with respect to the conditional distribution $p^\sigma(\theta, a_{-i} | t_i)$, the training sample can be viewed as a “representative sample” in the spirit of Danenberg and Spiegler (2025).

For expositional convenience, we will assume here that ν has finite support (in applications, we employ continuous noise as well). Thus, given $\bar{x}_i^\sigma = (\bar{x}_i^\sigma(t_i))_{t_i \in T_i}$, we identify the sample with $\varepsilon_i = (\varepsilon_{t_i})_{t_i \in T_i}$. For any $M \subseteq T_i$, let

$$\hat{p}^{\varepsilon_i}(x_i | M) = \sum_{t_i \in M} \mu(t_i | t_i \in M) \cdot \mathbf{1}\{\bar{x}_i^\sigma(t_i) + \varepsilon_{t_i} = x_i\} \quad (4)$$

be the frequency of x_i in the sub-sample $\{\varepsilon_{t_i}\}_{t_i \in M}$, weighted according to the distribution of t_i conditional on $t_i \in M$. Likewise, let

$$\hat{x}^{\varepsilon_i}(M) = \sum_{t_i \in M} \mu(t_i | t_i \in M) (\bar{x}_i^\sigma(t_i) + \varepsilon_{t_i}) \quad (5)$$

be the weighted average value of x_i in the sub-sample.

Given the sample ε_i , the ML agent chooses a partition $\Pi_i(\varepsilon_i)$ of T_i from some family \mathcal{P}_i of feasible partitions. Let $\pi(t_i) \in \Pi_i(\varepsilon_i)$ be the partition cell that includes t_i . The ML agent evaluates the partition's out-of-sample predictive performance with a *validation sample*, which observes $\bar{x}_i^\sigma(t_i)$ with *i.i.d* zero-mean noise η_{t_i} for each t_i . The ML agent selects the partition with the best average out-of-sample predictive performance. This partition generates a belief, and player i best-responds to this belief. In what follows, $g_i : T_i \rightarrow A_i$ denotes a pure ex-ante strategy for player i

Definition 1. A pair (Π_i, g_i) is compatible with σ_{-i} and ε_i if

$$\Pi_i \in \arg \min_{\Pi'_i \in \mathcal{P}_i} \mathbb{E}_\eta \sum_{t_i \in T_i} \mu(t_i) [\hat{x}^{\varepsilon_i}(\pi'(t_i)) - (\bar{x}_i^\sigma(t_i) + \eta_{t_i})]^2 \quad (6)$$

and, for every $t_i \in T_i$,

$$g_i(t_i) \in \arg \max_{a_i \in A_i} \sum_{x_i \in X_i} \hat{p}^{\varepsilon_i}(x_i | \pi(t_i)) u_i(a_i, t_i, x_i) \quad (7)$$

The first condition requires that given the training sample ε_i , the partition Π_i minimizes the expected out-of-sample mean squared prediction error, with respect to the true expected payoff-relevant outcomes in the validation sample given by \bar{x}_i^σ . The expectation with respect to η means that the ML agent evaluates a partition's predictive performance across many validation samples. The second condition is that player i 's ex-ante pure strategy prescribes a best-reply to the belief induced by the partition Π . Note that since u_i is assumed to be linear in x_i , we can rewrite player i 's expected payoff in (7) as

$$u_i(a_i, t_i, \sum_{x_i \in X_i} \hat{p}^{\varepsilon_i}(x_i | \pi(t_i)) x_i). \quad (8)$$

We are now able to close the model and define when a strategy profile constitutes an equilibrium.

Definition 2. A strategy profile σ is a *Cross-Validation Equilibrium (CVE)* if for every $i \in N$ and every training sample realization ε_i , there is a mixture β_{ε_i} over pairs (Π_i, g_i) that are compatible with σ_{-i} and ε_i , such that for every $t_i \in T_i$,

$$\sigma_i(a_i | t_i) = \sum_{\varepsilon_i} v(\varepsilon_i) \sum_{(\Pi_i, g_i) | g_i(t_i) = a_i} \beta_{\varepsilon_i}(\Pi_i, g_i) \quad (9)$$

Thus, in CVE, a player’s mixed strategy arises from her ML-based belief-formation procedure. The player’s ML agent has access to a random sample of the type-dependent payoff-relevant outcome \bar{x}_i^σ , forms an out-of-sample-optimal partition of T_i that induces a type-dependent belief over payoff-relevant outcome. The player best-responds to this belief. Because the player’s sample is random, so is her best-responding behavior, and this randomness is described by the player’s mixed strategy.

A simple illustrative example Consider the following symmetric, binary-action, two-person game with complete information. Players’ action set is $\{0, 1\}$. The set of states of Nature is $\{0, 1\}$. The two states are equally likely. The state is commonly known so that $t_i = \theta$ with probability 1. When $\theta = 0$, player i ’s payoff is $-a_i$. When $\theta = 1$, player i ’s payoff is $a_i(\frac{4}{3}a_j - 1)$. In the notation of the model,

$$u_i(a_i, \theta, x) = \begin{cases} -a_i & \theta = 0 \\ a_i(\frac{4}{3}x - 1) & \theta = 1 \end{cases}$$

and $f_i(\theta, a_j) = a_j$. Consequently, we can interpret $\bar{x}_i^\sigma(\theta) = \sigma_j(1 | \theta)$ as the probability with which player j plays action $a_j = 1$ when the state is θ .

Clearly, the game has a Nash equilibrium in which both players play $a = \theta$.¹ This is the efficient outcome. Note that player i ’s best-reply in state $\theta = 1$ is $a_i = 1$ if she assigns probability above $\frac{3}{4}$ to player j choosing $a_j = 1$, i.e., if $\bar{x}_i^\sigma(1) \geq \frac{3}{4}$.

Assume the training sample noise distribution is uniform over $\{-d, d\}$, where $d > \frac{1}{2}$. Use V to denote the variance of the validation sample noise. We solve for symmetric CVE. Denote $\bar{x}_\theta = \bar{x}_i^\sigma(\theta)$.

In state $\theta = 0$, each player i strictly prefers playing $a = 0$, independently of her belief regarding x_i . Consequently, $\bar{x}_0 = 0$. Therefore, we only need to derive the probability that each player chooses $a = 1$ in state $\theta = 1$. This probability equals \bar{x}_1 .

The expected out-of-sample prediction error when adopting the fine partition $0, 1$ is $\frac{1}{2}(\varepsilon_0^2 + \varepsilon_1^2) + V$, whereas the out-of-sample prediction error from adopting the coarse partition $0, 1$ is $(\frac{\bar{x}_1}{2} + \frac{\varepsilon_0 + \varepsilon_1}{2})^2 + V$. When $(\varepsilon_1 - \varepsilon_0)^2 \geq (\bar{x}_1 - \bar{x}_0)^2 = \bar{x}_1^2$, player i ’s ML agent adopts the coarse partition, and the player’s estimate of x_i is $\frac{1}{2}(\bar{x}_1 + \varepsilon_0 + \varepsilon_1)$. When $(\varepsilon_1 - \varepsilon_0)^2 < \bar{x}_1^2$, player i ’s ML agent adopts the fine partition, and the player’s estimate of x_i is $\bar{x}_1 + \varepsilon_1$. The player will play $a = 1$ only if her estimate of x_i is weakly above $\frac{3}{4}$.

¹In state $\theta = 0$, this is the only Nash equilibrium. In state $\theta = 1$, there are two more Nash equilibria: $a_1 = a_2 = 0$ and $\sigma_1(1) = \sigma_2(1) = 3/4$.

Let us now use these conditions, combined with the noise distribution, to derive \bar{x}_1 . When $\varepsilon_1 = \varepsilon_0$, the ML agent necessarily adopts the fine partition. Thus, the coarse partition is adopted only if $\varepsilon_1 = -\varepsilon_0$. But in this case, the player plays $a = 1$ only if

$$\frac{1}{2}(\bar{x}_1 + \varepsilon_0 + \varepsilon_1) = \frac{1}{2}\bar{x}_1 > \frac{3}{4}$$

which this is impossible. It follows that \bar{x}_1 is bounded from above by the probability that an ML agent adopts the fine partition.

When a player's ML agent adopts the fine partition, she plays $a = 1$ only if $\bar{x}_1 + \varepsilon_1 \geq \frac{3}{4}$. Since $d > \frac{1}{2}$, it follows that $\varepsilon_1 = d$ is necessary for playing $a = 1$. Thus, $\bar{x}_1 \leq \frac{1}{2}$. If $\varepsilon_0 = -d$, then

$$(\varepsilon_1 - \varepsilon_0)^2 = 4d^2 > \frac{1}{4} > \bar{x}_1^2$$

such that the ML agent adopts the coarse partition. It follows that players choose $a = 1$ only if $\varepsilon_1 = \varepsilon_0 = d$, such that $\bar{x}_1 \leq \frac{1}{4}$. It follows that there exists a CVE with $\bar{x}_1 = \frac{1}{4}$.

Discussion We conclude this section with a discussion of some of our modeling choices.

In our framework, the player's belief formation is entirely delegated to the ML agent. The player herself is not assumed to have a thorough understanding of the game she is playing. Indeed, she need not perceive the situation as a game; she is interested in quantifying the mapping from t_i to x_i . In this sense, our model follows the approach of Osborne and Rubinstein (1998). But while in that paper, players relied on naive sampling to evaluate such mappings, in our model the player employs an ML agent to quantify the mapping. The player does understand that x_i should be attributed to t_i only, i.e., that it is independent of a_i . She may also have certain preconceptions that limit the set of partitions she allows the ML agent to explore. But that is the only extent of her understanding of the situation, as far as belief formation is concerned.

The space of models that a player's ML agent explores is a set of partitions of the player's type space. Of course, one could imagine other classes of predictive models; our use of partitions can be regarded as a "proof of concept." It has two advantages. First, it is non-parametric and can be applied to any Bayesian game (whereas parametric predictive models would require knowledge of the specifics of the Bayesian game). Second, it enables us to form a link to the concept of Analogy-Based Expectations Equilibrium (ABEE, due to Jehiel (2005)), as we will see in Section 3.

Note that the domain of partitions that the ML agent considers may be restricted. There are two motivations for this feature of our framework. First, the player may have

a-priori theoretical reasons to impose structure on the set of predictive models. For example, when T_i is two-dimensional, i.e., it is the product of two sets, the player may believe that this product structure should be preserved in any model that predicts x_i . Second, we rely on restricted domains in examples for tractability, e.g. by assuming that the only feasible partitions are the fully coarse and fully fine ones. Tractability considerations also lie behind our treatment of the training sample noise — specifically, our assumption that the noise is additive and independent of the true expected value of the predicted payoff-relevant outcome.

Our assumption that x_i is a scalar is made for expositional simplicity. A natural extension would be that $x \in \mathbb{R}^n$, and that out-of-sample predictive fit would be evaluated by the sum of squared prediction errors across dimensions. Our assumption that player i 's payoff is linear in x_i should be viewed in this context. Linearity implies that when player i has uncertainty about x_i , the mean with respect to her belief is a sufficient statistic for predicting her payoff — i.e., she only needs to predict a scalar variable in order to predict her payoff. If we relaxed linearity, the restriction to a scalar x_i would be less well-founded. Without linearity, the tight link between CVE and ABEE would break down.

CVE is based on a criterion for out-of-sample predictive success that involves taking an expectation over the validation sample noise. One interpretation is that the ML agent's estimated model is pitted against multiple (independent) validation samples, and the model is evaluated according to its average predictive success across these samples. However, our motivation behind this assumption is tractability: as we will see below, it makes the analysis of CVE more tractable. Note that the exact shape of the validation sample noise is immaterial for our purposes.

3. General Results

In this section, we present results that hold for all finite games. We show that a CVE exists. Furthermore, we consider convergence of CVE as the training sample becomes perfect. We conclude with two results that illustrate properties of the out-of-sample prediction criterion. All proofs are contained in the [Appendix](#).

Say that the game is finite if each player's action set, type space, the set of nature, and the support of the training sample is finite.

Proposition 1. *In every finite game, a CVE exists.*

The proof applies a standard fixed-point argument to a correspondence Ψ that maps, for each realization of the training sample, a distribution over partitions λ and a profile of mixed strategies σ into itself. The correspondence returns the partitions that minimize the out-of-sample prediction error, given the training sample and the strategy profile σ , as well as the mixed strategies that are compatible with the distribution over partitions λ and the strategy profile σ . A fixed point of Ψ corresponds to a CVE.

Next, we consider the convergence properties of CVE as the training sample becomes perfect. We first recall the definition of an Analogy Based Expectations Equilibrium, adapted to the notation in our setting.²

Definition 3. *ABEE, Jehiel (2005)*

σ is an Analogy Based Expectations Equilibrium (ABEE) with respect to the partitions $\{\Pi_i\}$ if for each player i and type t_i , for all a_i with $\sigma_i(a_i|t_i) > 0$,

$$u_i \left(a_i, t_i, \sum_{t'_i \in \pi(t_i)} \mu(t'_i | t'_i \in \pi(t_i)) \bar{x}_i^\sigma(t'_i) \right) \geq u_i \left(a'_i, t_i, \sum_{t'_i \in \pi(t_i)} \mu(t'_i | t'_i \in \pi(t_i)) \bar{x}_i^\sigma(t'_i) \right)$$

holds for all $a'_i \in A_i$.

Say that player i 's set of feasible partitions \mathcal{P}_i contains a *finest partition* if there is $\Pi \in \mathcal{P}_i$ such that Π refines each $\Pi' \in \mathcal{P}_i$. Denote the finest partition by Π_i^{finest} .

For each n , let v^n be a distribution that has finite support and is symmetric around 0. Assume that $\text{Var}_{\varepsilon \sim v^n}(\varepsilon) \rightarrow 0$ as $n \rightarrow \infty$. Let (ε^n) be a sequence of independent random variables where $\varepsilon^n \sim v^n$. Let σ^n be a CVE when the training sample is distributed according to v^n .

Proposition 2. *Assume the game is finite. Assume each player's set of feasible partitions contains a finest partition. If the limit $\sigma = \lim_{n \rightarrow \infty} \sigma^n$ exists, then σ is an ABEE with respect to each player's finest partition Π_i^{finest} .*

Proposition 2 states that an accumulation point of CVEs as the noise of the training sample vanishes is an ABEE with respect to the finest partition. Consequently, if each

²Note that we use linearity of u_i in x here.

player's finest partition is maximally fine (such that each type is in different cell), then CVE converges to a Nash equilibrium when the training sample noise vanishes.

The key step in the proof of Proposition 2 is to show that each player's belief about the sufficient statistic converges to the belief a player holds when (i) she adopts the finest partition, and (ii) there is no noise in the belief formation process. While intuitive, this property is not a priori obvious as, for any non-degenerate training sample noise, the partition chosen in a CVE need not coincide with the finest partition. Moreover, the partitions chosen by each player need not converge to the finest partition. In fact, such a convergence fails, in general.³ Nonetheless, beliefs converge; see Lemma 1.

Next, we make an observation regarding which partitions are chosen in CVE. The next Proposition states that the chosen partition does not condition on irrelevant variables.

Proposition 3. *Let Π and Π' be two feasible partitions for player i , and assume that Π' refines Π . Suppose that σ_{-i} is such that $\bar{x}_i^\sigma(\cdot)$ is measurable with respect to Π .⁴ Then, the expected out-of-sample prediction error from partition Π is weakly lower than from partition Π' for any realization of the training sample noise ε_i .*

Corollary 1. *Assume that for each player i , \mathcal{P}_i includes the coarsest partition $\Pi_i^{\text{coarse}} = \{T_i\}$.*

If types are independently distributed across players and players have private values (i.e., $f(\theta, a_{-i})$ does not depend on θ), then every CVE is equivalent to a CVE in which each player is forced to choose the coarsest partition.

The intuition behind Proposition 3 is the following. When the expected sufficient statistic \bar{x}_i^σ is measurable with respect to Π , the prediction error arises solely from noise in the training sample. It is then never strictly optimal to choose a finer partition Π' , as this would result in a higher variance by de-pooling the noise in the training sample. The proof is a straightforward application of Jensen's inequality.

Moreover, the following is also true: if the finer partition Π' is chosen, the beliefs it generates are the same as those generated by the partition Π . We remark that the finer partition Π' is chosen with probability 0 if the training sample noise has an atomless distribution.

³See Corollary 1.

⁴I.e., $\bar{x}_i^\sigma(\cdot)$ is constant on each cell of Π .

Proposition 3 does *not* imply that the out-of-sample-fit criterion selects a partition Π such that the sufficient statistic $\bar{x}_i^\sigma(t_i)$ is measurable with respect to Π . To put it differently: even if $\bar{x}_i^\sigma(t_i) = \bar{x}_i^\sigma(t'_i)$, the types t_i and t'_i need not belong to the same cell. A counterexample is the following. Suppose that there are three types that occur with probability $1/3$. The vector of sufficient statistics is given by $\bar{x}_i^\sigma = (0, 0, 1)$ and the realized training sample noise is $\varepsilon_i = (0, -1, 1)$. Then, the partition that uniquely minimizes the out-of-sample prediction error, $\Pi = \{\{t^1\}, \{t^2, t^3\}\}$, puts types t^1 and t^2 in different cells. Intuitively, the prediction error can be reduced by introducing some bias if this pools extreme realizations of the training sample noise.

We next make an observation that focuses on the partitions $\Pi_i(\varepsilon_i)$ that are part of a compatible response to σ_{-i} and ε_i . While the training sample noise has mean zero, it is not obvious a priori that the expected noise conditional on a particular selected partition is also zero. For example, in Eliaz and Spiegler (2019), the crucial behavioral effects arise precisely because the noise distribution conditional on model selection can be biased. However, the following Proposition establishes that thanks to the validation sample noise and the validation criterion, this type of bias does not exist in our model.

Proposition 4. *For every $t_i \in T_i$ and any partition $\Pi_i \in \mathcal{P}_i$, the expectation of ε_{t_i} conditional on the event that Π_i satisfies (6) is zero.*

We emphasize that this result is due to the partition being selected through cross-validation, via *expected* out-of-sample predictive loss with a quadratic loss-function. For other loss functions, the result would fail. Moreover, the assumption that the training-sample noise is symmetrically distributed around 0 is essential for this result.

We conclude this section with an observation regarding partition choice, when only the maximally fine and the maximally coarse partitions are feasible. Let $\text{Var}(\varepsilon_i)$ be the variance of the realized training sample, i.e., the variance of the random variable that takes the value ε_{t_i} with probability $\mu(t_i)$. Similarly, let $\text{Var}(\bar{x}_i^\sigma)$ be the variance of the average sufficient statistic.

Proposition 5. *Assume the only feasible partitions are the maximally fine partition $\Pi^{fine} = \{\{t_i\}_{t_i \in T_i}\}$ and the maximally coarse partition $\Pi^{coarse} = \{T_i\}$. Given a realization of the training sample ε_i , the fine partition is adopted if*

$$\text{Var}(\varepsilon_i) < \text{Var}(\bar{x}_i^\sigma) \tag{10}$$

Proposition 5 provides a tractable partition-selection criterion: if the variance of the realized training sample is less than the variance of the sufficient statistic, the fully fine partition is favored over the fully coarse partition.

4. Applications

In this section we apply CVE to a variety of games, and illustrate the qualitative effects that the equilibrium concept generates.

4.1. Quadratic-Payoff Games

When players' action set is the set of real numbers and their payoffs are quadratic in their own action, CVE becomes especially tractable, because each player i 's best-reply becomes linear in her point estimate of x_i .

Consider a two-person game with complete information and an arbitrarily large (finite) set of states, in which players' action set is \mathbb{R} , $t_1 = t_2 = \theta$, and the payoff function takes the quadratic form:

$$u_i(t_i, a_i, x_i) = a_i(\gamma\theta + \beta x_i) - \frac{1}{2}a_i^2 \quad (11)$$

where $x_i = a_j$. Normalize $Var(\theta) = 1$. The only feasible partitions are the fully coarse and fully fine ones. Finally, assume $\varepsilon_\theta \sim N(0, \sigma^2)$.

Clearly, the symmetry of the game and the linearity of players' best-reply functions implies that the game's unique Nash equilibrium is symmetric.

We proceed to characterize CVE. First, when the mean of player i 's belief over x_i in state θ is $\hat{x}_i(\theta)$, her best-reply is

$$a_i = \gamma\theta + \beta\hat{x}_i(\theta) \quad (12)$$

Consider some partition that player i 's ML agent adopts, such that $\pi(\theta)$ is the partition cell that includes θ . By definition,

$$\hat{x}_i(\theta) = \sum_{\theta' \in \pi(\theta)} \mu(\theta')(\bar{x}_i(\theta') + \varepsilon_{\theta'}).$$

By Proposition 4, the expectation of $\varepsilon_{\theta'}$ conditional on $\theta' \in \pi(\theta)$ is *zero*, regardless of the partition that was selected.

In CVE, the expectation of player i 's best-reply in θ (which itself is given by (12))

coincides with $\bar{x}_j(\theta)$. Denote

$$x_i^* = \sum_{\theta'} \mu(\theta') \bar{x}_i(\theta') \quad \theta^* = \sum_{\theta'} \mu(\theta') \theta'$$

Let λ_i be the equilibrium probability with which player i 's ML agent adopts the fine partition. Then, since $\bar{x}_i(\theta)$ coincides with the expectation of a_j in state θ , we can use (12) to obtain

$$\bar{x}_i(\theta) = \gamma\theta + \beta [\lambda_j \bar{x}_j(\theta) + (1 - \lambda_j)x_j^*].$$

Taking expectations of both sides of this equation w.r.t μ for both $i = 1, 2$, we obtain $x_i^* = \gamma\theta^*/(1 - \beta)$. It follows that

$$\bar{x}_i(\theta) = \frac{\gamma(1 + \lambda_j\beta)}{1 - \lambda_i\lambda_j\beta^2}\theta + \kappa\theta^*, \quad (13)$$

where κ is a coefficient that involves β , λ_i , and λ_j .

Using the normalization $Var(\theta) = 1$, we obtain

$$Var(\bar{x}_i) = \left(\frac{\gamma(1 + \lambda_j\beta)}{1 - \lambda_i\lambda_j\beta^2} \right)^2 \quad (14)$$

By Proposition 5, $\lambda_i = \mathbb{P}\{Var(\varepsilon) \leq Var(\bar{x}_i)\}$. Plugging (14) into this equation gives us a pair of equations for λ_1 and λ_2 .

When we restrict attention to symmetric CVE, the characterization is reduced to a single equation for λ :⁵

$$\lambda_i = \mathbb{P} \left\{ Var(\varepsilon) \leq \left(\frac{\gamma}{1 - \beta\lambda} \right)^2 \right\} \quad (15)$$

Let us now apply this characterization to two special cases of interest.

Strategic substitutability

Suppose $\beta < 0$. In this case, the game exhibits *strategic substitutability*. The standard Cournot competition with linear demand (and allowing prices to get negative values) is a familiar instance. In this case, equation (15) has a *unique interior solution*. The reason is that the R.H.S. is continuously decreasing in λ when $\beta < 0$; the R.H.S. is above the L.H.S. at $\lambda = 0$; and the R.H.S. is below the L.H.S. at $\lambda = 1$.

Moreover, when the variance σ^2 of the training sample noise goes up, the R.H.S. of

⁵When θ is uniformly distributed, the R.H.S. of this equation can be conveniently written in terms of the *cdf* of a χ^2 distribution with $|\Theta|$ degrees of freedom.

(15) decreases, and therefore so does the solution λ . This means that the coarse partition is more likely to be adopted in symmetric CVE when the training sample is noisier. As a result, while the ex-ante expected equilibrium action x^* is invariant to σ^2 , its *volatility* $Var(\bar{x})$ increases with σ^2 when $\beta < 0$, because $Var(\bar{x})$ is decreasing in λ — which itself is decreasing in σ^2 .

Thus, when the game exhibits strategic substitutability, symmetric-equilibrium behavior is more volatile relative to Nash equilibrium. In the Cournot-competition case, this means that equilibrium prices are more volatile than in Nash equilibrium. The intuition is that when a player's belief is based on the coarse partition, she underestimates the opponent's reaction to the fluctuations in θ ; strategic substitutability then implies that the player's own reaction to θ is amplified. This effect increases σ^2 , through its effect on the equilibrium probability of adopting the coarse partition.

Unlike Nash equilibrium, for some parameter values there exist *asymmetric* CVE. For example, for $\beta = -0.85$ and $\gamma = 1.5\sigma$, (14) has the asymmetric solutions (0.473, 0.938) and (0.938, 0.473) for (λ_1, λ_2) , in addition to the unique symmetric solution $\lambda = 0.794$. The intuition behind the possibility of asymmetric CVE is as follows. Suppose player 1's strategy has high volatility while player 2's strategy has low volatility. Then, player 2 is likely to adopt the fine partition while player 1 is likely to adopt the coarse partition. As a result, player 2's mean belief is relatively responsive to the state while player 1's belief is relatively unresponsive. Strategic substitutability then implies that player 2's best-reply is relatively unresponsive to the state while player 1's best-reply is relatively responsive. This is consistent with the initial guess regarding the volatility of the two players' behavior.

Team effort

We now examine an example of a game that exhibits strategic complementarities. Consider a team of two agents whose effort choices affect a project's outcome. If the project is successful, each agent receives a payoff of 1, whereas an unsuccessful project yields a payoff of 0. Agents choose effort $a_i \in \mathbb{R}$ simultaneously after commonly observing a uniformly distributed state of Nature $\theta \in \{\theta_1, \theta_2\}$. Let

$$u_i(a_i, a_j, \theta) = \delta\theta(a_i + a_j) + (1 - \delta)a_i a_j - \frac{1}{2}a_i^2.$$

The first two terms represent the probability that the project is successful, where $\delta \in (0, 1)$ measures the relative importance of complementarities in effort. The third term represents the cost of effort. The unique Nash equilibrium in state θ is $a_1 = a_2 = \theta$,

yielding a common ex-ante payoff of $\frac{1}{2}(\frac{1}{2} + \delta)(\theta_1^2 + \theta_2^2)$.

We now turn to symmetric CVE. Let $X = \mathbb{R}$, and identify x_i with a_j — i.e., each agent’s ML agent aims to predict the opponent’s action. Note that the expression for u_i does not literally fit into (11), because it contains an additional term that is linear in a_j . However, this term matters for welfare analysis and interpretation only; it is behaviorally irrelevant, and therefore we can apply our previous characterization.

Because there are only two states, $Var(\varepsilon) \leq Var(\bar{x})$ if and only if $(\Delta\varepsilon)^2 < (\Delta\bar{x}_i)^2$, where $\Delta\bar{x} = \bar{x}(\theta_2) - \bar{x}(\theta_1)$ and $\Delta\varepsilon = \varepsilon_{\theta_1} - \varepsilon_{\theta_2}$. Equation (13) becomes

$$\bar{x}(\theta) = \delta\theta + (1 - \delta) \left[\lambda\bar{x}(\theta) + (1 - \lambda) \frac{\bar{x}(\theta_1) + \bar{x}(\theta_2)}{2} \right]$$

such that

$$\Delta\bar{x} = \frac{\delta\Delta(\theta)}{1 - \lambda(1 - \delta)}$$

Consequently, equation (15) becomes

$$\lambda = \mathbb{P} \left(|\Delta\varepsilon| < \frac{\delta\Delta\theta}{1 - \lambda(1 - \delta)} \right)$$

Using the fact that $\Delta\varepsilon \sim N(0, 2\sigma_\varepsilon^2)$, we can rewrite this equation as follows:

$$\lambda = 2\Phi \left(\frac{\delta\Delta\theta}{\sqrt{2}\sigma_\varepsilon[1 - \lambda(1 - \delta)]} \right) - 1 \quad (16)$$

where $\Phi(\cdot)$ is the *cdf* of the standard normal distribution.

A solution to (16) always exists. However, in general, the solution is not unique; see Figure 1. The reason such multiplicity can arise is that the R.H.S. of (16) is increasing in λ . The intuition behind this feature is that if ML agents adopt the fine partition more frequently, team members’ equilibrium beliefs tend to become less coarse and more volatile. The team effort game’s strategic complementarities then imply that agents’ best-reply also becomes more volatile, which implies that the fine partition is more likely to prevail.

Let us explore (16) from additional angles. First, fix δ . When $\sigma_\varepsilon^2 \rightarrow \infty$, $\lambda \rightarrow 0$, such that $\Delta\bar{x} \rightarrow \delta\Delta\theta$. Conversely, when $\sigma_\varepsilon^2 \rightarrow 0$, $\lambda \rightarrow 1$, such that $\Delta\bar{x} \rightarrow \Delta\theta$. Second, fix $\sigma_\varepsilon > 0$. As $\delta \rightarrow 0$, $\lambda \rightarrow 0$ — i.e., agents’ ML agents almost always adopt the coarse partition. And since δ is small, agents’ equilibrium behavior is arbitrarily rigid — i.e., $x(\theta) \approx (\theta_1 + \theta_2)/2$ for every θ — and when σ_ε is small, their equilibrium payoff in each




Figure 1: This figure plots $f(\lambda) = 2\Phi\left(\frac{\delta(\theta_2 - \theta_1)}{\sqrt{2\sigma_\varepsilon[1 - \lambda(1 - \delta)]}}\right) - \lambda - 1$ for different values of δ and $\Delta = \frac{\theta_2 - \theta_1}{\sqrt{2\sigma_\varepsilon}}$. A root of $f(\lambda)$ corresponds to a solution to Equation (16). Note that there are three roots in $(0, 1)$ for each parameter specification.

state is approximately $(\theta_1 + \theta_2)^2/8$. Compare this with the Nash equilibrium payoff $(\theta_1 + \theta_2)^2/4$ in the same limit. In this sense, CVE induces a substantial deviation from Nash equilibrium for any training sample noise variance.⁶ In general, agents' expected payoff in CVE is below the Nash equilibrium level.

4.2. Unanimity Voting

The following is a pedagogically convenient specification of the familiar jury model (Feddersen and Pesendorfer (1998)). There are two players who vote on a reform. Voting

⁶The order of limits clearly matters here. For any given δ , if we send σ_ε^2 to zero, agents almost always select the fine partition and their equilibrium behavior converges to the Nash benchmark.

is unanimous so that the reform is adopted if and only if both players vote for it. The voters have aligned preferences and receive a state-dependent payoff from the reform: this payoff is 1 in state θ_1 and $-c < 0$ in state θ_2 . The payoff is 0 in both states if the reform is not adopted. The two states are equally likely.

Before voting, players receive a signal $t_i \in \{g, b\}$ that is informative of the state. Signals are independent across voters conditional on the state. Assume $\mathbb{P}[t_i = g | \theta = \theta_1] = \mathbb{P}[t_i = b | \theta = \theta_2] = q \in (1/2, 1)$.

Denote voting for the reform by Y and voting against by N . The sufficient statistic of player i is $f(\theta_1, a_j) = \mathbb{1}\{a_j = Y\}$ and $f(\theta_2, a_j) = -c\mathbb{1}\{a_j = Y\}$. Player i 's payoff function is then

$$u_i(t_i, Y, x) = x \quad u_i(t_i, N, x) = 0.$$

Thus, the sufficient statistic is the payoff from voting for reform.

Denote by $\sigma_i(t_i)$ the probability that a player votes Y after signal t_i . Then, the expected sufficient statistic is

$$\begin{aligned} \bar{x}(g) &= (q^2 - c(1-q)^2)\sigma_j(g) + q(1-q)(1-c)\sigma_j(b) \\ \bar{x}(b) &= (q^2 - c(1-q)^2)\sigma_j(b) + q(1-q)(1-c)\sigma_j(g) \end{aligned}$$

Assume that $\frac{1-q}{q} < c < 1$. Under this parametric restriction, there are two symmetric Nash equilibria: $\sigma(g) = \sigma(b) = 0$ and $\sigma(g) = 1, \sigma(b) = q(1-q)(1-c)/(cq^2 - (1-q)^2)$.

Let us explore symmetric CVE, under the assumption that the training sample noise is distributed uniformly on $\{-d, d\}$.

First, we argue that $\sigma(g) = \sigma(b) = 0$ is not a CVE for any positive level of noise. To see this, observe that a player votes in favour of the reform if her belief about the sufficient statistic is positive. For $\sigma(g) = \sigma(b) = 0$, $\bar{x}(g) = \bar{x}(b) = 0$. Consequently, a player votes Y if the noise realization is positive. This rules out the uninformative, always-vote-against Nash equilibrium.

Second, we argue that the probability of voting Y can be *higher* in CVE than in the efficient Nash equilibrium. For ease of exposition, assume $q = 2/3$ and $c = 5/8$. This parameter selection implies that in the efficient Nash equilibrium, $\sigma(b) = 1/2$. Let $d = 0.22$. The following is the unique symmetric CVE: $\sigma(g) = 1, \sigma(b) = 8d - 7/6$. Note that $1/2 < \sigma(b) < 3/4$. Given these strategies, $\bar{x}(g) - \bar{x}(b) = 2d$. Thus, the fine partition is chosen when

$$(\varepsilon_g, \varepsilon_b) \in \{(d, d), (-d, -d), (d, -d)\},$$

and the coarse partition is chosen when $(\varepsilon_g, \varepsilon_b) = (d, -d)$. At the realisation $(\varepsilon_g, \varepsilon_b) = (-d, d)$, the fine and the coarse partition achieve the same out-of-sample prediction error. The coarse partition is chosen with probability $32d - 20/3$. Players' best-responses are to vote Y after the g signal, irrespective of the chosen partition. When the fine partition is chosen, a player votes Y after the b signal if and only if $\varepsilon_b = d$; for the noise realization $\varepsilon_b = -d$, the player votes N after the b signal when the fine partition is chosen. When the coarse partition is chosen, a player votes Y after the b signal as well.

Thus, although one could expect that the training sample noise would make it harder for voters to coordinate on unanimous support for reform, the endogenous model selection ends up pushing them to vote for reform more frequently than in the efficient Nash equilibrium. Indeed, endogenous model selection plays a key role in this effect. If players were forced to use the fine partition, the unique CVE would be $\sigma(g) = 1, \sigma(b) = 1/2$. In contrast, in the CVE we constructed, players vote Y when the coarse partition is selected.

4.3. Speculative Bets

Consider a two player speculative betting game. There are two commonly known states of the world, $\theta_1 = 1$ and $\theta_2 = -1$. Both states are ex-ante equally likely. Player i can either accept the bet $a_i = A$ or decline $a_i = D$. The bet is executed if and only if both players accept. The payout to player 1 if the bet is accepted equals the state θ ; player 2 gets the negative.

The sufficient statistic is the probability x_i with which player j accepts the bet, i.e., $x_i = \mathbb{1}\{a_j = A\}$. That is, x_i is the position player j takes in the bet. We do not restrict $x_i \in [0, 1]$ so as to allow players to adopt short and long positions.

The payoffs then are $u_1(\theta, A, x_1) = \theta x_1$ and $u_1(\theta, D, x_1) = 0$ and $u_2(\theta, A, x_2) = -\theta x_2$ and $u_2(\theta, D, x_2) = 0$.

Player 1 has access only to the fine partition $\Pi^f = \{\{\theta_1\}, \{\theta_2\}\}$. Player 2 has access to the fine partition and the coarse partition $\Pi^c = \{\{\theta_1, \theta_2\}\}$. Assume the training sample noise ε_θ is distributed uniformly on $\{-d, d\}$ for a $d > 0$.

When the noise d is between $1/2$ and $3/4$, the unique CVE has

$$\begin{aligned}\sigma_1(A|\theta_1) &= \frac{1}{2}\sigma_1(A|\theta_2) = 0 \\ \sigma_1(A|\theta_2) &= \frac{1}{4}\sigma_1(A|\theta_2) = \frac{3}{4}.\end{aligned}$$

The corresponding payoffs for player 1 are $u_1(\theta_1) = 1/8$ and $u_1(\theta_2) = 0$, and player 2 gets the negative. Player 1 thus receives a higher payoff than player 2 and a higher payoff than in Nash equilibrium.

We now argue that the proposed strategies constitute a CVE.⁷ Since Player 1 always uses the fine partition, she plays A in state θ_1 iff $\sigma_2(A | \theta_1) + \varepsilon_{\theta_1} \geq 0$, and plays A in state θ_2 iff $\sigma_2(A | \theta_2) + \varepsilon_{\theta_2} \leq 0$. Since $1/2 < d < 3/4$, this implies $\sigma_1(A|\theta_1) = 1/2$ and $\sigma_1(A|\theta_2) = 0$.

With two equally likely states, player 2's fine partition is selected over the coarse partition iff $|\varepsilon_{\theta_1} - \varepsilon_{\theta_2}| < |\sigma_1(A | \theta_1) - \sigma_1(A | \theta_2)|$. Since $d > \frac{1}{2}$, then the coarse partition is selected whenever the two noise realisations have opposite signs. Player 2 selects the fine partition exactly when $\varepsilon_{\theta_1} = \varepsilon_{\theta_2}$, and selects the coarse partition when $\varepsilon_{\theta_1} \neq \varepsilon_{\theta_2}$. When the noise realisation is $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (d, d)$, the fine partition is selected, and player 2 plays D in state θ_1 and A in state θ_2 . When the noise realisation is $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (-d, -d)$, the fine partition is selected, and player 2 plays A in state θ_1 and D in state θ_2 . For the realisation $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (d, -d)$, the coarse partition is selected. The pooled estimate is $\frac{1}{2}(\frac{1}{2} + d - d) = \frac{1}{4} > 0$, hence Player 2 plays D in state θ_1 and A in state θ_2 . The same holds for the realisation $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (-d, d)$. Thus Player 2 plays A in state θ_1 only in the realisation $(-d, -d)$, which has probability $1/4$.

When the noise satisfies $d > \frac{3}{4}$, the unique CVE has

$$\begin{aligned}\sigma_1(A|\theta_1) &= \frac{1}{2}\sigma_1(A|\theta_2) = \frac{1}{2} \\ \sigma_1(A|\theta_2) &= \frac{1}{4}\sigma_1(A|\theta_2) = \frac{3}{4}.\end{aligned}$$

The corresponding payoffs for player 1 are $u_1(\theta_1) = 1/8$ and $u_1(\theta_2) = -3/8$, and player 2 gets the negative. In expectation over the state, player 1 thus receives a strictly lower payoff than player 2 and a lower payoff than in Nash equilibrium.

Player 1's estimate in state θ_1 is $\frac{1}{4} + \varepsilon_{\theta_1}$, hence Player 1 plays A in state θ_1 exactly when $\varepsilon_{\theta_1} = d$. Therefore $\sigma_1(A | \theta_1) = \frac{1}{2}$. In state θ_2 , Player 1's estimate is $\frac{3}{4} + \varepsilon_{\theta_2}$, so player 1 plays A in state θ_2 exactly when $\varepsilon_{\theta_2} = -d$.

Player 2's statistic is constant across states. Thus the fine partition is never strictly better than the coarse partition. If the noise realisations have opposite signs, the coarse partition is strictly better. If they have the same sign, the two partitions generate the same estimate, so either partition is optimal.

Again consider the four equally likely noise realisations. For the realisation $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) =$

⁷The proof of uniqueness is straightforward but tedious.

(d, d) , player 2 plays D in state θ_1 and A in state θ_2 . For the realisation $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (-d, -d)$, player 2 plays A in state θ_1 and D in state θ_2 .

For the realisations $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (d, -d)$ and $(\varepsilon_{\theta_1}, \varepsilon_{\theta_2}) = (-d, d)$, the coarse partition is selected. The average sufficient statistic is

$$\frac{1}{2} \left(\frac{1}{2} + d + \frac{1}{2} - d \right) = \frac{1}{2} > 0.$$

Thus Player 2 plays D in state θ_1 and A in state θ_2 . Thus Player 2 plays A in state θ_1 only in the realisation $(-d, -d)$, which has probability $1/4$.

The example shows that access only to the correct fine partition need not benefit a player, even in a zero-sum setting. At high noise levels, player 1 obtains a lower payoff than player 2. Player 2's ability to choose the coarse partition therefore raises her payoff in this case. The intuition is as follows. Player i accepts the bet in the state that is unfavourable to her, θ_j , only if she believes that the other player is taking a short position. When beliefs are very noisy, this occurs with positive probability and leads to a negative payoff. By contrast, when player 2 chooses the coarse partition—which happens when the noise realizations have opposite signs in the two states—the coarse partition averages out the noise. Her belief is then correct, and she chooses the weakly dominant action in each state. The example shows that this effect can be strong enough to generate non-zero payoffs.

4.4. An Illusion of Control

We use the final example in this section to explore a natural extension of the framework, and a notable effect that this extension generates.

Consider a binary-action, two-person game with a commonly observed binary state. Let $a_i \in \{0, 1\}$. Identify x_i with the probability that $a_j = 1$. Let $u_i(\theta, a_i, x_i) = 2\theta x_i - a_i$, where $\theta \sim U\{0, 1\}$. Under rationality, $a_i = 0$ is strictly dominant.

We extend the main model by allowing players to form beliefs that are based on partitions of $\Theta \times A$ — whereas in the main model, only partitions of Θ are allowed. Specifically, we allow for three possible partitions of all four pairs (θ, a) : The maximally fine partition, and two coarse partitions: $\{(0, 0), (0, 1)\}, \{(1, 0), (1, 1)\}$ and $\{(0, 0), (1, 0)\}, \{(0, 1), (1, 1)\}$. The former coarse partition is correct, because player j 's action is purely a function of θ . The latter coarse partition is wrong, because it falsely attributes a_j to a_i . By Proposition 3, we can ignore the maximally fine partition because it is a refinement of the correct coarse partition. Therefore, in symmetric CVE,

players' ML agents will only mix between the two coarse partition.

Since players' ML agents may attribute outcomes to both states and actions, player i 's training sample noise is associated with pairs (θ, a_i) . Suppose $\varepsilon_{\theta a} \sim U\{-\frac{1}{2}, \frac{1}{2}\}$.

In CVE, $a_i = 0$ with certainty when $\theta = 0$, independently of the partition player i 's ML agent adopts. Let α denote the probability that $a_i = 1$ at $\theta = 1$ in CVE. When player i 's ML agent adopts the correct coarse partition, she chooses the rational action $a_i = 0$, because she does not think that a_i affects a_j . It follows that the player can only play $a_i = 1$ with positive probability when $\theta = 1$, provided that her ML agent adopts the wrong coarse partition.

Let us derive the equilibrium partition selection rule. The expected out-of-sample prediction error induced by the correct coarse partition is

$$L_\theta = \frac{1}{2} \left(\frac{\varepsilon_{11} + \varepsilon_{10}}{2} \right)^2 + \frac{1}{2} \left(\frac{\varepsilon_{01} + \varepsilon_{00}}{2} \right)^2$$

To derive the expected out-of-sample prediction error induced by the wrong coarse partition, note first that $a_i = 1$ with probability $\frac{1}{2}\alpha$ and $a_i = 0$ with probability $1 - \frac{1}{2}\alpha$. Let us now calculate the joint equilibrium distribution over θ and x_i conditional on any given a_i :

$$\mathbb{P}(\theta, x_i \mid a_i = 1) = \begin{cases} \alpha & \text{for } (\theta, x_i) = (1, 1) \\ 1 - \alpha & \text{for } (\theta, x_i) = (1, 0) \end{cases}$$

and

$$\mathbb{P}(\theta, x_i \mid a_i = 0) = \begin{cases} \frac{\frac{1}{2}(1-\alpha)}{1-\frac{1}{2}\alpha} \alpha & \text{for } (\theta, x_i) = (1, 1) \\ \frac{\frac{1}{2}(1-\alpha)}{1-\frac{1}{2}\alpha} (1 - \alpha) & \text{for } (\theta, x_i) = (1, 0) \\ \frac{\frac{1}{2}}{1-\frac{1}{2}\alpha} & \text{for } (\theta, x_i) = (0, 0) \end{cases}$$

The expected out-of-sample prediction error of the wrong coarse partition is thus

$$L_a = \frac{\alpha}{2} \varepsilon_{11}^2 + \frac{1-\alpha}{2} \left(\varepsilon_{10} - \frac{\alpha(1-\alpha)}{2-\alpha} \right)^2 + \frac{1}{2} \left(\varepsilon_{00} - \frac{\alpha(1-\alpha)}{2-\alpha} \right)^2$$

We can now calculate α : It is the probability that the following inequalities hold jointly: $L_a < L_\theta$ and $2(\alpha + \varepsilon_{11}) - 1 > 0$. The logic behind the latter inequality is that conditional on $a_i = 1$, $\theta = 1$ and $x_i = \alpha$; and so when player i adopts the wrong coarse partition, her estimate of x_i at $a = 1$ is $\alpha + \varepsilon_{11}$.

It can be checked that the configuration $\varepsilon_{\theta a} = \frac{1}{2}$ for all (θa) (which occurs with probability $\frac{1}{16}$) together with $\alpha = 1/16$ is the only one that satisfies both inequalities.

Therefore, this specification constitutes a CVE. We have thus found a CVE in which players choose the dominated action $a = 1$ with positive probability. This is an “illusion of control” effect: Although a_i has no objective causal effect on a_i , players sometimes choose the costly action $a = 1$ because the model their ML agent adopts sometimes attributes a successful outcome to this action.

5. Related Literature

Our main contribution is to the literature on how ML plays out in strategic interactions, where a key feature is that the training data is endogenous. While we incorporate ML into belief equilibrium formation and assume that decisions are made by humans in response to ML-generated beliefs, most of the literature in economics focuses on reinforcement learning. Calvano et al. (2020), Hansen, Misra, and Pai (2021), Asker, Fershtman, and Pakes (2024), among many others, numerically study reinforcement learning in oligopoly games. Banchio and Mantegazza (2024), Dolgoplov (2024), Possnig (2024), and Cartea et al. (2025) analytically characterize long-run outcomes in such environments. Waizmann (2025) studied an interaction between a long-run player who obeys Q-learning and a sequence of rational short-run players.

More recently, Jehiel and Weber (2025), Jehiel and Mohlin (2026) and Spiegler (2026) focused on ML-based predictive models that serve human players in games. All three use partitions of a set of contingencies to conceptualize predictive models. Moreover, in contrast to our model, these three papers assume that the data which determines model selection is noiseless. In Jehiel and Weber (2025), the partition size is fixed, and the equilibrium belief formation model obeys a form of k-means clustering. Their *clustered distributional analogy-based expectations equilibrium* obtains as a special case of our cross-validation equilibrium for a particular choice of feasible models and noiseless training samples.

The model of Jehiel and Mohlin (2026), applied to static Bayesian games, differs from ours in two important aspects. First, they impose an exogenous feedback constraint, which implies that certain situations cannot be distinguished. Second, their model-selection criterion relies on an exogenous notion of similarity between situations. Neither feature is present in our framework. Under perfect feedback, their categorization equilibrium coincides with Nash equilibrium. In Spiegler (2026), the equilibrium partition trades off complexity and expected predictive accuracy: a player incurs a physical cost that increases in the number of cells of the chosen partitions. In contrast, in the

present model, there is no explicit cost of complex models. Rather, simple models can sometimes prevail because they improve out-of-sample predictive accuracy, sacrificing bias for lowering variance.

At a broader level, this paper contributes to the literature on game-theoretic equilibrium concepts that incorporate non-rational belief formation; see Jehiel (2005), Eyster and Rabin (2005), Esponda (2008), Esponda and Pouzo (2016), Spiegel (2016), and **clyde2025proxy**. In each of these cases, equilibrium behavior involves subjective optimization against a belief that arises from fitting (in some sense) a subjective model to the equilibrium distribution.

A key feature that distinguishes CVE from these works is that agents' subjective models in CVE are randomly and endogenously determined, following a procedure that involves taking samples from the equilibrium distribution. In this respect, CVE is related to equilibrium concepts in which agents best-reply to endogenously drawn samples; see Osborne and Rubinstein (1998), Salant and Cherry (2020), and Danenberg and Spiegel (2025). These papers, however, feature no element of ML-based model selection.

A. Proofs

In the proofs, it will be convenient to let $h_i(t_i, \varepsilon_i, \Pi_i) \in \Delta(A_i)$ denote a player i 's mixed strategy when her type is t_i , the realized training sample is ε_i and her ML agent adopts the partition Π_i . Moreover, we let $\lambda_i(\varepsilon_i, \Pi_i)$ denote the probability with which player i 's ML agent adopts the partition Π_i given ε_i . Clearly, given σ_{-i} , if $\lambda_i(\Pi, \varepsilon_i)$ is supported on partitions satisfying (6) and $h_i(t_i, \varepsilon_i, \Pi)$ is supported on actions satisfying (7), one can find a distribution β_{ε_i} over pairs (Π_i, g_i) that are compatible with σ_{-i} and ε_i according to Definition 1, and vice versa.

A.1. Proof of Proposition 1

For each player i and vector of realized type-dependent shocks $\varepsilon_i = (\varepsilon_{t_i})$, let $\lambda_i(\varepsilon_i) \in \Delta\mathcal{P}_i$ be a distribution over the set of feasible models. Denote $\lambda_i = \lambda_i(\varepsilon_i)_{\varepsilon_i}$ and $\lambda = (\lambda_i)_i$. Let $\sigma_i(\cdot | t_i) \in \Delta A_i$ be a mixed strategy for type t_i , $\sigma_i = (\sigma_i(\cdot | t_i))_{t_i}$ and $\sigma = (\sigma_i)_i$.

Construct a correspondence

$$\Psi : \prod_{i=1}^N (\Delta \mathcal{P}_i)^{|\text{supp} \varepsilon| \times |T_i|} \times \prod_{i=1}^N (\Delta A_i)^{|T_i|} \rightrightarrows \prod_{i=1}^N (\Delta \mathcal{P}_i)^{|\text{supp} \varepsilon| \times |T_i|} \times \prod_{i=1}^N (\Delta A_i)^{|T_i|}$$

$$(\lambda, \sigma) \mapsto \Psi(\lambda, \sigma)$$

The goal is to define Ψ such that, if $(\lambda^*, \sigma^*) \in \Psi(\lambda^*, \sigma^*)$ is a fixed point of this correspondence, σ^* is a CVE.

Given σ , for any i and t_i define $p^\sigma(\theta, t_i, a)$ as in Equation (1) and $\bar{x}_i^\sigma(t_i)$ as in Equation (2). $\bar{x}_i^\sigma(t_i)$ is a continuous function of σ , for any fixed t_i . For each $t_i, \varepsilon_i = (\varepsilon_{t_i}) \in \text{supp} \varepsilon_i^{|T_i|}$, and $T_i \supset M_i = \pi(t_i) \in \Pi_i \in \mathcal{P}_i$ define

$$\hat{x}^{\varepsilon_i, \sigma}(M_i) = \sum_{t_i \in M_i} \mu(t_i | t_i \in M) (\bar{x}_i^\sigma(t_i) + \varepsilon_{t_i})$$

$\hat{x}^{\varepsilon_i, \sigma}(M_i)$ is a continuous function of σ .

Denote by $L(\Pi_i, \varepsilon_i)(\sigma)$ the squared prediction error of a partition $\Pi_i \in \mathcal{P}_i$ given a vector of shocks ε_i

$$L_i(\Pi, \varepsilon_i)(\sigma) = \mathbb{E}_\eta \left[\sum_{t_i \in T_i} \mu(t_i) [\bar{x}^{\varepsilon_i, \sigma}(\pi(t_i)) - \bar{x}_i^\sigma(t_i) - \eta_{t_i}]^2 \right].$$

Note that for fixed $\Pi, \varepsilon_i, L(\Pi, \varepsilon_i)(\sigma)$ is a continuous function of σ .

Define $\psi_1(i, \varepsilon_i, \cdot)(\sigma)$ to be

$$\psi_1(i, \varepsilon_i)(\sigma) = \left\{ \lambda'_i(\varepsilon_i) \in \Delta \mathcal{P}_i \left| \lambda'_i(\varepsilon_i)(\Pi_i) > 0 \Rightarrow \Pi_i \in \arg \min_{\tilde{\Pi}_i \in \mathcal{P}_i} L_i(\tilde{\Pi}_i, \varepsilon_i)(\sigma) \right. \right\}.$$

Note that the continuity of $L_i(\Pi, \varepsilon_i)(\cdot)$ and the finiteness of \mathcal{P}_i imply that $\psi_1(i, \varepsilon_i)(\cdot)$ is an upper hemi-continuous, non-empty -, convex - and compact-valued correspondence.

Define $\Psi_1(\lambda, \sigma) = (\psi_1(i, \varepsilon_i)(\sigma))_{\varepsilon_i, i}$. Clearly, $\Psi_1(\lambda, \sigma)$ is an upper hemi-continuous, non-empty-, convex-and compact-valued correspondence.

Now we define $\Psi_2(\lambda, \sigma) \in \prod_i (\Delta A_i)^{|T_i|}$. Fix $\Pi_i \in \mathcal{P}_i, \varepsilon_i \in (\text{supp} \varepsilon)^{|T_i|}$ and t_i . Define $\hat{p}^{\varepsilon_i}(y | M_i)(\sigma)$ as in Equation (4).

Define $H_2(i, t_i, \varepsilon_i, M)(\sigma)$ as

$$\left\{ h_i(t_i, \varepsilon_i, M) \in \Delta(A_i) \mid h_i(t_i, \varepsilon_i, M)(a_i) > 0 \Rightarrow a_i \in \arg \max_{a'_i \in A_i} \sum_{y \in X_i} \hat{p}^{\varepsilon_i}(y \mid M)(\sigma) u_i(a', t_i, y) \right\}$$

Claim 1. For all fixed $i, \varepsilon_i, t_i, M_i, H_2(i, \varepsilon_i, t_i, M_i)(\sigma)$ is a non-empty-, convex- and compact-valued and upper-hemicontinuous correspondence.

Proof. Only the claim that the correspondence is uhc needs proof. Take $\sigma^n \rightarrow \sigma, H(i, t_i, \varepsilon_i, M)(\sigma^n) \ni \gamma^n \rightarrow \gamma$. Need to show that $\gamma \in H(i, t_i, \varepsilon_i, M)(\sigma)$. Note that for each $n, \hat{p}^\varepsilon(y \mid M)(\sigma^n)$ is positive for at most $|M|$ many points $y \in X_i$. Assume wlog that for each $n, \hat{p}^n(y \mid M)C_{0^n}$ is positive for exactly $|M|$ -many points y (if not, pass to subsequences such that $\hat{p}^\varepsilon(y \mid M)(\sigma^n)$ is positive for exactly k -many, $k \leq |M|$, points, and repeat the argument for each k). Note that there exists $|M|$ sequences (x_m^n) such that $\hat{p}^\varepsilon(x_m^n \mid M)(\sigma^n) = \hat{p}^\varepsilon(x_m^{n+1} \mid M)(\sigma^{n+1}) =: p_m$ for all $n \in \mathbb{N}, m = 1, \dots, |M|$. Moreover, as $\sigma^n \rightarrow \sigma, x_m^n \rightarrow x_m$ for each m .

Hence, one can identify γ^n with a probability distribution over

$$\arg \max_{a'_i \in A_i} \sum_{m=1}^{|M|} p_m u_i(a'_i, t_i, x_m^n).$$

Because $u_i(a_i, t_i, x)$ is continuous in x , the limit γ of γ^n puts probability 1 on

$$\arg \max_{a'_i \in A_i} \sum_{m=1}^{|M|} p_m u_i(a'_i, t_i, x_m)$$

which implies $\gamma \in H_2(i, \varepsilon_i, t_i, M_i)(\sigma)$. □

Now define

$$\Psi_2(i, t_i)(\lambda, \sigma) = \left\{ \sigma_i(\cdot | t_i)' \in \Delta(A_i) \mid \exists h_i(t_i, \varepsilon_i, M_i) \in H_2(i, \varepsilon_i, t_i, M_i)(\sigma) \text{ such that, for all } a_i \in A_i, \right. \\ \left. \sigma'_i(a_i | t_i) = \sum_{\varepsilon_i \in \text{supp}(\varepsilon_i) \times |T_i|} v(\varepsilon_i) \sum_{\Pi_i \in \mathcal{P}_i} \lambda_i(\varepsilon_i)(\Pi_i) h_i(t'_i, \varepsilon_i, \pi_i(t'_i))(a_i) \right\}.$$

Claim 2. The correspondence Ψ admits a fixed point $(\lambda^*, \sigma^*) \in \Psi(\lambda^*, \sigma^*)$.

Proof. Ψ is non-empty-, compact- and convex-valued. Moreover, Ψ_1 is upper hemicontinuous. Ψ_2 is upper-hemi-continuous as well: since $H_2(i, \varepsilon_i, t_i, M_i)$ is upper hemicontinuous.

continuous by the previous Claim, $\Psi_2(i, t_i)$ is upper hemi-continuous as well. Consequently, Ψ admits a fixed point by Kakutani's Fixed Point Theorem. \square

Claim 3. *Let $(\lambda^*, \sigma^*) \in \Psi(\lambda^*, \sigma^*)$ be a fixed point of the correspondence Ψ . Then σ^* is a CVE.*

Proof. Let $(\lambda^*, \sigma^*) \in \Psi(\lambda^*, \sigma^*)$. By construction, for every profile (ε_i) , $\lambda_i^*(\varepsilon_i)(\Pi_i) > 0$ only if Π_i solves (6). Moreover, for every i and t_i , $\sigma_i^*(\cdot|t_i)$ satisfies (9) since it is composed of functions $h_i(t_i, \varepsilon_i, M_i) \in H_2(i, \varepsilon_i, t_i, M_i)(\sigma^*)$. \square

A.2. Lemma 1

For a player i let

$$\Gamma(\Pi, \bar{x}, \varepsilon, t_i) = \sum_{t'_i \in \pi(t_i)} \mu(t'_i|t_i \in \pi(t_i)) (\bar{x}_i(t'_i) + \varepsilon_{t'_i})$$

be the belief of player i when adopting partition Π , given the training sample noise ε , the vector of sufficient statistics \bar{x} , and the type t_i .

Lemma 1. *Let $\Pi^n(\varepsilon)$ be the chosen partition when the realization of the training sample noise is ε . For each player i and type t_i ,*

$$\Gamma(\Pi^n(\varepsilon), \bar{x}^{\sigma^n}, \varepsilon, t_i) \xrightarrow[n \rightarrow \infty]{} \Gamma(\Pi^{\min}, \bar{x}^\sigma, 0, t_i)$$

in probability.

Proof. Denote by $L(\Pi, \bar{x}, \varepsilon)$ the expected out-of-sample prediction error for partition Π given the realized training sample ε and a vector of sufficient statistics \bar{x} .

First, for every n and partition $\Pi \in \mathcal{P}_i$,

$$L(\Pi^n(\varepsilon), \bar{x}^{\sigma^n}, \varepsilon) \leq L(\Pi, \bar{x}^{\sigma^n}, \varepsilon) \tag{17}$$

almost surely. Note that for every fixed partition Π , the prediction error $L(\Pi, \bar{x}, \varepsilon)$ is a polynomial of degree two in \bar{x} and ε . Because $\bar{x}^{\sigma^n} \rightarrow \bar{x}^\sigma$ and $\text{Var}(\varepsilon^n) \rightarrow 0$, the RHS of equation (17) converges to $L(\Pi, \bar{x}^\sigma, 0)$ in probability. Consequently,

$$\mathbb{P} \left[L(\Pi^n(\varepsilon), \bar{x}^{\sigma^n}, \varepsilon) \leq \min_{\Pi \in \mathcal{P}_i} L(\Pi, \bar{x}^\sigma, 0) \right] \xrightarrow[n \rightarrow \infty]{} 1.$$

Moreover, since \mathcal{P}_i is finite,

$$\mathbb{P} \left[\Pi^n(\varepsilon) \in \arg \min_{\Pi \in \mathcal{P}_i} L(\Pi, \bar{x}^\sigma, 0) \right] \xrightarrow{n \rightarrow \infty} 1.$$

Next, Lemma 2 and the hypothesis that Π^{finest} refines every $\Pi \in \mathcal{P}_i$ imply that,

$$\Pi^{\text{finest}} \in \arg \min_{\Pi' \in \mathcal{P}_i} L(\Pi', \bar{x}^\sigma, 0).$$

Moreover, if

$$\Pi \in \arg \min_{\Pi' \in \mathcal{P}_i} L(\Pi', \bar{x}^\sigma, 0),$$

then for all types t_i

$$\Gamma(\Pi, \bar{x}^\sigma, 0, t_i) = \Gamma(\Pi^{\text{finest}}, \bar{x}^\sigma, 0, t_i).$$

□

Lemma 2. *Let $X = \{x_i\}$ be a finite set of numbers and $\alpha_i \in (0, 1)$, $\sum_i \alpha_i = 1$. For any $S \subset X$, denote $m(S) = \sum_{i:x_i \in S} \frac{\alpha_i x_i}{\sum_{i:x_i \in S} \alpha_i}$. Then for any $S \neq \emptyset$*

$$\sum_{i:x_i \in S} \frac{\alpha_i}{\sum_{i:x_i \in S} \alpha_i} (x_i - m(S))^2 + \sum_{i:x_i \in X \setminus S} \frac{\alpha_i}{\sum_{i:x_i \in X \setminus S} \alpha_i} (x_i - m(X \setminus S))^2 \leq \sum_{i:x_i \in X} \frac{\alpha_i}{\sum_{i:x_i \in X} \alpha_i} (x_i - m(X))^2$$

with equality only if $m(S) = m(X \setminus S)$.

Proof. The proof is standard and therefore omitted. □

A.3. Proof of Proposition 2

Let σ^n be a CVE for each n . Denote by $\Pi^n(\varepsilon)$ the partition chosen given the realized training sample. Note that $\Pi^n(\varepsilon)$ is a random variable because the training sample ε is random and because the choice of partition is random as well.

Assume toward a contradiction that for some player i and type t_i , $\sigma_i(t_i)(a_i) = 3\beta > 0$ but

$$\begin{aligned} & u_i \left(a_i, t_i, \sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t'_i \in \Pi^{\text{finest}}(t_i)) (\bar{x}_i^\sigma(t'_i)) \right) \\ & \leq u_i \left(a_i^\dagger, t_i, \sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t'_i \in \Pi^{\text{finest}}(t_i)) (\bar{x}_i^\sigma(t'_i)) \right) - 2\delta \end{aligned}$$

for some $a_i^\dagger \in A_i$ and $\delta > 0$.

Let γ be such that $u_i(a_i, t_i, x) \leq u_i(a_i^\dagger, t_i, x) - \delta$ for all x within γ of

$$\sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t_i \in \Pi^{\text{finest}}(t_i)) (\bar{x}_i^\sigma(t'_i)).$$

Denote by B_γ^n the set

$$B_\gamma^n = \left\{ (\Pi, \varepsilon) : \left| \sum_{t'_i \in \pi(t_i)} \mu(t'_i | t'_i \in \pi(t_i)) (\bar{x}_i^{\sigma^n}(t'_i) + \varepsilon_{t'_i}) - \sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t'_i \in \Pi^{\text{finest}}(t_i)) \bar{x}_i^\sigma(t'_i) \right| \leq \gamma \right\}$$

By Lemma 1, for all n large enough,

$$\sum_{(\Pi, \varepsilon) \in B_\gamma^n} v^n(\varepsilon) \lambda_i^n(\Pi, \varepsilon) \geq 1 - \beta.$$

Since $\sigma_i^n \rightarrow \sigma$, it must be that for all n large enough

$$2\beta \leq \sum_{\varepsilon} \sum_{\Pi} v^n(\varepsilon) \lambda^n(\Pi, \varepsilon) h_i^n(t_i, \Pi, \varepsilon)(a_i) \leq \sum_{(\varepsilon, \Pi) \in B_\gamma^n} v^n(\varepsilon) \lambda^n(\Pi, \varepsilon) h_i^n(t_i, \Pi, \varepsilon)(a_i) + \beta$$

The inequality requires that $h_i^n(t_i, \Pi, \varepsilon)(a_i) > 0$ on B_γ^n , contradicting the best-response condition, equation (7) in Definition 1.

We conclude that for each player i and type t_i , $\sigma_i(t_i)(a_i) > 0$ implies

$$u_i \left(a_i, t_i, \sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t'_i \in \Pi^{\text{finest}}(t_i)) (\bar{x}_i^\sigma(t'_i)) \right) \geq u_i \left(a_i^\dagger, t_i, \sum_{t'_i \in \Pi^{\text{finest}}(t_i)} \mu(t'_i | t'_i \in \Pi^{\text{finest}}(t_i)) (\bar{x}_i^\sigma(t'_i)) \right)$$

$\forall a_i^\dagger \in A_i$, as was to be shown.

A.4. Proof of Proposition 3

Given the realized training sample ε_i , the expected out-of-sample prediction error of partition Π equals

$$\mathbb{E}_\eta \sum_{t_i \in T_i} \mu(t_i) \left[\left(\sum_{t'_i \in \pi(t_i)} \mu(t'_i | t'_i \in \pi(t_i)) \varepsilon_{t'_i} \right) - \eta_{t_i} \right]^2.$$

By Jensen's inequality, the finer partition Π'_i achieves the same out-of-sample prediction error by splitting the cell $\pi(t_i)$ into the cells π'_1 and π'_2 only if⁸

$$\sum_{t'_i \in \pi(t_i)} \mu(t'_i | t'_i \in \pi(t_i)) \varepsilon_{t'_i} = \sum_{t'_i \in \pi'_1} \mu(t'_i | t'_i \in \pi'_1) \varepsilon_{t'_i} = \sum_{t'_i \in \pi'_2} \mu(t'_i | t'_i \in \pi'_2) \varepsilon_{t'_i}.$$

This implies that choosing the finer partition Π' induces the same beliefs as Π for any type t_i .

A.5. Proof of Proposition 4

Fix ε_i and Π_i . Consider a partition cell $M \in \Pi_i$. Denote

$$\mathbb{E}_\eta \sum_{t_i \in M} \mu(t_i | t_i \in M) [\hat{x}^{\varepsilon_i}(M) - (\bar{x}_i^\sigma(t_i) + \eta_{t_i})]^2 = c_M(\varepsilon_i)$$

Plugging the definition of \hat{x}^{ε_i} into the L.H.S., we obtain

$$\mathbb{E}_\eta \sum_{t_i \in M} \mu(t_i | t_i \in M) \left[\sum_{t'_i \in M} \mu(t'_i | t'_i \in M) (\bar{x}_i^\sigma(t'_i) + \varepsilon_{t'_i}) - \bar{x}_i^\sigma(t_i) - \eta_{t_i} \right]^2$$

Expand this expression. Since η_{t_i} is independently distributed with $E_\eta(\eta_{t_i}) = 0$, the only terms in the elaborated expression that includes ε_i are

$$\left(\sum_{t'_i \in M} \mu(t'_i | t'_i \in M) \varepsilon_{t'_i} \right)^2$$

⁸The argument when splitting $\pi(t_i)$ into more than two cells is the same.

and

$$2 \sum_{t_i \in M} \mu(t_i | t_i \in M) \left(\sum_{t'_i \in M} \mu(t'_i | t'_i \in M) \varepsilon_{t'_i} \right) \left(\sum_{t'_i \in M} \mu(t'_i | t'_i \in M) \bar{x}_i^\sigma(t'_i) - \bar{x}_i^\sigma(t_i) \right)$$

The second term is equal to zero. As to the first term, it remains the same if we replace ε_i with $-\varepsilon_i$. It follows that $c_M(-\varepsilon_i) = c_M(\varepsilon_i)$ for every partition cell M . Therefore, if Π_i satisfies (6) under ε_i , then it also does so under $-\varepsilon_i$. Since v is symmetric around zero, it follows that the expectation of ε_i conditional on Π_i satisfying (6) is zero.

A.6. Proof of Proposition 5

The expected out-of-sample prediction error of the fine partition is

$$\sum_{t_i \in T_i} \mu(t_i) \varepsilon_{t_i}^2 + \sigma_\eta^2$$

and the expected out-of-sample prediction error of the coarse partition is

$$\sum_{t_i \in T_i} \mu(t_i) \left(\sum_{t'_i \in T_i} \mu(t'_i) \varepsilon_{t'_i} + \sum_{t'_i \in T_i} \mu(t'_i) \bar{x}_i^\sigma(t'_i) - \bar{x}_i^\sigma(t_i) \right)^2 + \sigma_\eta^2.$$

Simple algebra shows that the fine partition is adopted if the inequality (10) holds.

References

- Asker, John, Chaim Fershtman, and Ariel Pakes (2024). “The impact of artificial intelligence design on pricing”. In: *Journal of Economics & Management Strategy* 33.2, pp. 276–304.
- Banchio, Martino and Giacomo Mantegazza (2024). “Adaptive Algorithms and Collusion via Coupling.” In: *working paper*.
- Calvano, Emilio et al. (2020). “Artificial intelligence, algorithmic pricing, and collusion”. In: *American Economic Review* 110.10, pp. 3267–3297.
- Cartea, Alvaro et al. (2025). “Algorithmic Collusion and a Folk Theorem from Learning with Bounded Rationality”. In: *Games and Economic Behavior*.
- Danenberg, Tuval and Ran Spiegler (2025). “A Representative-Sampling Model of Stochastic Choice”. In: *JPE: Micro (forthcoming)*.

- Dolgoplov, Arthur (2024). “Reinforcement Learning in a Prisoner’s Dilemma”. In: *Games and Economic Behavior* 144, pp. 84–103.
- Eliasz, Kfir and Ran Spiegler (2019). “The model selection curse”. In: *American Economic Review: Insights* 1.2, pp. 127–140.
- Esponda, Ignacio (2008). “Behavioral equilibrium in economies with adverse selection”. In: *American Economic Review* 98.4, pp. 1269–1291.
- Esponda, Ignacio and Demian Pouzo (2016). “Berk–Nash equilibrium: A framework for modeling agents with misspecified models”. In: *Econometrica* 84.3, pp. 1093–1130.
- Eyster, Erik and Matthew Rabin (2005). “Cursed equilibrium”. In: *Econometrica* 73.5, pp. 1623–1672.
- Feddersen, Timothy and Wolfgang Pesendorfer (1998). “Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting”. In: *American Political Science Review* 92.1, pp. 23–35.
- Hansen, Karsten T, Kanishka Misra, and Malleesh M Pai (2021). “Algorithmic collusion: Supra-competitive prices via independent algorithms”. In: *Marketing Science* 40.1, pp. 1–12.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Jehiel, Philippe (2005). “Analogy-based expectation equilibrium”. In: *Journal of Economic Theory* 123.2, pp. 81–104.
- Jehiel, Philippe and Erik Mohlin (2026). “Categorization in Games: A Bias-Variance Perspective”. In: *mimeo*.
- Jehiel, Philippe and Giacomo Weber (2025). “Endogenous clustering and analogy-based expectation equilibrium”. In: *Review of Economic Studies*, rda054.
- Osborne, Martin J and Ariel Rubinstein (1998). “Games with procedurally rational players”. In: *American Economic Review*, pp. 834–847.
- Possnig, Clemens (2024). “Reinforcement Learning and Collusion”. In: *mimeo*.
- Salant, Yuval and Josh Cherry (2020). “Statistical inference in games”. In: *Econometrica* 88.4, pp. 1725–1752.
- Spiegler, Ran (2016). “Bayesian networks and boundedly rational expectations”. In: *The Quarterly Journal of Economics* 131.3, pp. 1243–1290.
- (2026). “Machine-Learning to Trust”. In: *working paper*.
- Waizmann, Stephan (2025). “AI in Action: Algorithmic Learning with Strategic Consumers”. In: *mimeo*.