Forecast Dispersion in Finite-Player Forecasting Games*

Jin Yeub Kim† Myungkyu Shim‡

October 25, 2017

Abstract

We study forecast dispersion in a finite-player forecasting game modeled as an aggregate game with payoff externalities and dispersed information. In the game, each agent cares about being accurate as well as about the distance of his forecast from the average forecast; and with a finite number of agents, the agents can strategically influence that average. We show that the finiteness of the number of agents weakens the strategic effect induced by the underlying preference. We find that when each agent prefers to be close to the average forecast, the presence of strategic manipulation of the average forecast contributes to a higher forecast dispersion; when instead each agent wants to be distinctive from the average, the opposite is true.

Keywords: forecast dispersion, finite-player, aggregate games, coordination, incomplete information.

JEL Classification Codes: C72, D82, D83, E37.

*We appreciate the very helpful suggestions from two anonymous referees. We also thank many for helpful comments and discussions, in particular, Lars Stole, Wilbert van der Klaauw, Jaeok Park, and the seminar participants at Michigan State University, Sogang University, Yonsei University, the Fall 2014 Midwest Economic Theory Conference, the 84th Annual Meetings of the Southern Economic Association, and the 11th World Congress of the Econometric Society.

†Department of Economics, University of Nebraska-Lincoln, 730 N. 14th Street, Lincoln, NE 68588-0489, USA, E-mail: shiningjin@gmail.com

‡School of Economics, Sogang University, 35 Baekbeom-Ro, Mapo-Gu, Seoul, 04107, South Korea, E-mail: audrb917@gmail.com

1 Introduction

In economic forecasting, an agent desires to give an accurate prediction that correctly estimates some underlying fundamentals. Different agents may obtain different information about the unknown state of the economy from news media reports, public announcements of government agencies or central banks, and private sources.[1] With dispersed information, an agent often cares about the aggregate (or average) prediction that reflects the agents' different information about the fundamentals, generating a coordination motive. The strategic effects in such environments are well-documented in the context of games with a continuum of agents (e.g., Angeletos and Pavan 2007).

But if there is a finite number of agents, so that an agent can influence the average forecast by changing his own forecast, then there is an additional channel of strategic effect.[2] The equilibrium behavior will reflect not only the coordination motive—i.e, the agents' desire to move toward or away from the average forecast—but also their ability to strategically manipulate that average. The study of finite-player forecasting games enables us to disentangle these two aspects of strategic behavior and their effects on forecast dispersion. The goal of this paper is to theoretically examine the *finite-player* equilibrium properties in forecasting games.

Our approach provides a micro foundation for explaining forecast dispersion in oligopolistic settings with a relatively small number of forecasters. This is important for the following reasons. When market decisions—such as production, investment, or budget planning decisions—are based on economic forecasts, the dispersion of forecasts can be undesirable. In this sense, understanding why agents disagree on their forecasts is of direct interest for economic policy. Moreover, in many situations there is only a limited number of agents who have information or credibility to provide forecasts. Then a better understanding of different strategic factors that affect forecast dispersion helps to clarify how much of the degree to which agents disagree is attributed to the limited number of agents.

For example, one can think of the forecasting game as being played by a group of financial

[1] Agents may give different predictions not because of differences in information but because of differences in their priors or models (Patton and Timmermann 2010).

[2] For ease of exposition, we use male pronouns for the agent.

analysts, each of whom desires to accurately estimate a company's earnings over the coming years. Holding the accuracy fixed, each analyst might prefer that his forecast is further away from the so-called consensus forecast—the average of all the forecasts from individual analysts tracking a particular stock. In some other cases, analysts might prefer to herd toward the consensus forecast.[3] What is important is that the consensus forecast could be the average of 30 analyst forecasts or, for a smaller company's stock, the average of just two forecasts.[4] The smaller the number of analysts tracking the company, the greater the extent to which each analyst can strategically influence the consensus forecast. We study such strategic effects on the dispersion of forecasts.

To model our finite-player forecasting game, we use the framework of aggregate games in which each agent's payoff is a function of his own strategy and some aggregator of the strategy profile of all agents.[5] In our forecasting game, each agent receives two signals—private and public—about the true fundamentals. An agent's strategy is a mapping from his private and public signals to a prediction of the fundamentals. Each agent's payoff depends on the distance of his prediction from the fundamentals and from the average forecast across the population. The agents prefer to be close to the fundamentals, but they may want either to be close to or to be distinct from the "herd."

In this setting, the optimal forecast takes the form of a linear combination of private and public signals, where the weights on the signals depend on the degree of agents' private value to coordinate. Importantly, this private value of coordination subsumes both the agents' incentive to herd or stand out and their incentive to strategically influence the average forecast. The former incentive, which we call the *herding motive*, arises from the preference structure of our game; whereas the latter, which we call the *market power motive*, is only present in games with a finite set of players.[6] To highlight the differences in the nature of

[3] Croushore (1997, 3) notes that "some participants might shade their forecasts more toward the consensus (to avoid unfavorable publicity when wrong), while others might make unusually bold forecasts, hoping to stand out from the crowd."

[4] See Ben McClure, "Earnings Forecasts: A Primer," accessed August 22, 2017, http://www.investopedia.com/articles/stocks/06/earningsforecasts.asp.

[5] Aggregate games with additively separable aggregators, such as the mean, are studied in Acemoglu and Jensen (2013) and Cornes and Hartley (2012). The general definition of aggregate games with a linear aggregate is treated in Martimort and Stole (2012).

[6] The term "private value of coordination" is borrowed from Angeletos and Pavan (2007) in which it is used to capture only the herding motive; in their environment with a continuum of agents, there is no market

strategic behavior, we use the limit game where agents have no market power as a benchmark.

We find that in a game where the agents want to herd around the average forecast, each agent assigns *less* value to aligning his forecast with others; in equilibrium, each agent adjusts upward his reliance on private information, as compared to the benchmark. This adjustment represents the market power distortion of the agent's behavior that weakens the underlying degree of herding motive. That is, any agent's inclination to herd is not entirely transferred to his reliance on public information given the non-negligible influence that any agent's forecast, thus his private information, can have on the average forecast. Consequently, the dispersion of forecasts in the finite-player game with preference for herding is higher than that in the benchmark. The opposite is true in a game where the agents want to be distinctive from the average forecast.

2 Finite-Player Forecasting Game

In the economy there is a finite number of agents (forecasters), each of whom is indexed by i, and the number of agents is $n \in \mathbb{N}$ where $n \geq 2$. We represent the true fundamentals of the economy with an exogenous random variable $\theta \in \mathbb{R}$ drawn from the (improper) uniform distribution over the real line.[7] Each agent i chooses a prediction of θ, which we denote as forecast $a_i \in \mathbb{R}$, and receives a payoff u_i. This payoff depends on the agent's own forecast and an aggregate of all agents' forecasts, the property of which defines an aggregate game. We consider the average forecast across the population $A_n \equiv \frac{1}{n}\sum_{i=1}^{n} a_i$ as the aggregator.

Each agent cares both about being correct, generating a fundamental motive to be close to the true θ, and about his distance to the aggregator A_n, which generates a co-ordination motive. For tractability, we assume that agent's preferences are quadratic to ensure linearity in the best responses. Formally, agent i's payoff is given by $u_i(a_i, A_n, \theta) =$

power motive. While they examine the effects of payoff externalities, we focus on the effects of market power. Bergemann, Heumann and Morris (2015) also study games with a finite number of agents and noisy signals, but their focus is on the interaction of market power and information.

[7] In other words, θ is drawn according to the Lebesgue measure on the real line. This specification of the common prior on θ is for simplicity of analysis; improper priors are well-behaved when concerned with conditional beliefs. See Morris and Shin (2003) for a brief discussion of improper priors.

$-\frac{1}{2}\left((1-r)(a_i-\theta)+r(a_i-A_n)\right)^2$ or, equivalently,

$$u_i(a_i, A_n, \theta) = -\frac{1}{2}\left(a_i-(1-r)\theta-rA_n\right)^2 \tag{1}$$

where the parameter r, which is the same across all agents, gives the weight that the agent puts on the aggregator relative to the fundamentals.[8] We assume $r<1$ to guarantee uniqueness of equilibrium. When $r=0$, each agent cares only about being close to the true θ. When $r>0$, each agent benefits the most by making an accurate prediction that is closer to the average forecast. When $r<0$, each agent desires to correctly predict the fundamentals but also prefers to be distinctive from the "herd."

Agents do not observe the realization of the true θ but instead observe noisy signals that are informative about the underlying fundamentals. Each agent i observes a public signal $p=\theta+(\alpha_p)^{-1/2}\varepsilon$ and a private signal $x_i=\theta+(\alpha_x)^{-1/2}\varepsilon_i$. The common noise ε follows $N(0,1)$, independent of θ. The idiosyncratic noises ε_i follow $N(0,1)$, independent of each other as well as of θ and ε. We let α_p and α_x denote the precision of public and private signals, respectively.

Because a change in a_i exerts a non-negligible effect on the average forecast A_n in the finite-player game, A_n is a function of $(\mathbf{x}, p)$ where $\mathbf{x}=(x_1,\ldots,x_n)$ is an arbitrary private signal profile. The following definition of a linear equilibrium is analogous to that in Angeletos and Pavan (2007) and Morris and Shin (2002).

Definition 1. *A linear equilibrium is a strategy profile* $\mathbf{a}^*=(a_1^*,\cdots,a_n^*)$ *where* $a_i^*:\mathbb{R}^2\to\mathbb{R}$ *for each* $i\in\{1,\ldots,n\}$ *such that* a_i^* *is linear in* x_i *and* p*, and satisfies* $a_i^*(x_i,p)=\arg\max_{a'}\mathbb{E}\big[u_i(a', A_n(\mathbf{x},p),\theta)|x_i,p\big]$ *for all* (x_i,p)*.*

Agent i's best response is then determined by the first order condition $a_i(x_i,p)=$

[8] Other specifications of quadratic loss utilities can be used as long as the preference features of our interest are preserved with appropriate assumptions on partial derivatives for tractability. For example, we may adopt $u_i(a_i,A_n,\theta)=-(1-r)(a_i-\theta)^2-r(a_i-A_n)^2$ under which all of our results continue to hold. Our use of (1) is only to simplify the exposition of the slope of best responses with respect to the average of other agents' forecasts.

$\mathbb{E}\left[(1-r)\theta + rA_n(\mathbf{x},p)|x_i,p\right]$ for all (x_i,p), which can be rewritten as:

$$\begin{aligned} a_i(x_i,p) &= \mathbb{E}\Big[(1-r)\theta + \frac{r}{n}a_i(x_i,p) + \frac{r}{n}(n-1)A_{n-i}((x_j)_{j\neq i},p)\Big|x_i,p\Big], \forall(x_i,p) \\ \leftrightarrow\ a_i(x_i,p) &= \mathbb{E}\left[(1-\gamma)\theta + \gamma A_{n-i}((x_j)_{j\neq i},p)|x_i,p\right],\ \forall(x_i,p) \end{aligned} \tag{2}$$

where $A_{n-i}((x_j)_{j\neq i},p) \equiv \frac{1}{n-1}\sum_{j\neq i} a_j(x_j,p)$ and $\gamma \equiv \frac{r(n-1)}{n-r}$. The parameter γ denotes the *equilibrium degree of coordination*, which measures how agents value aligning their forecasts with the forecasts of others in equilibrium, à la Angeletos and Pavan (2007).

Proposition 1. *For any given value of* (θ,p)*, a linear equilibrium exists and is the unique equilibrium. The equilibrium forecast of agent* i *is given by*

$$a_i^*(x_i,p) = \lambda_n x_i + (1-\lambda_n)p,\ \forall i \in \{1,\ldots,n\} \tag{3}$$

where $\lambda_n = \frac{\alpha_x}{\alpha_x + \frac{1}{1-\gamma}\alpha_p}$.

Proof. The proof follows the similar arguments in Angeletos and Pavan (2007) but with a finite number of agents. Given linearity in (2), it is natural to look for an equilibrium strategy that is linear in x_i and p so that $a_i = \kappa_0 x_i + \kappa_1 p$, where κ_0 and κ_1 are constants determined in equilibrium. Substituting $\mathbb{E}[\theta|x_i,p] = \delta x_i + (1-\delta)p$, where $\delta = \frac{\alpha_x}{\alpha_x+\alpha_p}$ (because the prior on θ is improper; see Morris and Shin 2003), and $A_{n-i}((x_j)_{j\neq i},p) = \frac{1}{n-1}\sum_{j\neq i} a_j$ in (2), we have $a_i(x_i,p) = (1-\gamma)(\delta x_i + (1-\delta)p) + \frac{\gamma}{n-1}\sum_{j\neq i}\mathbb{E}[a_j|x_i,p]$. Plugging the candidate equilibrium strategy in a_j and using $\mathbb{E}[x_j|x_i,p] = \mathbb{E}[\theta|x_i,p]$ (also see Morris and Shin 2003), we obtain $a_i(x_i,p) = \big((1-\gamma)\delta + \gamma\delta\kappa_0\big)x_i + \big((1-\gamma)(1-\delta) + \gamma\kappa_1 + \gamma(1-\delta)\kappa_0\big)p$. It follows that $a_i(x_i,p) = \kappa_0 x_i + \kappa_1 p$ constitutes a linear equilibrium if and only if $\kappa_0 = \alpha_x/\big[\alpha_x + \frac{1}{1-\gamma}\alpha_p\big]$ and $\kappa_1 = \big[\frac{1}{1-\gamma}\alpha_p\big]/\big[\alpha_x + \frac{1}{1-\gamma}\alpha_p\big]$. Let $\lambda_n \equiv \frac{\alpha_x}{\alpha_x+\frac{1}{1-\gamma}\alpha_p}$; then we obtain (3). Uniqueness can be verified as in Morris and Shin (2002). □

The equilibrium strategy (3) dictates how the agents allocate their use of private and public information. The equilibrium allocation is characterized by the coefficient λ_n, which measures the sensitivity of the equilibrium forecasts to private information relative to public information.

Our goal is to disentangle the strategic effects due to the finiteness of the number of agents from those due to the intrinsic payoff structure of the game; thus gauging to what extent the finiteness affects forecast dispersion. For this purpose, we find it useful to identify two benchmarks. The first benchmark is the forecasting game with $r = 0$, referred to as a simple-prediction game. In this benchmark, there is no strategic behavior featured in the equilibrium use of information where the two types of information are given weights that are commensurate with their precision: $\lambda_n = \frac{\alpha_x}{\alpha_x+\alpha_p} \equiv \delta$. The second benchmark is the forecasting game as $n \to \infty$, referred to as an infinite-player game.[9] As n goes to infinity the γ converges to r, and so $\lim_{n\to\infty} \lambda_n = \frac{\alpha_x}{\alpha_x+\frac{1}{1-r}\alpha_p} \equiv \lambda$.

3 Strategic Behavior and Forecast Dispersion

Our finite-player forecasting game imparts two types of strategic incentives: the herding motive and the market power motive. In our present context, we say that an agent exercises market power if, by changing his own forecast, he can strategically manipulate the average forecast. So we use the partial derivative $\partial A_n/\partial a_i = 1/n$ as a measure for market power of each agent.

The equilibrium degree of coordination γ reflects both the agents' herding motive and their market power motive. To distinguish the two notions of strategic incentives, we decompose γ as follows:

$$\gamma = r + r^* \tag{4}$$

where $r^* \equiv \frac{-r(1-r)}{n-r} \to 0$ as $n \to \infty$. So the parameter r denotes the equilibrium degree of coordination in the infinite-player game, whereas r captures the degree of the herding motive—the importance that the agents are inclined to attach to the average forecast—that is intrinsic to our finite-player game.

Lemma 1. *For any given n such that $2 \leq n < \infty$,*

(i) When $r > 0$: $\gamma < r$ and $\frac{\partial \gamma}{\partial n} > 0$.

(ii) When $r < 0$: $\gamma > r$ and $\frac{\partial \gamma}{\partial n} < 0$.

[9] As the number of agents increases, we obtain a well-defined limit forecasting game. Hence, the infinite-player game is an approximation to the forecasting game with a large number of agents.

Proof. $\gamma = \frac{r(n-1)}{n-r} \lesseqgtr r$ iff $r \gtreqless 0$ because $r < 1$ and $n \geq 2$. Also, $\frac{\partial \gamma}{\partial n} = \frac{r(1-r)}{(n-r)^2}$ is strictly positive when $r > 0$ and strictly negative when $r < 0$. $\square$

In the finite-player forecasting game, a change in agent i's choice of forecast exerts a non-negligible effect on the average forecast of the population. In equilibrium each agent takes into account such externality by appropriately adjusting how he privately values aligning his forecast with those of others. The forecasts of others are summarized by A_{n-i}, which is contaminated with the other agents' private noises. Consequently, the externality causes each agent to care less about what other agents do, thus weakening the underlying herding motive. Given the degree of the herding motive, stronger market power (in terms of a higher $1/n$) increases $|r^*|$, which effectively gauges the degree of the market power motive.

The equilibrium strategy characterized by the coefficient $\lambda_n = \frac{\alpha_x}{\alpha_x + \frac{1}{1-\gamma}\alpha_p}$ then respects both the herding and market power motives that are together captured by γ. We can disentangle the herding and market power effects on the equilibrium behavior by decomposing λ_n as follows:

$$\lambda_n = \lambda + \underbrace{\frac{\frac{\gamma}{n-1}\lambda(1-\lambda)}{1 + \frac{\gamma}{n-1}\lambda}}_{\equiv \lambda^*} \tag{5}$$

where $\lambda = \frac{\alpha_x}{\alpha_x + \frac{1}{1-r}\alpha_p}$ is the sensitivity of the equilibrium forecasts to private information in the infinite-player game. We can further break down λ into two terms as $\lambda = \delta + \left[-\frac{r\delta(1-\delta)}{1-r\delta}\right]$, where $\delta \equiv \frac{\alpha_x}{\alpha_x + \alpha_p}$ is the sensitivity of the equilibrium forecasts to private information in the simple-prediction game.

We formalize the comparison of the relative sensitivity of the equilibrium to private information between our game and the two benchmarks.

Lemma 2. *For any given α_x and α_p and for any given n such that $2 \leq n < \infty$,*

(i) When $r > 0$: $\lambda < \lambda_n < \delta$, and $\frac{\partial \lambda_n}{\partial n} < 0$.

(ii) When $r < 0$: $\delta < \lambda_n < \lambda$, and $\frac{\partial \lambda_n}{\partial n} > 0$.

Proof. We can see from (5) that $\lambda_n \lesseqgtr \lambda$ iff $\gamma \lesseqgtr 0$, which holds iff $r \lesseqgtr 0$. Also, $\frac{\partial \lambda_n}{\partial n} = \frac{-r\frac{\alpha_x \alpha_p}{n^2(1-r)}}{\left(\alpha_x + \frac{n-r}{n(1-r)}\alpha_p\right)^2}$ is strictly positive if $r < 0$ and is strictly negative if $r > 0$. Further, $\lambda_n = \frac{\alpha_x}{\alpha_x + \frac{1}{1-\gamma}\alpha_p} \lesseqgtr \frac{\alpha_x}{\alpha_x + \alpha_p} = \delta$ iff $\gamma \gtreqless 0$ (or $r \gtreqless 0$). $\square$

Lemma 2 establishes that in the finite-player forecasting game where the agents want to herd (resp. be distinctive), each agent adjusts upward (resp. downward) his reliance on private information relative to the infinite-player game, but not above (resp. below) the level of reliance in the simple-prediction game. This adjustment measured by λ^* in (5) represents what we call the *market power distortion* in the equilibrium choice of information. The size of such distortion is larger with fewer agents, thus with stronger market power. Figure 1 illustrates Lemma 2 for the cases of $r = 0.5$ and $r = -0.5$ with $\delta = 0.5$.

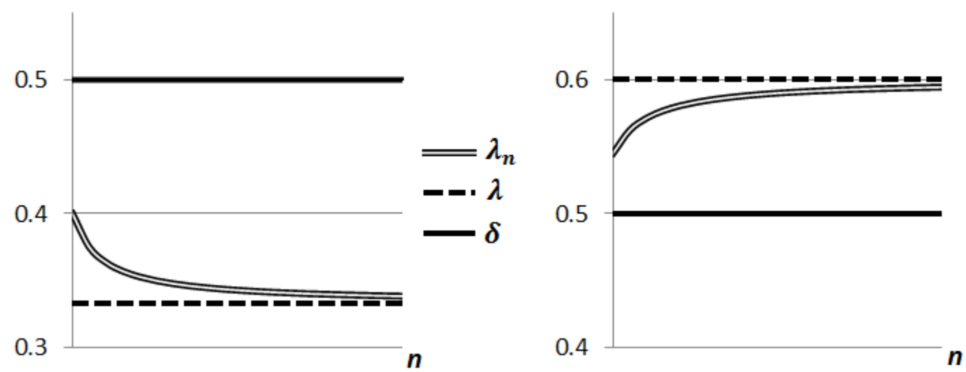

Figure 1. The sensitivity to private information when $r = 0.5$ (left) and $r = -0.5$ (right)

In the infinite-player game, the agents are infinitesimal so that each agent cannot influence the average forecast in which the private noises disappear as n goes to infinity.[10] However, with a finite number of agents, the average forecast is contaminated with the agents' private noises. Accordingly, when the agents want to herd ($r > 0$), the importance that any agent is inclined to attach to the average forecast is not entirely transferred to the agent's weight on public information, given the influence that any agent's private information can have on the average forecast. Hence, all agents strategically use more private information in equilibrium compared to the infinite-player game. The opposite happens when the agents want to move away from the herd ($r < 0$). The finiteness of the number of agents induces each of them to assign a higher weight on public information relative to the infinite-player game. These results can be interpreted as the weakening of the "pre-existent" herding effect on the equilibrium use of information due to the presence of market power.

[10]As n goes to infinity, the average error in the private signals of the agents converges to zero by the law of large numbers.

We now consider the equilibrium level of dispersion that is measured by the variation in the equilibrium forecasts across agents. The equilibrium forecast (3) can be rewritten as $a_i^* = \theta + \lambda_n(\alpha_x)^{-1/2}\varepsilon_i + (1-\lambda_n)(\alpha_p)^{-1/2}\varepsilon$. Therefore, the equilibrium level of forecast dispersion for any given realizations of θ and p is:

$$Var(a_i^*|\theta, p) = \left(\lambda_n(\alpha_x)^{-1/2}\right)^2 \tag{6}$$

The equilibrium level of forecast dispersion depends on the weight λ_n that incorporates the market power distortion. We examine the market power effect by showing how the number of agents affects the dispersion of forecasts.

Proposition 2. *For any given α_x, α_p, and r, the finite-player game's forecast dispersion falls with an increase in n if and only if $r > 0$; and increases with an increase in n if and only if $r < 0$.*

Proof. An observation of (6) yields $\frac{\partial Var(a_i^*|\theta,p)}{\partial n} \propto \frac{\partial \lambda_n}{\partial n}$. Then the proof follows from Lemma 2. □

Proposition 2 implies that the forecast dispersion in the finite-player game is higher when $r > 0$ but lower when $r < 0$ compared to the forecast dispersion in the infinite-player game given by $\left(\lambda(\alpha_x)^{-1/2}\right)^2$.[11] Public information is a relatively better predictor of the average forecast than private information. But the average forecast contains the agents' private noises given a finite number of agents. Hence, as the market power (measured in terms of $1/n$) increases, the agents who wish to herd around the average forecast find it optimal to rely less on public information, which leads to a higher disagreement among agents; when instead the agents want to move away from the average forecast, they find it optimal to rely more on public information, generating a lower disagreement among agents. In sum, the market power effect on forecast dispersion runs counter to the role of the underlying motive for herding or being distinctive in the forecasting game.

[11] In the simple-prediction game where $r = 0$, a change in the number of agents has no effect on the equilibrium level of forecast dispersion given by $\left(\delta(\alpha_x)^{-1/2}\right)^2$. Following from Lemma 2, the forecast dispersion in the finite-player game is lower (resp. higher) when $r > 0$ (resp. $r < 0$) as compared to the dispersion in the simple-prediction game.

4 Concluding Remarks

In this paper, we investigate the finite-player equilibrium properties for the class of forecasting games in which each forecaster cares both about being correct on the true fundamentals and about his distance to the average forecast. We find that the market power motive that is only present in the finite-player game causes each forecaster to behave strategically in a way that weakens the role played by the underlying herding motive in the game. Such consideration leads to the discrepancy between the forecast dispersion that should be observed in a setting with a small number of forecasters and the forecast dispersion observed in a setting with a large number of forecasters. The presence of market power contributes to a higher (resp. lower) forecast dispersion when the forecasters' preferences are for herding (resp. being distinctive). Our analysis offers a micro-theoretic framework for explaining forecast dispersion. For future work, a fruitful analysis is to consider heterogeneous agents, as well as to estimate the key parameters in our game by using survey data on n number of professional forecasters and simulate the model to match some empirical patterns of forecast dispersion observed in the data.

References

Acemoglu, Daron and Martin Kaae Jensen. 2013. "Aggregate Comparative Statics." *Games and Economic Behavior* 81:27–49.

Angeletos, George-Marios and Alessandro Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75(4):1103–1142.

Bergemann, Dirk, Tibor Heumann and Stephen Morris. 2015. "Information and Market Power." *Working Paper* .

Cornes, Richard and Roger Hartley. 2012. "Fully Aggregative Games." *Economics Letters* 116(3):631–633.

Croushore, Dean. 1997. "The Livingston Survey: Still Useful After All These Years." *Business Review - Federal Reserve Bank of Philadelphia* 2:1–12.

Martimort, David and Lars Stole. 2012. "Representing Equilibrium Aggregate in Aggregate Games with Applications to Common Agency." *Games and Economic Behavior* 76:753–772.

Morris, Stephen and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review* 92(5):1521–1534.

Morris, Stephen and Hyun Song Shin. 2003. Global Games: Theory and Applications. In *Advances in Economics and Econometrics: Theory and Applications*, ed. Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky. Vol. 1 Cambridge University Press pp. 57–114.

Patton, Andrew J. and Allan Timmermann. 2010. "Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion." *Journal of Monetary Economics* 57:803–820.