EQUILIBRIUM INFORMATION AGGREGATION UNDER MACHINE LEARNING

ANDREW ELLIS, MICHELE PICCIONE, SHENGXING ZHANG

Abstract. We introduce a framework for studying the equilibrium effects of machine learning. Agents process information using a Chow and Liu (1968) tree, a widely-used machine learning procedure that admits a closed-form solution. We apply the model to an asset market with dispersed information based on Hellwig (1980). The price mechanism fails to aggregate the information extracted by the algorithm, even approximately. While there are partial equilibrium benefits from access to algorithms, the price aggregates less information and more slowly than the rational equilibrium. Equilibrium typically features diverse world-models, demands, and utilities, even with ex ante identical agents.

1. Introduction

The large and increasing amount of information available in the economy makes behaving as a good Bayesian would difficult. One increasingly common solution is to rely on machine learning algorithms to process this information. While machine learning excels at prediction, it operates as a "black box," and agents often do not understand why their algorithms generate the forecasts they do. Despite this, algorithmic trading systems, robo-advisors, and automated market makers now account for a substantial share of financial market activity. As more and more of the economy relies on these algorithms, their use may affect endogenous, informative variables like prices. This can lead to a feedback loop: the information that the algorithm extracts from endogenous variables affects their distribution which affects the information extracted. This paper seeks to understand the consequences for how much information the price aggregates in equilibrium.

We develop a framework for analyzing competitive equilibria when agents form beliefs using a machine learning algorithm. Agents form beliefs using the *Chow and Liu (1968) Tree* (CLT) algorithm. It approximates the true probability distribution with one that factorizes according to a tree-structured Bayesian network, and finds the distribution that minimizes the Kullback-Leibler divergence from the truth. CLT is used in statistics and machine learning for dimensionality reduction, classification, and as a foundation for more complex dependence structures (Jiao et al., 2013). Unlike neural networks or random forests, CLT has a closed-form solution, making it amenable to equilibrium analysis while preserving essential features of machine learning: model selection, parsimony, and the potential for misspecification when agents face high-dimensional data.

We embed the algorithm in the canonical asset market setting of Hellwig (1980). Risk-averse traders trade a safe asset and a risky asset with unknown supply in a competitive market. The equilibrium price adjusts to clear the market. They can, in principle, learn about the return of the risky asset from both private information and the price. Traders form beliefs using a CLT tree calibrated to the equilibrium price distribution rather than Bayes rule.

The competitive equilibrium under machine learning has markedly different properties than the Bayesian one. In Grossman (1976), equilibrium prices fully reveal private information, but traders demands are measurable with respect to the equilibrium price alone. Hellwig (1980) showed that uncertain supply prevents full revelation but preserves *approximate* efficiency: as uncertainty about supply vanishes, prices reveal all private information. We show that even approximate informational efficiency is impossible, regardless of how small noise trading becomes.

Bayesian updating is computationally very demanding. CLT simplifies the computation by approximating a complex joint probability distribution with a simpler one. This solves the key computational difficulty, the curse of dimensionality. While a full joint distribution over many variables requires tracking an enormous number of parameters (exponential in the number of variables for discrete data, or a full covariance matrix for Gaussian data), a tree-structured distribution requires only tracking pairwise relationships along the branches of the tree.

Several features make this algorithm attractive as a model of machine learning. First, it is computationally efficient, running in polynomial time. Second, it produces sparse models: a small subset of relationships suffices for prediction. Third, and crucially for our purposes, the optimal tree has a closed-form characterization, enabling equilibrium analysis. Fourth, the algorithm is non-parametric in the sense that it does not require the analyst to specify a functional form for relationships. The main limitation relative to richer machine learning methods is that CLTs restrict each variable to have at most one "parent" in the dependence structure, precluding the deeper architectures of neural networks. Finding optimal higher-order trees is NP-hard (Chickering, 1996), which provides a computational justification for focusing on first-order trees.

We embed this algorithm in the canonical model of learning from an endogenous price. Traders each observe private signals about the value of a risky asset and submit demands for the asset after observing the market-clearing price. Traders have CARA utility over terminal wealth, and market clearing requires aggregate demand equal to a random supply. Each trader's algorithm processes training data consisting of all variables the trader eventually

observes: the realized asset value, the market price, and the trader's private signals. Given their equilibrium joint distribution, the algorithm constructs a CLT for it. The trader then uses the CLT, her private information, and the public price to update her beliefs, and submits the demand that maximizes expected utility.

A CLT equilibrium consists of a distribution of prices as well as (a distribution over) trees for each trader so that the following three conditions hold. First, each trader's demand must maximize utility given her tree and information. Second, each trader's tree must be a Chow-Liu Tree, that is, the best approximation to the equilibrium distribution of feedback. Third, the market must clear for every realization of signals, price, and supply shock.

The tree selected by the algorithm determines how the trader's demand responds to information. When more traders use trees that link a private signal to the fundamental value, they trade more aggressively on that signal, so the price must be more responsive to it. Ceteris paribus, the price becomes more correlated with it as well as with the underlying value. This changes the quality of the approximation that each tree yields, feeding back into which one the algorithm selects. In equilibrium, the distribution of the price that clears the market for the trees in the population makes each a CLT. Our first result shows that this feedback loop has a steady state: an equilibrium exists.

We then turn to the question of how the equilibrium differs from its rational counterpart. The most important distinction is that the equilibrium correlation between price and fundamental value is bounded strictly below that in the Bayesian equilibrium. In particular, the bound does not vanish as noise trading goes to zero. The logic runs as follows. Suppose, toward a contradiction, that prices were approximately informationally efficient – highly correlated with the fundamental value, nearly revealing all private information. Then by the data-processing inequality, the price would also be more correlated with each private signal than the signal is with the fundamental. The Chow-Liu algorithm would then connect all signals to the price node rather than to the value node, since those edges have higher mutual information. But traders using such trees would condition only on price, effectively ignoring their private signals. With no one trading on private information, the market cannot clear, as demand would not vary with the supply shock.

In equilibrium, this tension resolves through the distribution of trees adjusting to maintain market clearing. Some traders must use trees that incorporate private signals, which requires that the price not be too informative. The result is an upper bound on price informativeness that does not depend on the level of noise trading, a stark departure from Hellwig (1980), where prices become arbitrarily informative as noise vanishes, and Grossman (1976) where the price is a sufficient statistic for all private information in the economy.

The CLT equilibrium exhibits comparative statics that often reverse the predictions of rational expectations models. In Hellwig (1980), increasing the number of traders improves price informativeness through diversification of private information; in our model, increasing the number of traders can *reduce* welfare as traders increasingly rely on the (crowded) price signal rather than diverse private information. Similarly, while Grossman and Stiglitz (1980) predict that higher risk aversion reduces price informativeness, in our framework higher risk aversion can *increase* informativeness by shifting traders toward trees that incorporate private signals. These reversed comparative statics suggest that conventional policy intuitions may not apply when markets are populated by algorithmic traders using machine learning.

This finding has important implications. It suggests a non-vanishing private value of machine learning: even in liquid markets with minimal noise, traders who deploy sophisticated algorithms retain an informational advantage that is not competed away through prices. Since training such algorithms is costly, this provides an economic rationale for the persistent investment in algorithmic trading that we observe. The result also speaks to the Grossman-Stiglitz paradox—the puzzle of why traders invest in information acquisition when prices should reveal that information (Grossman and Stiglitz, 1980). In our framework, the paradox dissolves: traders have a persistent incentive to acquire and process information because the price never fully reveals it. In addition, ex ante identical traders may use different CLTs in equilibrium, leading to asymmetric reactions to the same information. With purely public but high-dimensional information, traders may still update beliefs in response to price changes even though prices are merely a garbling of the public signals.

Calvano et al. (2020); Klein (2021); Dolgopolov (2024) and others show that reinforcement learning, a different form of machine learning, leads to non-Nash outcomes in pricing games

and the prisoners dilemma.[1] Methodologically, these papers are essentially experimental with the algorithm taking the role of subjects. In contrast to these papers, we study a general equilibrium outcome and analytically characterize the outcome. Eliaz et al. (2021) is the only economics paper of which we are aware that explicitly uses CLT. They show that the Chow-Liu tree will select the tree that gets maximally tricked. Since the algorithm potentially introduces misspecification, our paper is also related to the recent literature on misspecified models (e.g., Spiegler, 2016; Esponda and Pouzo, 2016; Fudenberg et al., 2017; Heidhues et al., 2018; Bohren and Hauser, 2021; Frick et al., 2020). The CLT selects the best misspecification. This makes it particularly related to papers like Cho and Kasa (2015) and Ba (2024) that incorporate model selection and misspecification tests.

2. Chow-Liu Trees

A directed acyclic graph (DAG) $G = (N, R)$ is a set of nodes N and edges $R \subset N \times N$ with no directed cycles. We write xRy for $(x, y) \in R$. Consider an n-dimension random variable $X = (X_1, \ldots, X_n)$. An K-order dependence tree is a DAG $(\{1, \ldots, n\}, R)$, sometimes just denoted R, with the following two properties. First, if iRj and kRj then either iRk or kRi. Second, $|R(i)| \leq K$ for every node i where $R(i) = \{k : kRi\}$. The dependence tree R determines an approximation ν_R of a distribution ν using the formula

$$\nu_R(x_1, \ldots x_n) = \prod_{j=1}^{n} \nu\left(x_j | x_{R(j)}\right)$$

where $x_\emptyset$ is an arbitrary constant, and ν_R, ν are the probability density (or mass) functions of ν_R and ν, respectively.

Dependence trees have several computational advantages. First, the set

$$\{\nu_R : \nu \in \Delta(\mathcal{X}_1 \times \cdots \times \mathcal{X}_n)\}$$

has lower dimension than $\Delta(\mathcal{X}_1 \times \cdots \times \mathcal{X}_n)$. Second ν_R is fully determined by the marginal distribution on the cliques of R (subsets of nodes on which R is complete). Third, it leads to simple and computationally quicker updating. These capture salient and important features

[1]Banchio and Mantegazza (2023) argue this is due "spontaneous coupling", an endogenous statistical linkage between variables.

of modern machine learning models. Finally, increasing the order corresponds to an increase in the depth of the network. This increases the quality of the approximation at the cost of decreasing the amount of dimension reduction. This captures the inherent tradeoff in machine learning: higher predictive power requires more computation and entails a danger of overfitting. For textbook treatments of these topics, see Chapter 5 of Hajek et al. (1992) or Chapter 6 of Koski and Noble (2009).

In our analysis, we use on Chow-Liu Trees (CLT) (Chow and Liu, 1968), 1-order dependence trees that best approximate the underlying distribution of variables. Let $\mathcal{R}$ be the set of possible 1-order dependence trees. Formally, R^* is a CLT if

$$R^* \in \arg\min_{R\in\mathcal{R}} D_{KL}(\nu||\nu_R),$$

that is, ν_R minimizes KL-divergence from ν within that set, and where

$$D_{KL}(p||q) = \sum_x p(x)\log\left(\frac{p(x)}{q(x)}\right)$$

when X is discrete or

$$D_{KL}(P||Q) = \int \log\left(\frac{f(x)}{g(x)}\right) P(dx)$$

when X is continuous, $Q \ll P$, f is pdf of P, and g is the pdf of Q.

Chow and Liu (1968) show that the optimal CLT can be found using a Greedy Algorithm in finite time. The algorithm is simple. Consider undirected edges between each pair of variables. First, calculate the mutual information between the variables related by each edge, and order the edges according to their mutual. Start with an graph consisting of the first two edges. Check whether adding the next edge would create a cycle. If not, then add it to the graph. Otherwise, discard it, and proceed to the next edge. Continue until each edge is either added or discarded. The resulting graph can be oriented to form a 1-order dependence tree, and this tree has the lowest divergence amongst all such trees.

2.1. **Normal variables and CLT.** With normally distributed variables, the divergence can be reduced to correlations.

$$\begin{array}{ccccc} v & \rightarrow & p & & \\ \downarrow & & & & \\ s^i & & & & \end{array} \qquad \begin{array}{ccccc} s^i & \rightarrow & p & & \\ \downarrow & & & & \\ v & & & & \end{array} \qquad \begin{array}{ccccc} p & \rightarrow & s^i & & \\ \downarrow & & & & \\ v & & & & \end{array}$$

$$R_v \qquad\qquad R_s \qquad\qquad R_p$$

FIGURE 1. Possible 1-order trees with 3 variables

Lemma 1. *For any normal distribution p, there exists a constant $\kappa > 0$ so that*

$$D_{KL}\left(P||P_R\right) = \kappa + \sum_{(i,j)\in R} \frac{1}{2} \ln\left(1 - \rho\left(x_i, x_j\right)^2\right)$$

for every 1-order dependence tree R.

The divergence minimizing tree is the one that maximizes correlations between linked variables. To illustrate with 3 normal variables, every 1-order dependence tree takes form $R = i \rightarrow j \rightarrow k$. Then, the divergence can be written

$$\begin{aligned} D_{KL}(P||P_R) &= \frac{1}{2} \ln\left[\left(1 - \rho(i,j)^2\right) \sigma_i^2 \sigma_j^2\right] - \frac{1}{2} \ln\left[\left(1 - \rho(j,k)^2\right) \sigma_k^2 \sigma_j^2\right] \\ &\quad + \frac{1}{2} \ln\left[\sigma_j^2\right] - \frac{3}{2} \ln\left(2\pi e\right) - h(p) \\ &= \frac{1}{2} \ln\left(1 - \rho(i,j)^2\right) + \frac{1}{2} \ln\left(1 - \rho(j,k)^2\right) \\ &\quad - \ln\sigma_1 - \ln\sigma_2 - \ln\sigma_3 + \frac{3}{2} \ln\left(2\pi e\right) - h(p) \end{aligned}$$

where $h(p)$ is the Shannon entropy of the distribution p. Here, $\kappa = -\ln\sigma_1 - \ln\sigma_2 - \ln\sigma_3 + \frac{3}{2}\ln\left(2\pi e\right) - h(p)$. Inspecting the above reveals that the CLT links every pair of variables except the pair with minimal correlation.

With the three variables v, s^i, p, there are three possible 1-order dependence trees (Figure 1).Suppose agent tries to predict v. On the one hand, R_v leaves out the correlation between p and s^i conditional on v, so beliefs about v overreact to s given p. However, it is *precise*, in that it results in beliefs about v with relatively low variance. On the other hand, R_s (and R_p, with p replacing s in what follows) leaves out the correlation between v and p. Therefore, beliefs underreact to p given s relative to a Bayesian. However, beliefs are *accurate*, in the sense $E_{R_S}[v|s,p] = E[v|s]$.

2.2. **Higher order trees.** In general, the determination of how large K should be involves tradeoffs. First, finding best K-order tree is NP-hard for $K \geq 2$ (Chickering (1996)). Therefore, actually calculating the optimum may not be feasible and one may have to satisfice instead. Second, with limited (though still large) numbers of observation, the ideal K is determined by the bias-variance tradeoff in a process called regularization. A smaller K suffers from the potential for bias due to omitted relationships. However, larger K leads to more variance in estimates. This is due to the risk of overfitting because of outliers. In particular, the model must be estimated on each $K+1$ tuple of variables, and for a fixed dataset size, every particular realization of these variables occurs less often. In general, algorithms trade-off the greater expressive power of deeper learning (i.e., larger K) with the stability of estimates by using penalty functions. Our use of CLT amounts to assuming that the dataset is such that the penalty for $K > 1$ is is large enough that the algorithm chooses $K = 1$ as optimal.

3. Asset market model

There is a mass $I \in \{1, 2, 3, \dots\}$ of agents indexed by $i' \in (0, I]$. Trader $i' \in (i-1, i]$ is referred to as a trader of type i. Traders form a portfolio from a risky asset and a safe asset, and are able to take arbitrarily large short and long positions provided that they satisfy the budget constraint. The risky asset has a common value v where $v \sim N(0, \sigma_v^2)$. The safe asset has a return normalized to unity. Trader $i' \in (i-1, i]$ maximizes CARA utility with risk aversion parameter $\frac{1}{r^i} > 0$, so if she purchases x units of the risky asset at price p, her utility is $-\exp\left(-\frac{1}{r^i}x(v-p)\right)$

There are I signals, one for each trader type, and all traders of type i observe $s^i \in \mathbb{S}^i = \mathbb{R}^J$. Each s^i is a J dimensional vector with $s^i_j = v + \epsilon^i_j$, where $\epsilon^i \sim N(0, \Sigma^i)$ for $i = 1, \dots, I$, where each Σ^i is a symmetric, positive definite $J \times J$ matrix, independently of all other variables. Trader $i' \in (i-1, i]$ observes the price p of the risky asset and s^i. Let $\sigma^2_{i,j} = \Sigma^i_{jj}$ and $\vec{s} = (s^1, \dots, s^I)$. Supply of the asset equals u, where $u \sim N(0, \sigma_u^2)$ with $\sigma_u^2 > 0$ and is independent of all other variables.

Equilibrium consists of three endogenous objects. The first two are standard: a mapping from fundamentals to the price and a demand function for each trader that clears the market. The remaining object is a CLT for each trader, which we describe via a distribution over 1-order dependence trees for each type i. We consider linear equilibria where the price p is an affine function of the vector $(s,u) = \left(s^1,\ldots,s^I,u\right)$, i.e., there is an $\alpha \in \mathbb{R}^{IJ+2}$ so that $p = \alpha \cdot (1,\vec{s},u)$. For such an α, denote by $\nu(\cdot|\alpha)$ the distribution over all variables that results, and for each i, $\nu^i(\cdot|\alpha)$ the distribution over $(v,s^i_1,\ldots,s^i_J,p)$. Denote the 1-order dependence trees over $N^i = \{s^i_1,\ldots,s^i_J,v,p\}$ including $\{v,x\}$ for at least one $x \in N^i \setminus \{v\}$ by $\mathcal{R}^i$.

Definition 1. *A CLT equilibrium* is an $\alpha \in \mathbb{R}^{IJ+2}$ and a vector $\vec{\mu} = \left(\mu^1,\ldots,\mu^I\right)$ where μ^i is a distribution over $\mathcal{R}^i$ so that:

(1) for every $u \in \mathbb{R}$ and $\vec{s} \in \mathbb{R}^{IJ}$, $p = \alpha \cdot (1,\vec{s},u)$;

(2) for each trader i and tree $R \in \mathcal{R}^i$, every price $p \in \mathbb{R}$, and all signals $s^i \in \mathbb{R}^J$,

$$x^i_R\left(s^i,p\right) = \arg\max_x -\int \exp\left(-\frac{1}{r^i}x\left(v-p\right)\right)\nu_R\left(dv|\alpha,s^i,p\right);$$

(3) the distribution over trees clear the market for every $(\vec{s},u) \in \mathbb{R}^{IJ+1}$

$$\sum_{i=1}^{I}\sum_{R\in\mathcal{R}_i}\mu^i(R)x^i_R(s^i,p) = u;$$

(4) and each trader uses a CLT: for every i,

$$\mu^i(R) > 0 \implies R \in \arg\min_{R\in\mathcal{R}^i} D_{KL}\left(\nu^i(\cdot|\alpha)||\nu^i_R(\cdot|\alpha)\right).$$

First, the price is an affine function of the vector of signals and noise trader demand. Second, each trader maximizes expected utility given her tree, private information s^i, and the market price. Consequently, the demand from a trader of type i who uses the tree R is

$$x^i_R(s^i,p) = r^i\frac{E_R[v|s^i,p]-p}{Var_R[v|s^i,p]}$$

following the usual formulas. This means that the total demand from all traders of that type is

$$\sum_{R \in \mathcal{R}_i} \mu^i(R) x^i_R(s^i, p).$$

Third, given the demand functions above and the distribution of trees $\vec{\mu}$, the market clears for every realization of signals and noise trader demand. Finally, μ^i attaches support only to trees that minimize divergence given the equilibrium distribution.

4. Analysis

We begin by showing that an equilibrium exists for every economy as above. Then, we turn to the properties of this equilibrium. In every equilibrium of every economy, even approximate informational efficiency is impossible.

4.1. **Existence.** First, we show that an equilibrium exists.

Theorem 1. *A CLT equilibrium exists for any collection of parameters with $\sigma_u^2 > 0$.*

Given a fixed distribution of models μ, there exists a unique price that clears the market. The relative weights on different signals are determined by the ratio

$$\frac{\alpha_{i,j}}{\alpha_{k,l}} = \frac{\mu^i\left(\left\{R : \{v, s^i_j\} \in R\right\}\right)\sigma^{-2}_{i,j} r^i}{\mu^k\left(\left\{R : \{v, s^k_l\} \in R\right\}\right)\sigma^{-2}_{k,l} r^k}.$$

Signals receive higher weights when they are more frequently used, more precise, or utilized by more risk-loving traders. Notably, and unlike in standard Bayesian frameworks, the correlation between the signal and the price does not influence these weights. The resulting price determines the correlation between the fundamental value v and the price p, as well as between p and each signal s^i_j. The distribution of trees adjusts endogenously so that only the best approximations are used. This creates feedback effects on pricing and correlations.

The equilibrium balances two considerations. The more that the price responds to signals, the higher the correlation between the two. Therefore, fewer traders' trees use that signal, and demand becomes less responsive to changes in the signal. Conversely, the less that the price responds to a given signal, the lower the correlation between the two. Consequently,

more traders' trees link it to the value, and so demand becomes more responsive to it. In equilibrium, the price must respond just enough so that the market clears. If the correlation with price is too large for every signal, then all traders rely exclusively on information from the price. If it is instead too low for enough signals, then all traders use only their private information. Either poses difficulty for market clearing.

We prove existence of equilibrium using a nested fixed point problem. First, we characterize the unique price that clears the market for a given $\vec{\mu}$ via a continuous function, $\alpha(\vec{\mu})$. Then, we define a correspondence from a vector of distributions over trees to itself. The correspondence evaluated at $\vec{\mu}$ puts weight only on the trees that minimize divergence given the price is generated by $\alpha(\vec{\mu})$. Any fixed point of the correspondence is an equilibrium.

4.2. **Equilibrium information.** Replacing s^i with $\hat{s}^i = 1'(\Sigma^i)^{-1} s^i$, our setup is identical to that of Hellwig (1980). In that model, all traders are Bayesian and there is a unique equilibrium price p that reflects all private signals. In particular, the equilibrium price is *approximately informationally efficient*, in the following sense. As noise σ_u goes to zero,

$$E[v|p] \to E[v|\vec{s}].$$

That is, beliefs approach what they would be if all private signals were public. Given the normal distribution, this is equivalent to

$$\rho(v,p)^2 \to \rho(E[v|\vec{s}],v)^2 \equiv \rho^*.$$

The limiting price is a sufficient statistic for signals and so maximizes the correlation between the price and value.

In our setting, approximate informational efficiency fails.

Theorem 2. *Fixing $I > 1$, σ_v, and Σ^i for all i, let $\rho^* = \rho(E[v|\vec{s}],v)^2$. There exists $\delta > 0$ so that for any $\sigma_u > 0$, $\rho(p,v)^2 < \rho^* - \delta$ in any CLT equilibrium.*

The result shows that not all private information is aggregated in the market price, even approximately. After trading, private information still has value to a Bayesian decision maker. Moreover, traders get at least as much uncertainty reduction from using the tree

as a Bayesian would from the price. Put differently, traders perceive a non-negative and sometimes strictly positive value of using the algorithm and private information relative to just the public price. Since training a machine learning algorithm is typically costly, this result is consistent with persistent use of them in a market.

The proof can be found in the appendix. The key step utilizes the data processing inequality. Recall that $\mathbb{E}[v|\vec{s}]$ is a sufficient statistic for all private information, and that for pairwise normally distributed variables, mutual information is an increasing function of correlation squared. The data-processing inequality (Theorem 2.8.1 of Cover and Thomas, 2006) says that if $X \to Y \to Z$ is a Markov chain, then the mutual information between X and Y exceeds that between X and Z, with equality only if $X \to Z \to Y$ is also a Markov chain. If they are normally distributed, then mutual information is a strictly increasing function of the squared-correlation. Therefore, $\rho(\mathbb{E}[v|\vec{s}], s^i_j) > \rho(v, s^i_j)$ and $\rho(v, \mathbb{E}[v|\vec{s}]) > \rho(v, s^i_j)$. If the correlation between price and v is sufficiently close to $\rho(v, \mathbb{E}[v|\vec{s}])$, then $\rho(p, s^i_j) > \rho(v, s^i_j)$ and $\rho(v, p) > \rho(v, s^i_j)$ as well. But then by Lemma 1, the unique CLT is

$$\begin{array}{ccc} v \leftarrow & p \rightarrow & s^i_1 \\ & \downarrow\searrow & \\ & s^i_J & \cdots \end{array}$$

for every type of trader, so demand is invariant to private signals given the price. But then for markets to clear, $p = \alpha_u u$, implying $\rho(p, v) = 0$ and $\rho(p, s^i_j) = 0$, a contradiction.

5. Examples

To understand the forces in the model, we consider several example economies. The first considers information aggregation in the classic Hellwig (1980) setup with identical precisions and risk tolerance across traders. Then, the algorithm trades off between public information p and single-dimensional private information s^i. Our focus is on how the equilibrium differs from Bayesian benchmarks. The second focuses on what aspects of the information get used by the tree in a setting where all information is public. Then, the algorithm trades off between different dimensions of the information and potentially substitutes the garbling in

p. The general case has features of both, and the two are informative of what the general model looks like.

5.1. **Aggregation across signals.** First, suppose that all traders receive a one-dimensional signal ($J = 1$), that they all have the same risk aversion ($r^i = r$ for all i), and that signals are equally precise ($\Sigma_\epsilon^i = [\sigma_\varepsilon^2]$). Denote by R_j for $j \in \{v, s, p\}$ be the CLT that has two edges involving the node j. This is the model of Hellwig (1980), specialized so that agents are symmetric.

Theorem 3. *If $I > 1$, $J = 1$, $r^i = r$ for all i, and $\Sigma_\epsilon^i = [\sigma_\varepsilon^2]$ for all i, there is a symmetric CLT equilibrium $(\alpha, \vec{\mu})$. If $(I^2 - 2I)\,\sigma_v^2 > \sigma_\varepsilon^2$, then for every i,*

$$\mu^i(R_v) = \min\left\{1, \sqrt{\frac{\sigma_\varepsilon^2 \sigma_u^2 \sigma_v^2}{r^2\left(\sigma_\varepsilon^2 + I\sigma_v^2\right)}}\right\} = 1 - \mu^i(R_p).$$

If $(I^2 - 2I)\,\sigma_v^2 < \sigma_\varepsilon^2$, then for every i

$$\mu^i(R_s) = \min\left\{1, \sqrt{\frac{\sigma_u^2 \sigma_\varepsilon^2}{r^2 I\left(I-1\right)}}\right\} = 1 - \mu^i(R_p).$$

In the former case, $\rho(p, v) \geq \frac{I\sigma_v^2}{I\sigma_v^2 + \sigma_\varepsilon^2}$ and in the latter $\rho(p, v) \geq \frac{\sigma_v}{\sqrt{\sigma_v^2 + \sigma_\varepsilon^2}}$, with equality whenever $\mu^i(R_p) > 0$.

The result provides an explicit characterization of the equilibrium. The two cases obtain because for any α,

$$\left(I^2 - 2I\right)\sigma_v^2 \geq \sigma_\varepsilon^2 \iff \rho(v, p) \geq \rho(p, s^i),$$

so the inequality determines whether the CLT can contains an edge from v to both p and s^i or not. When σ_u^2 is large, the equilibrium price is not very informative about v nor highly correlated with any other variables. Consequently, the algorithm selects either R_s or R_v for all traders. When σ_u^2 is small enough, the price is more correlated with v. If the algorithm never selected the tree R_p, then the price would be sufficiently strongly correlated with both s^i and v that R_p would be the unique CLT, a contradiction. Consequently, the distribution of trees and price must adjust so that the algorithm finds both trees optimal.

5.1.1. *Benchmark: Hellwig (1980).* Consider first the economy where all traders are rational. Denote the equilibrium price $p^h = \alpha_s^h I^{-1} \sum s^i + \alpha_u^h u$. This special case of Hellwig (1980) is explicitly solved in Theorem 6.1 of Kyle (1989) via the system of equations

$$\begin{aligned}\frac{\varphi}{(1-\varphi)^3} &= r^2 \frac{I-1}{\sigma_u^2 \sigma_\varepsilon^2} \\ \alpha_s^h &= \frac{\sigma_\varepsilon^{-2}\left(1+\varphi\left(I-1\right)\right)}{\sigma_v^{-2}+\sigma_\varepsilon^{-2}\left(1+\varphi\left(I-1\right)\right)} \\ \alpha_u^h &= \alpha_s^h \sigma_\varepsilon^2 \frac{1}{Ir(1-\varphi)}.\end{aligned}$$

Letting $\tau = \sigma_v^{-2} + (1+\phi(I-1))\sigma_\varepsilon^{-2}$, equilibrium demand is

$$x^i(p, s^i) = r\mathbb{E}\left[v|s^i, p\right]\tau = (1-\phi)\sigma_\varepsilon^{-2} s^i - \left(\tau - I\phi\frac{\sigma_\varepsilon^{-2}}{\alpha_s}\right)p.$$

5.1.2. *Benchmark: Grossman and Stiglitz (1980) (GS).* We introduce an extension of GS to dispersed information. A trader $i' \in (i-1, i]$ can either give up $c > 0$ units of the safe asset to see the signal s^i in addition to the price or only submit a demand that is a function of the price.

Let $x_U^i(p)$ and $x_I^i(s^i, p)$ be the demands maximize the utility of a type i trader given p and (s^i, p) are observed, respectively. Equilibrium consists of a price $p^g = \alpha_s^g I^{-1} \sum_i s^i + \alpha_u^g u$ and a $\lambda \in [0, 1]$ fraction of traders that acquire the signal so that

$$\mathbb{E}\left[-\exp\left(-x_U^i(p)(v-p)/r\right)\right] \leq \mathbb{E}\left[-\exp\left(-(x_I^i(s^i,p)(v-p)-c\right)/r\right],$$

with equality whenever $\lambda < 1$, and so that

$$\sum_i [(1-\lambda)x_U^i(p) + \lambda x_I^i(s^i, p)] = u.$$

Defining

$$\begin{aligned}\varphi^I &= \frac{\left(\alpha_s^g\right)^2 (I-1)\,\sigma_\varepsilon^2}{(I-1)\left(\alpha_s^g\right)^2\sigma_\varepsilon^2 + I^2\left(\alpha_u^g\right)^2\sigma_u^2}\\ \varphi^U &= \frac{\left(\alpha_s^g\right)^2 I\sigma_\varepsilon^2}{\left(\alpha_s^g\right)^2 I\sigma_\varepsilon^2 + I^2\left(\alpha_u^g\right)^2\sigma_u^2}\\ \tau^I &= \sigma_v^{-2} + (1+\varphi^I(I-1))\sigma_\varepsilon^{-2}\\ \tau^U &= \sigma_v^{-2} + I\varphi^U\sigma_\varepsilon^{-2}\end{aligned}$$

we can determine that

$$x_I^i(p,s^i) = r\mathbb{E}\left[v-p|s^i,p\right]\tau^I = r(1-\varphi^I)\sigma_\varepsilon^{-2}s^i - r\left(\tau^I - I\varphi^I\frac{\sigma_\varepsilon^{-2}}{\alpha_s^g}\right)p.$$

and

$$x_U^i(p) = r\mathbb{E}\left[v-p|p\right]\tau^U = -r\left(\tau^U - I\varphi^U\frac{\sigma_\varepsilon^{-2}}{\alpha_s^g}\right)p.$$

Calculating expected utility, the indifference condition becomes

$$\sqrt{\frac{\tau^I}{\tau^U}} = \exp[c/r].$$

and this combined with

$$\begin{aligned}(I-1)\,\lambda^2 r^2 &= \sigma_\varepsilon^2\sigma_u^2\frac{\varphi^I}{(1-\varphi^I)^3}\\ \alpha_s^g &= \sigma_\varepsilon^{-2}\frac{\lambda\left(1-\varphi^I\right)+\varphi^* I}{\tau}\\ \alpha_u^g &= -\frac{\lambda\left(1-\varphi^I\right)+\varphi^* I}{Ir\tau\lambda\left(1-\varphi^I\right)}.\end{aligned}$$

pins down the model for sufficiently low cost.

5.1.3. *Comparison.* The CLT Equilibrium differs from that above Hellwig (1980) in several ways. Denote p^{CLT} the equilibrium price in a CLT equilibrium. In particular, for sufficiently small σ_u^2, the p^{CLT} is both less informative and more volatile than p^h.

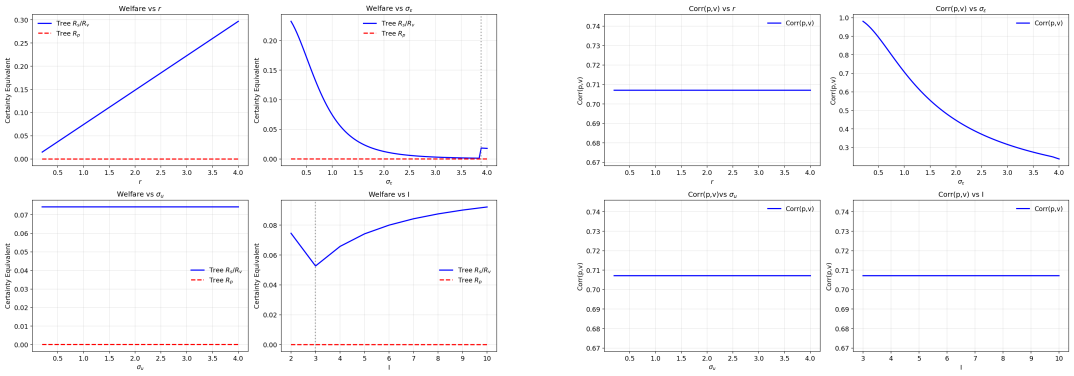


FIGURE 2. Welfare and correlation in CLT equilibrium

Corollary 1. *For σ_u^2 small enough, $\rho(v, p^{CLT}) < \rho(v, p^h)$ and $Var(p^{CLT}) > Var(p^h)$. For c small enough, the CLT equilibrium has lower $\rho(v, p^{CLT}) < \rho(v, p^g)$ and $Var(p^{CLT}) > Var(p^g)$.*

Proof. By Theorem 3, when $\mu^i(R_p) > 0$ the correlation $\rho(v, p^{CLT})$ is pinned at a constant that depends only on σ_v^2, σ_ε^2, and I, and in particular does not depend on σ_u^2. In the Hellwig equilibrium, $\rho(v, p^h) \to \rho^*$ as $\sigma_u^2 \to 0$, so for σ_u^2 small enough, $\rho(v, p^{CLT}) < \rho(v, p^h)$. The variance comparison follows from the explicit formulas for $Var(p^{CLT})$ and $Var(p^h)$ in the proof of Theorem 3 in the appendix, noting that $\alpha_s^{CLT} \geq \alpha_s^h$ while $h^{CLT} \geq h^h$ for small σ_u^2. The comparison with GS follows analogously, since $\rho(v, p^g) \to \rho^*$ as $c \to 0$ with σ_u^2 held fixed at any positive value. □

Moreover, it has significantly different comparative statics on both welfare and correlation between price and value than either of the two. See Figures 2, 3, and 4, which plot the equilibrium certainty equivalents of each trader by tree. In CLT these equivalents are calculated with respect to the objective equilibrium distribution of price and the demand functions that result from the equilibrium distribution of trees. These take the baseline parameters $I = 5$, $r = 1$, $\sigma_v^2 = 1$, $\sigma_\varepsilon^2 = 1$, and $\sigma_u = 0.01$.

Notice that welfare depends on the tree that the algorithm selects, and trees that are equally good approximations may yield different welfare (according to the true distribution

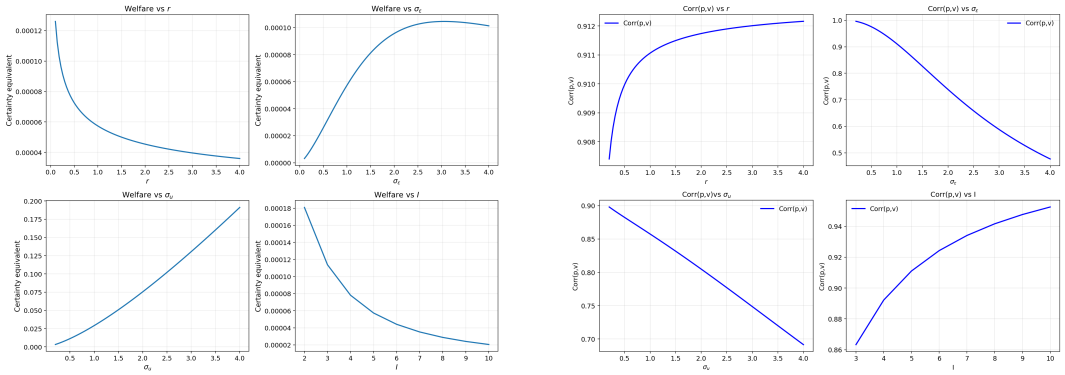


FIGURE 3. Welfare and correlation in Hellwig equilibrium

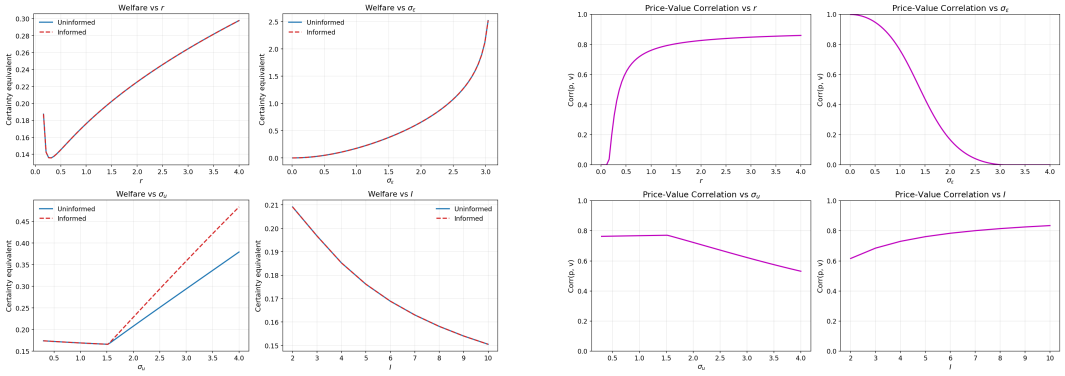


FIGURE 4. Welfare and correlation in GS equilibrium

of price). In particular, letting $X_R = x^i(p,s)/r$, then

$$\mathbb{E}_\nu\left[-\exp\left(X_R(v-p)\right)\right] = \frac{-1}{\sqrt{\left(1+Cov(X_R, v-p)\right)^2 - Var(X_R)Var(v-p)}}.$$

Utility is increasing in how correlated demand is with $v-p$ and decreasing in its own variance and that of p. We can show, for instance, that in equilibrium the utility of a trader using R_s is higher than that of one using R_p. For the former, $Var(X_R)$ is smaller and $Cov(X_R, v-p)$ is larger. The above plots use the certainty equivalent to adjust for risk aversion.

As I goes to infinity, both CLT and Hellwig prices become perfectly correlated with the state. However, the CLT equilibrium price converges more slowly than does the Hellwig equilibrium.

Corollary 2. *Let $p^{CLT}(I)$ and $p^h(I)$ be equilibrium prices in CLT and Hellwig models, respectively, as a function of the number of traders. Then, $\rho\left(v, \mathbb{E}[v|(s_i)_{i=1}^I]\right)^2 - \rho\left(v, p^{CLT}(I)\right)^2 = O(I^{-1})$ while $\rho\left(v, \mathbb{E}[v|(s_i)_{i=1}^I]\right)^2 - \rho\left(v, p^h(I)\right)^2 = O(I^{-4/3})$*

Proof. Let $K = \frac{r^2}{\sigma_u^2\sigma_\varepsilon^2}$ and $V(I) = K^{-1/3}I^{-1/3}$. Claim that $(1-\varphi)/V(I) \to 1$. This follows from

$$\frac{1-V(I)}{V(I)^3}\frac{(1-\varphi^I)^3}{\varphi^I} = \frac{1-K^{-1/3}I^{-1/3}}{KI^{-1}}\frac{1}{(I-1)K} = \frac{I-K^{-1/3}I^{2/3}}{(I-1)} \to 1$$

For I, the extra variance in p^h/α_s^h conditional on v relative to $E[v|\vec{s}]$ is

$$\phi^h(I) = \left(\frac{\alpha_u^h}{\alpha_s^h}\right)^2\sigma_u^2 = \frac{\sigma_\varepsilon^4\sigma_u^2}{I^2r^2(1-\varphi)^2}$$

By the above, $\phi(I) = O(I^{-4/3})$. Since $\rho\left(\mathbb{E}[v|(s_i)_{i=1}^I], v\right)^2 = \frac{\sigma_v^2}{\sigma_v^2+I^{-1}\sigma_\varepsilon^2}$, we have

$$\begin{aligned}\rho\left(\mathbb{E}[v|(s_i)_{i=1}^I], v\right)^2 - \rho(p^h, v)^2 =& \frac{\sigma_v^2 + \phi(I)\frac{\sigma_v^2}{\sigma_v^2+I^{-1}\sigma_\varepsilon^2}}{\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \phi(I)} - \frac{\sigma_v^2}{\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \phi(I)} \\ =& \frac{\phi(I)\frac{2\sigma_v^2+I^{-1}\sigma_\varepsilon^2}{\sigma_v^2+I^{-1}\sigma_\varepsilon^2}}{\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \phi(I)} = O(I^{-4/3})\end{aligned}$$

By Theorem 3, for I large, $\rho(v, p^{CLT}) = \frac{\sigma_v^2}{\sigma_v^2+I^{-1}\sigma_\varepsilon^2} = \rho\left(\mathbb{E}[v|(s_i)_{i=1}^I], v\right)^2$. Therefore,

$$\rho\left(\mathbb{E}[v|(s_i)_{i=1}^I], v\right)^2 - \rho(v, p^{CLT})^2 = \frac{\left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2\right)\sigma_v^2 - \sigma_v^4}{\left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2\right)^2} = \frac{I^{-1}\sigma_\varepsilon^2}{\left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2\right)^2} = O(I^{-1})$$

□

5.2. **Aggregation within signals.** Our second main benchmark concerns high-dimensional public information. The simplest case is captured by $I = 1$, $J = 2$, and

$$s^i|v \sim \mathcal{N}\left(v\vec{1}, \begin{bmatrix} \sigma_1^2 & \bar{\rho}\sigma_1\sigma_2 \\ \bar{\rho}\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}\right)$$

where $\bar{\rho} \geq 0$. WLOG, $\sigma_1^2 \leq \sigma_2^2$. This implies that $\rho(v, s_1) > \rho(v, s_2)$. There are more trees than before. Those in Figure 5) are representative ones.

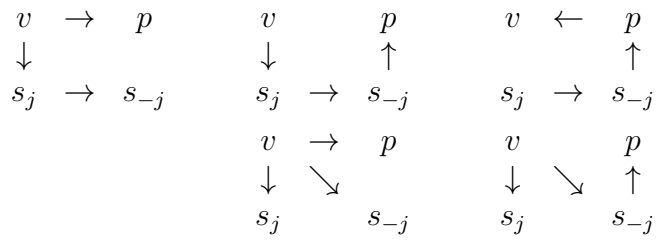


FIGURE 5. Possible CLTs with $J = 2$

Theorem 4. *There exists $\hat{\rho}(\sigma_v, \sigma_1, \sigma_2) > 0$ so that the CLT equilibrium μ has*

$$\mu\left(\left\{R : \left\{\{v, s_1\}, \{v, s_2\}\right\} \subset R\right\}\right) > 0$$

only if $\bar{\rho} \leq \hat{\rho}(\sigma_v, \sigma_1, \sigma_2)$.

Even when all traders observe the same signals, equilibrium features specialization when $\bar{\rho}$ is large. Only when the correlation between the two dimensions, conditional on v, is sufficiently low does any trader focus on both dimensions at the same time. When signals are sufficiently correlated with each other, the algorithm faces a tradeoff: each the two signals may be both strongly correlated with the value but they are also strongly correlated with each other. The optimal tree trades off more precise beliefs by having more edges with more accuracy in not double counting some correlated information. For certain parameters, the algorithm may pick an edge between the value and price instead of the edges between the signal and the value.

This contrasts sharply with the rational benchmark. There, all traders combine all available signals according to Bayes' rule, weighting each signal by its precision relative to its correlation with other signals. Information is fully utilized, and there is no role for either specialization or learning from the price.

Consider now a J-dimensional signal where each dimension has the same variance and has the same correlation with each other dimension. That is, $(\Sigma_{\varepsilon,J})_{j,j} = \sigma_\varepsilon^2$ and $(\Sigma_{\varepsilon,J})_{j,k} = \bar{\rho}\sigma_\varepsilon^2$ for $j \neq k$ and $\bar{\rho} \geq 0$. We let J go to infinity and study the limit. Then we have

$$\rho_J^* \equiv \rho(\mathbb{E}[v \mid s_1, \ldots, s_J], v)^2 = \frac{J\sigma_v^2}{J\sigma_v^2 + \sigma_\varepsilon^2(1 + (J-1)\bar{\rho})}.$$

The price with Bayesian trades is

$$p_J^h = \mathbb{E}[v|s_1, \ldots, s_J] - r(1-\rho_J^*)\sigma_v^2 u$$

and so we have

$$\rho(v, p_J^h)^2 = \frac{(\rho_J^*)^2\sigma_v^2}{\sigma^2(p)} = \frac{(\rho_J^*)^2\sigma_v^2}{(\rho_J^*)^2\left(\sigma_v^2 + \frac{1+(J-1)\bar{\rho}}{J}\sigma_\varepsilon^2\right) + \sigma_u^2 r^2 \sigma_v^4 (1-\rho_J^*)^2}$$

which tends to ρ_J^* as $\sigma_u^2 \to 0$.

Theorem 5. *If $\sigma_j^2 = \sigma_\varepsilon^2$ and $E[\varepsilon_k \varepsilon_j] = \bar{\rho}\sigma_\varepsilon^2$ for all $k \neq j$, then for any J,*

$$\rho(v, p^{CLT}) \leq \rho_J^*,$$

in any CLT equilibrium.

Proof. Note $\rho_J^* \geq \rho(v, s_j)^2$. So $\rho(v,p) \leq \rho(v, s_j)$ implies that the inequality holds. If $\rho(v,p) > \rho(v, s_j)$, then equilibrium requires that $\rho(v, s_j) \geq \rho(p, s_j)$. Otherwise, the divergence minimizing tree has no edge between v and s_j for any j. This follows from the greedy algorithm. The edges (v,p), (p, s_j), and possibly (s_j, s_j') all have more mutual information than (v, s_j), so the divergence minimizing tree either has a chain of edges between dimensions of the signals s_j (when $\rho(s_j, s_{j'}) \geq \rho(v,p)$) or has an edge between p and s_j for every j (when $\rho(s_j, s_{j'}) \leq \rho(v,p)$). (A hybrid is also possible when the correlations are equal.) But this tree makes market clearing impossible, or leads to $\rho(v,p) = 0$. Therefore,

$$\frac{\sigma_v}{\sqrt{\sigma_v^2 + \sigma_\varepsilon^2}} = \rho(v, s_j) \geq \rho(p, s_j) = \frac{\sigma_v^2 + \frac{1+(J-1)\bar{\rho}}{J}\sigma_\varepsilon^2}{\sqrt{\sigma_v^2 + \sigma_\varepsilon^2}\sigma(\alpha_s^{-1}p)} = \frac{\sigma_v^2 + \frac{1+(J-1)\bar{\rho}}{J}\sigma_\varepsilon^2}{\sqrt{\sigma_v^2 + \sigma_\varepsilon^2}\sigma_v}\rho(v,p)$$

(in any symmetric equilibrium), so

$$\frac{\sigma_v^2}{\sigma_v^2 + \frac{1+(J-1)\bar{\rho}}{J}\sigma_\varepsilon^2} \geq \rho(v,p)$$

must hold. □

Appendix A. Proofs

A.1. **Proof of Lemma 1.** Let $E \subset \{1, \ldots, n\}$ be such that $R(i) = \emptyset$ if and only if $i \in E$, and $J = \{1, \ldots, n\} \setminus E$. Slightly abusing notation, let $R(j) = \{R(j)\}$ for $j \in J$. Then we can write the density of P_R as

$$p_R(x) = \prod_{i \in E} p\left(x_i\right) \prod_{j \in J} p\left(x_j | x_{R(j)}\right)$$

where p is the density of P. Then, we have

$$\begin{aligned}
D_{KL}\left(P||P_R\right) &= \int p(x) \ln \frac{p(x)}{p_R(x)} dx = -\int p(x) \ln p_R(x) dx - h(p) \\
&= -\int p(x) \left(\sum_{i \in E} \ln p(x_i) + \sum_{j \in J} \ln p\left(x_j | x_{R(j)}\right) \right) dx - h(p) \\
&= -\sum_{i \in E} \int p\left(x_i\right) \ln p\left(x_i\right) dx_i - h(p) \\
&\quad - \sum_{j \in J} \int \int p\left(x_j, x_{R(j)}\right) \left[\ln p\left(x_j, x_{R(j)}\right) - \ln p\left(x_{R(j)}\right) \right] dx_j dx_{R(j)} \\
&= \sum_{i \in E} \left[\frac{1}{2}\left(1 + \ln 2\pi\right) + \frac{1}{2} \ln \sigma_i^2 \right] - h(p) \\
&\quad + \sum_{j \in J} \left[\left[\left(1 + \ln 2\pi\right) + \frac{1}{2} \ln \sigma_j^2 \sigma_{R(j)}^2 \left(1 - \rho\left(x_j, x_{R(j)}\right)^2\right) \right] - \left[\frac{1}{2}\left(1 + \ln 2\pi\right) + \frac{1}{2} \ln \sigma_{R(j)}^2 \right] \right] \\
&= \sum_{k=1}^{n} \left[\frac{1}{2}\left(1 + \ln 2\pi\right) + \frac{1}{2} \ln \sigma_k^2 \right] - h(p) + \sum_{j \in J} \left[\frac{1}{2} \ln \left(1 - \rho\left(x_j, x_{R(j)}\right)^2\right) \right].
\end{aligned}$$

The fourth equality comes from $p\left(x_j | x_{R(j)}\right) = p\left(x_j, x_{R(j)}\right) / p\left(x_{R(j)}\right)$, and the fifth comes from the formulae for differential entropy of a normal distribution. Taking $\kappa = \frac{n}{2}\left(1 + \ln 2\pi\right) + \sum_{j=1}^{n} \frac{1}{2} \ln \sigma_j^2 - h(p)$ completes the proof.

A.2. **Proof of Theorem 1.** Rewrite

$$p = \sum_i \sum_j \alpha_j^i s_j^i + \alpha_u u = \sum_i \sum_j \alpha_j^i \left(v + \epsilon_j^i\right) + \alpha_u u = \alpha_s \left(v + \sum_i \sum_j \hat{\alpha}_j^i \epsilon_j^i + \hat{\alpha}_u u \right)$$

for $\alpha_s = \sum_i \sum_j \alpha_j^i$, $\hat{\alpha}_j^i = \alpha_s^{-1}\alpha_j^i$, and $\hat{\alpha}_u = \alpha_s^{-1}\alpha_u$ when $\alpha_s > 0$. Observe $\sum_i \sum_j \hat{\alpha}_j^i = 1$. Letting $\mathbb{I}$ be the indicator function,

$$E_R[v|s^i, p = \alpha \cdot (s,u)] = \frac{\sum_j \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\sigma_{i,j}^{-2}s_j^i + \mathbb{I}_R\left(\{p,v\}\right)\alpha_s^{-1}\sigma_{\tilde{p}}^{-2}p}{\sigma_v^{-2} + \sum_j \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\sigma_{i,j}^{-2} + \mathbb{I}_R\left(\{p,v\}\right)\sigma_{\tilde{p}}^{-2}}$$

$$Var_R(s^i, p) = \left(\sigma_v^{-2} + \sum_j \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\sigma_{i,j}^{-2} + \mathbb{I}_R\left(\{p,v\}\right)\sigma_{\tilde{p}}^{-2}\right)^{-1}$$

$$\sigma_{\tilde{p}}^{-2} = \begin{cases} \left(\hat{\alpha}'_{-u}\Sigma_\epsilon\hat{\alpha}_{-u} + \hat{\alpha}_u^2\sigma_u^2\right)^{-1} & if\ \alpha_s > 0 \\ 0 & if\ \alpha_s = 0 \end{cases}$$

$$\sigma_{i,j}^2 = Var\left(\epsilon_j^i\right)$$

where Σ_ϵ is the covariance matrix for $\vec{\epsilon}$. Therefore, we have

$$x_R^i(s,u) = r^i \sum_{j=1}^{J} \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\sigma_{i,j}^{-2}s_j^i + r^i\left[\mathbb{I}_R\left(\{p,v\}\right)\alpha_s^{-1}\sigma_{\tilde{p}}^{-2} - Var_R(v|s^i,p)^{-1}\right]\left(\sum_k \sum_j \alpha_j^k s_j^k + \alpha_u u\right)$$

and aggregated demand equals

$$x^i(s,u) = \sum_R \mu(R)x_R^i(s,u) = \sum_{j=1}^{J}\gamma_j^i s_j^i + \gamma_p^i\alpha_s^{-1}p - \beta^i p$$

for

$$\gamma_j^i = r^i\sigma_{i,j}^{-2}\sum_R \mu^i(R)\left(\sum_j \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\right)$$

$$\gamma_p^i = r^i\sum_R \mu^i(R)\mathbb{I}_R\left(\{p,v\}\right)\sigma_{\tilde{p}}^{-2}$$

$$\beta^i = r^i\sum_R \mu^i(R)\left(\sigma_v^{-2} + \sum_j \mathbb{I}_R\left(\left\{s_j^i, v\right\}\right)\sigma_{i,j}^{-2} + \mathbb{I}_R\left(\{p,v\}\right)\sigma_{\tilde{p}}^{-2}\right)$$

Note that $\sigma_{\tilde{p}}^{-2}$ and β^i are functions of the DM's prediction about α, and that each γ_j^i is a function of only μ^i.

We first consider μ so that $\mu^i(R) \geq \epsilon > 0$ for all i and R. We show that there exists $\alpha(\vec{\mu})$ that clears the market for any such $\vec{\mu}$. Denote $\gamma_p = \sum_i \gamma_p^i \geq 0$ and $\beta = \sum_i \beta^i$, noting both depend only on $\sigma_{\tilde{p}}^{-2}$ (which depends in turn on α) and $\vec{\mu}$. Assume $\alpha_s > 0$, which must be the case if $\gamma_j^i > 0$ for some i and j for the market to clear. This is guaranteed by $\mu^i(R) > 0$ for all i and R. Then, it is legitimate to write the aggregate demand as

$$x(s,u) = \sum_i \sum_j \left(\gamma_j^i + \gamma_p \alpha_s^{-1} \alpha_j^i - \beta \alpha_j^i\right) s_j^i + \left[\gamma_p \alpha_s^{-1} - \beta\right] \alpha_u u.$$

Since $x(s,u) = u$, we must have

$$\begin{aligned} \gamma_j^i + \left(\gamma_p \alpha_s^{-1} - \beta\right) \alpha_j^i &= 0 \\ \frac{\gamma_j^i}{\beta - \gamma_p \alpha_s^{-1}} &= \alpha_j^i \\ \left[\gamma_p \alpha_s^{-1} - \beta\right] \alpha_u &= 1 \end{aligned}$$

This implies that

$$-\frac{\alpha_j^i}{\alpha_u} = \gamma_j^i = \sigma_{i,j}^{-2} r^i \mu^i \left(\left\{R : \left\{s_j^i, v\right\} \in R\right\}\right).$$

Adding up,

$$-\frac{\alpha_u}{\alpha_s} = \left(\sum_{i,j} r^i \mu^i \left(\left\{R : \left\{s_j^i, v\right\} \in R\right\}\right) \sigma_{i,j}^{-2}\right)^{-1} \equiv h \tag{1}$$

which yields

$$\hat{\alpha}_j^i = \frac{\alpha_j^i}{\alpha_s} = \frac{\gamma_j^i}{\sum_{i',j'} \gamma_{j'}^{i'}} = \frac{\sigma_{i,j}^{-2} r^i \mu^i \left(\left\{R : \left\{s_j^i, v\right\} \in R\right\}\right)}{\sum_{i',j'} r^{i'} \mu^{i'} \left(\left\{R : \left\{s_j^{i'}, v\right\} \in R\right\}\right) \sigma_{i',j'}^{-2}}. \tag{2}$$

Writing out $\gamma = \left(\frac{\gamma_j^i}{\sum_{i',j'} \gamma_{j'}^{i'}}\right)_{i,j}$ as an appropriately ordered vector,

$$\sigma_{\tilde{p}}^{-2} = \left(\gamma' \Sigma_\epsilon \gamma + h^2 \sigma_u^2\right)^{-1}; \tag{3}$$

that is, $\sigma_{\tilde{p}}^{-2}$ is determined uniquely by h and γ, so it is uniquely determined by $\vec{\mu}$. Since β^i and γ_p^i depend only on μ and $\sigma_{\tilde{p}}^{-2}$, they are also uniquely determined by $\vec{\mu}$. Finally, notice

that

$$1 = \alpha_u \left[\gamma_p \alpha_s^{-1} - \beta\right] = -h\gamma_p - \beta\alpha_u$$

pins down α_u uniquely, determining a unique $\alpha\left(\vec{\mu}\right)$ that solves the above system.

On the other hand, no equilibrium exists if $\gamma_j^i = 0$ for all i and j. For then,

$$x(s,u) = \sum \left(\gamma_p \alpha_s^{-1} \alpha_j^i - \beta\alpha_j^i\right) s_j^i + \left[\gamma_p \alpha_s^{-1} - \beta\right] \alpha_u u.$$

Equilibrium requires $\gamma_p \alpha_s^{-1} = \beta$ so that aggregate demand does not depend on s_j^i, and also that $\gamma_p \alpha_s^{-1} \neq \beta$ so that aggregate demand equals u. Clearly, this is a contradiction. This cannot happen when $\mu^i(R) \geq \epsilon > 0$ for all R, since then $\gamma_j^i \geq r^i \epsilon \sigma_{i,j}^{-2} > 0$.

Fix $\epsilon \in \left(0, \min_i \frac{1}{|\mathcal{R}_i|}\right)$, and

$$\mathcal{M}_\epsilon = \left\{\left(\mu^1, \ldots, \mu^I\right) \in \prod \Delta\mathcal{R}_i : \mu^i(R) \geq \epsilon \forall R \in \mathcal{R}_i \forall i\right\}.$$

Note that $\gamma_j^i > 0$ whenever $\vec{\mu} \in \mathcal{M}_\epsilon$, so a unique $\alpha\left(\vec{\mu}\right)$ exists that clears the market. Moreover, α is continuous by the implicit function theorem. We search for an ϵ equilibrium,

$$\mu^i(R) > \epsilon \implies R \in \arg\min_{R' \in \mathcal{R}_i} D_{KL}\left(\nu^i(\cdot|\alpha(\vec{\mu}))||\nu_R^i(\cdot|\alpha(\vec{\mu}))\right)$$

for all i. Consider the correspondence Γ where

$$\Gamma(\mu) = \left\{\vec{m} \in \mathcal{M}_\epsilon : m^i(R) > \epsilon \implies R \in \arg\min_{R' \in \mathcal{R}_i} D_{KL}\left(\nu^i(\cdot|\alpha(\vec{\mu}))||\nu_R^i(\cdot|\alpha(\vec{\mu}))\right)\right\}.$$

Γ has a closed graph and is non-empty and convex-valued. Since $\mathcal{M}_\epsilon$ is compact and convex, and each $\mathcal{R}_i$ is finite, Kakutani implies that a fixed point exists; call it μ_ϵ.

Consider a sequence $\epsilon_1, \epsilon_2, \ldots$ with $\epsilon_t \to 0$ so that $\mu_t = \mu_{\epsilon_t}$ converges to μ^*, respectively. This subsequence exists because $\mathcal{M}_0 \supset \mathcal{M}_\epsilon$ is compact. We show that the pair $(\mu^*, \alpha^* = \alpha(\mu^*))$ corresponds to an equilibrium.

Let $\alpha_t = \alpha(\mu_t)$. Define $\mathcal{R}_p^i = \{R \in \mathcal{R}^i : \{z, v\} \in R \implies z = p\}$, the set of trees that only use the price to learn about v. For contradiction, suppose that $\mu^*(\mathcal{R}_p^i) = 1$ for all i. Then, $\gamma_j^i\left(\mu_t\right) \to 0$ for all i and j which in turn implies $\frac{(\alpha_t)_j^i}{(\alpha_t)_u} \to 0$. But then $\rho(p,v), \rho(p, s_j^i) \to 0$ for all i, j and in particular for t large, $\rho(p,v), \rho\left(s_j^i, p\right) < \rho\left(s_j^i, v\right)$ for all i, j, and so $\mu_t^i(R') < \epsilon_t$

for all $R' \in \mathcal{R}_p^i$, a contradiction. This means that $\mu^{i*}(\mathcal{R}_p^i) < 1$ for some i, and therefore $\gamma_j^i(\mu^*) > 0$ for at least one i, j pair. Consequently, $\alpha(\mu^*)$ is well-defined, and $\alpha_t \to \alpha(\mu^*) = \alpha^*$. By the above, α^* clears the market for μ^*, and by the Berge maximum theorem, μ^* attaches probability 1 to divergence minimizing trees.

A.3. **Proof of Theorem 2.** Fix I, σ_v, and Σ^i. Note that

$$\begin{pmatrix} v \\ \vec{s} \end{pmatrix} \sim N\left(0, \begin{bmatrix} \sigma_v^2 & \Sigma_{S,v} \\ \Sigma_{v,S} & \Sigma_S \end{bmatrix}\right)$$

where Σ_S is a $IJ \times IJ$ matrix, $\Sigma_{v,S} = \Sigma_{S,v}'$ is $IJ \times 1$ matrix. Then

$$v|\vec{s} \sim N(\Sigma_{S,v}\Sigma_S^{-1}\vec{s}, \sigma_v^2 - \Sigma_{S,v}\Sigma_S^{-1}\Sigma_{v,S})$$

by the projection theorem, and $v \perp \vec{s}|E[v|\vec{s}]$. Letting $f^*(\vec{s}) = \Sigma_{S,v}\Sigma_S^{-1}\vec{s} = E[v|\vec{s}]$, for any $f : \mathbb{R}^{(IJ)^2} \to \mathbb{R}$, $v \perp f(\vec{s})|f^*(\vec{s})$, so

$$I^* \equiv I(v, f^*) \geq I(v, f)$$

by Theorem 2.8.1 of Cover and Thomas.

Let $\Delta = [-1, 1]^{IJ+1}$. Any α can be rescaled to $\alpha' \in \Delta$ without changing the mutual information between any pair of variables. Consider $p : \Delta \times [0, 1] \to \mathbb{R}$ so that $p(\alpha, \sigma_u) = \alpha \cdot (\vec{s}, u)$ in distribution, given that $u \sim N(0, \sigma_u^2)$ independently of the other variables and fix any i, j. Let

$$g = (\alpha, \sigma_u) \mapsto I(v, p(\alpha, \sigma_u)).$$

Then, $g^{-1}(I^*) = C$ is closed since g is the composition of continuous functions and non-empty since $\left(\Sigma_{S,v}\Sigma_S^{-1}, 0, \sigma_u\right) \in C$ for every σ_u. Note $I(v, p(\cdot))$, $I(p(\cdot), s_j^i)$, and $I(v, s_j^i)$ are all continuous in (α, σ_u). Moreover, for any $(\alpha, \sigma_u) \in C$, $I\left(s_j^i, p(\alpha, \sigma_u)\right) > I(s_j^i, v)$ and $I(v, p(\alpha, \sigma_u)) > I\left(v, s_j^i\right)$. This follows from $I(v, p(\alpha, \sigma_u)) = I(v, f^*)$ if and only if $p(\alpha, \sigma_u)$ is a sufficient statistic for v given $\vec{s}$, and for any $(\alpha^*, \sigma_u^*) \in C$, $I(v, p(\alpha, \sigma_u)) \geq I(v, s_j^i)$ and

$I\left(s_j^i, p\left(\alpha, \sigma_u\right)\right) \geq I(s_j^i, v)$ by Theorem 2.8.1 of Cover and Thomas. Set

$$G = (\alpha, \sigma_u) \mapsto \left(I\left(v, p\left(\alpha, \sigma_u\right)\right) - I(v, s_j^i), I\left(s_j^i, p\left(\alpha, \sigma_u\right)\right) - I(s_j^i, v)\right)_{i,j}$$

Since $\Delta \times [0,1]$ is compact and $G\left(\alpha, \sigma_u\right) > 0$ for all $(\alpha, \sigma_u) \in C$, there exists $c > 0$ so that $G\left(\alpha, \sigma_u\right) \geq (c, c, \ldots, c)$ for all $(\alpha, \sigma_u) \in C$ by the Weierstraus theorem.

Since $\Delta \times [0,1]$ is compact, we can take G to be uniformly continuous, so there exists $\epsilon > 0$ so that $G\left(a, \sigma_u\right) > \left(\frac{c}{2}, \frac{c}{2}, \ldots, \frac{c}{2}\right)$ for every $(a, \sigma_u) \in \cup_{(a', \sigma_u) \in C} B_\epsilon\left(a', \sigma_u\right) = O$. We claim that there exists $\delta^* > 0$ so that $g\left(a, \sigma_u\right) > I^* - \delta^*$ implies that $(a, \sigma_u) \in O$. If not, then for all $n \in \mathbb{N}$, $\exists y_n \in \Delta \times [0,1]$ so that $g(y_n) > I^* - \frac{1}{n}$ and $y_n \notin O$ for all n. For some $y^* \in \Delta \times [0,1]$, $y_n \to y^*$ (after taking a subsequence). By continuity and that $g(y_n) \leq I^*$, $g(y^*) = I^*$. But then $y^* \in C \subset O$, so $y_n \in O$ for all n large enough, a contradiction.

Now, if (α, σ_u) is such that $g\left(\alpha, \sigma_u\right) > I^* - \delta^*$, the CLT R_i for i is such that $\{v, x\} \in R_i$ if and only if $x = p$. The edge with highest weight involving either v or p is $\{v, p\}$, so $\{v, p\} \in R_i$. Replacing the edge $\left\{s_j^i, v\right\}$ with the edge $\left\{s_j^i, p\right\}$ decreases the divergence of the tree. This creates a cycle only if there are $k_1, \ldots, k_m$ so that $\left\{s_j^i, s_{k_1}^i\right\}, \ldots, \left\{s_{k_{m-1}}^i, s_{k_m}^i\right\}, \left\{s_{k_m}^i, v\right\} \in R_i$ or $\left\{s_j^i, s_{k_1}^i\right\}, \ldots, \left\{s_{k_{m-1}}^i, s_{k_m}^i\right\}, \left\{s_{k_m}^i, p\right\} \in R_i$. Since $\{v, p\} \in R_i$, $\left\{s_j^i, v\right\}$ would also create a cycle, so the replacement does not create a cycle. Therefore, $\{v, p\}$ is the only edge R_i containing v. However, this implies that $E_{R_i}[v|s^i, p] = E[v|p]$ for all i, which means that demand does not depend on s^i, making market clearing impossible with $\sigma_u > 0$.

Using formulas for normal distribution, $I^* = -(1-\rho(v, f^*)^2)$ and $I(v, f) = -(1-\rho(v, f)^2)$, so $\rho(v, f)^2 \leq \rho(v, f^*)^2$. In particular, since p is a function of $(\vec{s}, u)$ and u is independent of $(v, \vec{s})$, we have $v \perp p | f(\vec{s})$, so $\rho(v, p)^2 \leq \rho(v, f^*)^2$, and there exists $\delta > 0$ so that if $\rho(v, p)^2 > \rho(v, f^*)^2 - \delta$, then $I(v, p) > I^* - \delta^*$, and so p cannot be an equilibrium.

A.4. **Proof of Theorem 3.** By the above,

$$x(s, u) = \sum_i \left(\gamma^i + \gamma_p \alpha_s^{-1} \alpha^i - \beta \alpha^i\right) s^i + \left[I \gamma_p \alpha_s^{-1} - \sum \beta^i\right] \alpha_u u$$

where

$$\gamma^i = r\sigma_\varepsilon^{-2}\left(1-\mu^i(R_p)\right)$$
$$\gamma_p = \sum_i r\sigma_{\tilde{p}}^{-2}\left(1-\mu^i(R_s)\right)$$
$$\beta = \sum_i r\left(\sigma_v^{-2}+\sigma_{\tilde{p}}^{-2}\left(1-\mu^i(R_s)\right)+\sigma_\varepsilon^{-2}\left(1-\mu^i(R_p)\right)\right).$$

Suppose that $\mu^i = \mu$ for all i. Since $x(s,u) = u$, substituting into Equations (1) and (3) gives

$$h = -\frac{\alpha_u}{\alpha_s} = \frac{\sigma_\varepsilon^2}{Ir\left(1-\mu(R_p)\right)}$$
$$\sigma_{\tilde{p}}^{-2} = \left(\gamma'\Sigma_\epsilon\gamma + h^2\sigma_u^2\right)^{-1} = \left(I^{-1}\sigma_\varepsilon^2 + \frac{\sigma_\epsilon^4}{r^2I^2\left(1-\mu(R_p)\right)^2}\sigma_u^2\right)^{-1}.$$

Notice that, independent of μ and α, we have

$$\rho(v,p) \leq \rho(p,s^i) \iff \left(I^2-2I\right)\sigma_v^2 \leq \sigma_\varepsilon^2.$$

If $\rho(v,p) < \rho(p,s^i)$, then $\mu(R_v) = 0$, and equilibrium requires that $\rho(v,s^i) \geq \rho(v,p)$ since otherwise $\mu(R_p) = 1$. Now, $\rho(v,p) \leq \rho(v,s^i)$ holds if and only if $1-\mu(R_p) \leq \sqrt{\frac{\sigma_u^2\sigma_\varepsilon^2}{I(I-1)r^2}}$. So if $\sqrt{\frac{\sigma_u^2\sigma_\varepsilon^2}{I(I-1)r^2}} \geq 1$, $\mu(R_s) = 1$ is an equilibrium. Otherwise, $1-\mu(R_p) = \sqrt{\frac{\sigma_u^2\sigma_\varepsilon^2}{I(I-1)r^2}} = \mu(R_s)$ characterizes the equilibrium.

Moreover, when $0 < \mu(R_p)$, $Var_{R_s}(v|s,p) = Var_{R_p}(v|s,p) = \frac{\sigma_v^2\sigma_\varepsilon^2}{\sigma_v^2+\sigma_\varepsilon^2} = rI\left(\beta\right)^{-1}$. Denote the equilibrium price $p^s = \alpha_s^s I^{-1}\sum s^i + \alpha_u^s u$. From above we have

$$h = \sqrt{\frac{I-1}{I}}\frac{\sigma_\varepsilon}{\sigma_u}$$

Solving for α_s^s gives

$$\alpha_s^s = \frac{\sigma_\varepsilon^{-2}}{\sigma_v^{-2}+\sigma_\varepsilon^{-2}}$$

Using the above formulas for h and $\mu(R_p)$, we have

$$h^2\sigma_u^2 = \frac{I-1}{I}\sigma_\varepsilon^2$$

and therefore

$$\begin{aligned} Var(p^s) &= \left(\alpha_s^s I^{-1}\right)^2 \left(I^2\sigma_v^2 + I\sigma_\varepsilon^2\right) + \left(\alpha_s^s h\right)^2 \sigma_u^2. \\ &= \frac{\sigma_\varepsilon^{-4}}{\left(\sigma_v^{-2}+\sigma_\varepsilon^{-2}\right)^2}\left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \frac{I-1}{I}\sigma_\varepsilon^2\right) \end{aligned}$$

If $\rho(v,p) > \rho(p,s^i)$, then $\mu(R_s) = 0$. Equilibrium requires that $\rho(v,s^i) \geq \rho(p,s^i)$ since otherwise $\mu(R_p) = 1$. Now, $\rho(s^i,p) \leq \rho(v,s^i)$ holds if and only if $\frac{\sigma_\varepsilon^2\sigma_u^2\sigma_v^2}{r^2(\sigma_\varepsilon^2+I\sigma_v^2)} \geq (1-\mu(R_p))^2$. So if $\frac{\sigma_\varepsilon^2\sigma_u^2\sigma_v^2}{r^2(\sigma_\varepsilon^2+I\sigma_v^2)} \geq 1$, $\mu(R_v) = 1$ is an equilibrium. Otherwise, $\sqrt{\frac{\sigma_\varepsilon^2\sigma_u^2\sigma_v^2}{r^2(\sigma_\varepsilon^2+I\sigma_v^2)}} = \mu(R_v) = 1-\mu(R_p)$ characterizes the equilibrium.

Moreover, when $\mu(R_p) > 0$,

$$\tau_{R_v}^{-1} = Var_{R_v}(v|s,p) = \left(\sigma_v^{-2}+\sigma_\varepsilon^{-2}+\sigma_{\tilde{p}}^{-2}\right)^{-1} < \left(\sigma_v^{-2}+\sigma_{\tilde{p}}^{-2}\right)^{-1} = Var_{R_p}(v|s,p) = \tau_{R_p}^{-1}$$

Denote the equilibrium price $p^v = \alpha_s^v\left(I^{-1}\sum s^i\right) + \alpha_u^v u$. Then,

$$x^i(s^i,p) = r(1-\mu(R^p))\sigma_\varepsilon^{-2}s^i + r\left(\sigma_{\tilde{p}}^{-2}\left(\alpha_s^v\right)^{-1} - \tau\right)p$$

with

$$\tau = \mu(R_p)\tau_{R_p}v + (1-\mu(R_p))\tau_{R_v} = \sigma_{\tilde{p}}^{-2} + \sigma_v^{-2} + (1-\mu(R_p))\sigma_\varepsilon^{-2}.$$

As $\sigma_u^2 \to 0$, $\mu(R_p) \to 1$ so

$$\begin{aligned} \alpha_s^v &= \frac{(1-\mu(R_p))\sigma_\varepsilon^{-2} + \sigma_{\tilde{p}}^{-2}}{\sigma_{\tilde{p}}^{-2} + \sigma_v^{-2} + (1-\mu(R_p))\sigma_\varepsilon^{-2}} \\ &\to \frac{\sigma_{\tilde{p}}^{-2}}{\sigma_v^{-2}+\sigma_{\tilde{p}}^{-2}} \end{aligned}$$

Also,

$$h(\sigma_u^2)^2\sigma_u^2 = \left(\frac{\sigma_\varepsilon^2}{r\left(1-\mu(R_p)\right)}\right)^2\sigma_u^2 = \frac{r^2\left(I^{-1}\sigma_\varepsilon^2+\sigma_v^2\right)}{I\sigma_\varepsilon^2\sigma_u^2\sigma_v^2}\left(\frac{\sigma_\varepsilon^2}{r}\right)^2\sigma_u^2 = \sigma_\varepsilon^2\frac{I^{-1}\sigma_\varepsilon^2+\sigma_v^2}{I\sigma_v^2}$$

so

$$Var(p^v) \to \left(\frac{\sigma_{\tilde{p}}^{-2}}{\sigma_{\tilde{p}}^{-2}+\sigma_v^{-2}}\right)^2\left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \sigma_\varepsilon^2\frac{I^{-1}\sigma_\varepsilon^2+\sigma_v^2}{I\sigma_v^2}\right).$$

Since $\rho(v,p) > \rho(s^i,p) = \rho(s^i,v)$, $\sigma_{\tilde{p}}^{-2} > \sigma_\varepsilon^{-2}$, and so

$$\lim_{\sigma_u^2 \to 0} Var(p^v) \geq \left(\frac{\sigma_\varepsilon^{-2}}{\sigma_\varepsilon^{-2} + \sigma_v^{-2}}\right)^2 \left(\sigma_v^2 + I^{-1}\sigma_\varepsilon^2 + \frac{I\sigma_\varepsilon^{-2} + \sigma_v^{-2}}{I^2\sigma_\varepsilon^{-2}}\right).$$

A.5. **Proof of Theorem 4.** The correlations are

$$\rho(s_1^i, s_2^i) = \frac{\sigma_v^2 + \bar{\rho}\sigma_1\sigma_2}{\sqrt{(\sigma_v^2 + \sigma_1^2)(\sigma_v^2 + \sigma_2^2)}}$$
$$\rho(s_j^i, v) = \frac{\sigma_v}{\sqrt{\sigma_v^2 + \sigma_j^2}}$$

Note $\rho(s_1^i, v) \geq \rho(s_2^i, v)$ and

$$\rho(s_1^i, s_2^i) \geq \rho(s_2^i, v) \iff \bar{\rho} \geq \hat{\rho} \equiv \frac{\sqrt{(\sigma_v^2 + \sigma_1^2)}\sigma_v - \sigma_v^2}{\sigma_1\sigma_2}.$$

For a subset $I \subset \{s_1, s_2, p\}$, denote $R_I = \{R \in \mathcal{R} : (v,i) \in R \text{ or } (i,v) \in R \forall i \in I\}$. Suppose that $\bar{\rho} > \hat{\rho}$, so $0 < \rho(s_2, v) < \rho(s_1, s_2)$. Consider a DAG $R \in R_{s_1,s_2} \cup R_{s_1,s_2}$. Let R' remove the edge $\{v, s_2\}$ and add the edge $\{s_1, s_2\}$. This creates a cycle only if $\{s_2, p\}$ and $\{s_1, p\}$ belong to R. But then $v - s_1 - p - s_2 - v$ is a cycle in R. Moreover, by Lemma 1, $D_{KL}(\nu, \nu_{R'}) < D_{KL}(\nu, \nu_R)$. Therefore, $\mu(R) = 0$.

References

Cuimin Ba. Robust misspecified models and paradigm shift. *Working Paper*, 2024.

Martino Banchio and Giacomo Mantegazza. Artificial intelligence and spontaneous collusion. *arXiv preprint*, 2023.

J. Aislinn Bohren and Daniel N. Hauser. Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica*, 89(6):3025–3077, 2021.

Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–3297, October 2020. doi: 10.1257/aer.20190623.

David Maxwell Chickering. Learning bayesian networks is np-complete. In *Learning from Data: Artificial Intelligence and Statistics V*, pages 121–130. Springer, 1996.

In-Koo Cho and Kenneth Kasa. Learning and model validation. *Review of Economic Studies*, 82(1):45–82, 2015.

C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006. ISBN ISBN Number.

Arthur Dolgopolov. Reinforcement learning in a prisoner's dilemma. *Games and Economic Behavior*, 144:84–103, March 2024. doi: 10.1016/j.geb.2024.01.004.

Kfir Eliaz, Ran Spiegler, and Yair Weiss. Cheating with models. *American Economic Review: Insights*, 3(4):417–434, December 2021. doi: 10.1257/aeri.20200635.

Ignacio Esponda and Demian Pouzo. Berk-nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(2):1093–1130, 2016.

Mira Frick, Ryota Iijima, and Yuhta Ishii. Misinterpreting others and the fragility of social learning. *Econometrica*, 88(6):pp. 2281–2328, 2020.

Drew Fudenberg, Gleb Romanyuk, and Philipp Strack. Active learning with a misspecified prior. *Theoretical Economics*, 12(3):1155–1189, 2017.

Sanford Grossman. On the efficiency of competitive stock markets where trades have diverse information. *Journal of Finance*, 31(2):573–585, 1976.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Petr Hajek, Tomas Havranek, and Radim Jirousek. *Uncertain Information Processing in Expert Systems*. CRC Press, 1992.

Paul Heidhues, Botond Koszegi, and Philipp Strack. Unrealistic expectations and misguided learning. *Econometrica*, 86(4):1159–1214, 2018.

Martin F Hellwig. On the aggregation of information in competitive markets. *Journal of economic theory*, 22(3):477–498, 1980.

Jiantao Jiao, Haim H. Permuter, Lei Zhao, Young-Han Kim, and Tsachy Weissman. Universal estimation of directed information. *IEEE Transactions on Information Theory*, 59 (10):6220–6242, 2013.

Timo Klein. Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3):538–558, September 2021. doi: 10.1111/1756-2171.12383.

Timo Koski and John M. Noble. *Bayesian Networks: An introduction*. Wiley, 2009.

Albert S. Kyle. Informed speculation with imperfect competition. *The Review of Economic Studies*, 56(3):317–355, 1989.

Ran Spiegler. Bayesian networks and boundedly rational expectations. *Quarterly Journal of Economics*, 131(3):1243–1290, 2016.