Information Percolation

Darrell Duffie
Stanford University

Gaston Giroux
unaffiliated

Gustavo Manso
Massachusetts Institute of Technology

October 28, 2008

1 Introduction

For a setting in which a large number of asymmetrically informed agents are randomly matched into groups over time, exchanging their information with each other when matched, we provide an explicit solution for the dynamics of the cross-sectional distribution of posterior beliefs. We also show that convergence of the cross-sectional distribution of beliefs to a common posterior is exponential and that the rate of convergence does not depend on the size of the groups of agents that meet. The rate of convergence is merely the mean rate at which an individual agent is matched.

For example, suppose that each agent has λ meetings per year, in expectation. At each meeting, say an auction, $m-1$ other agents are randomly selected to attend. Each agent at the meeting reveals to the others a summary statistic of his or her posterior, such as a bid for an asset, reflecting the agent's originally endowed information and any information learned from prior meetings. Over time, the conditional beliefs held across the population of agents regarding a variable of common concern (such as the payoff of the auctioned asset) converge to a common posterior. We construct an associated mathematical model of information transmission and calculate explicitly the cross-sectional distribution of the posterior beliefs held by the agents at each time. We show that convergence of these posteriors to a common posterior is exponential at the rate of λ, regardless of the number m of agents at each meeting.

An important role of markets and organizations, as argued for example by Hayek (1945) and Arrow (1974), is to facilitate the transmission of information that is dispersedly held by its participants. Our results suggest that varying the size of the groups in which individuals exchange information does not facilitate information transmission, at least in terms of the rate of convergence of posteriors. This point is further addressed at the end of the paper.

Previous studies have considered the problem of information aggregation in various contexts. For example, Grossman (1981) proposes the concept of rational-expectations equilibrium to capture the idea that prices aggregate information that is disperse in the economy. Wilson (1977), Milgrom (1981), Vives (1993), Pesendorfer and Swinkels (1997), and Reny and Perry (2006) provide strategic foundations for the rational-expectations equilibrium concept in centralized markets. In a number of important settings, however, agents learn from local interactions. For example, in over-the-counter markets, agents learn from the bids of other agents in privately held auctions. Wolinsky (1990) and Blouin and Serrano (2002) study information percolation in these markets.[1] In social learning settings, agents learn from direct interactions with other agents. Banerjee and Fudenberg (2004) study information percolation in a social learning context. In contrast to previous studies of learning through local interactions, we allow for meetings that have more than two agents, and we explicitly characterize the percolation of information and provide rates of convergence of the cross-sectional distribution of beliefs to a common posterior.

Our results extend those of Duffie and Manso (2007), who provided an explicit formula for the Fourier transform of the cross-sectional distribution of posterior beliefs in the same setting, but did not offer an explicit solution for the distribution itself, and did not characterize the rate of convergence of the distribution.

Section 2 provides the model setting. Section 3 provides our results for the traditional search-market setting of bilateral ($m = 2$) contacts. This also serves as an introduction to the results for the case of general m, which are presented in Section 4.

---

[1]Rubinstein and Wolinsky (1985) and Gale (1986a, 1986b) study decentralized markets without asymmetric information. Satterthwaite and Shneyerov (2003) study decentralized markets with private-value asymmetric information.

2 The Basic Model

The model of information percolation is that of Duffie and Manso (2007). A probability space $(\Omega, \mathcal{F}, \mathrm{P})$ and a "continuum" (a non-atomic finite measure space $(G, \mathcal{G}, \gamma)$) of agents are fixed. Without loss of generality, the total quantity $\gamma(G)$ of agents is 1. A random variable X of potential concern to all agents has two possible outcomes, H ("high") and L ("low"), with respective probabilities ν and $1 - \nu$.

Each agent is initially endowed with a sequence of signals that may be informative about X. The signals $\{s_1, \ldots, s_n\}$ primitively observed by a particular agent are, conditional on X, independent with outcomes 0 and 1 (Bernoulli trials). The number $n \geq 0$ of signals as well as the probability distributions of the signals may vary across agents. Without loss of generality, we suppose that $\mathrm{P}(s_i = 1 \mid H) \geq \mathrm{P}(s_i = 1 \mid L)$. A signal i is *informative* if $\mathrm{P}(s_i = 1 \mid H) > \mathrm{P}(s_i = 1 \mid L)$. For any pair of agents, their sets of originally endowed signals are independent.

By Bayes' rule, the logarithm of the likelihood ratio between states H and L conditional on signals $\{s_1, \ldots, s_n\}$ is

$$\log \frac{\mathrm{P}(X = H \mid s_1, \ldots, s_n)}{\mathrm{P}(X = L \mid s_1, \ldots, s_n)} = \log \frac{\nu}{1 - \nu} + \theta, \tag{1}$$

where the "type" θ of this set of signals is

$$\theta = \sum_{i=1}^{n} \log \frac{\mathrm{P}(s_i \mid H)}{\mathrm{P}(s_i \mid L)}. \tag{2}$$

The higher the type θ of the set of signals, the higher is the likelihood ratio between states H and L and the higher the posterior probability that X is high.

Any particular agent is matched to other agents at each of a sequence of Poisson arrival times with a mean arrival rate (intensity) λ, which is common across agents. At each meeting time, $m-1$ other agents are randomly selected from the population of agents.[2] The meeting group size m is a parameter

[2]That is, each of the $m - 1$ matched agents is chosen at random from the population, without replacement, with the uniform distribution, which we can take to be the agent-space measure γ. Duffie and Sun (2007) provide a complete construction for independent random matching from a large set (a non-atomic measure space) of agents, for the case $m = 2$.

of the information model that we shall vary. We assume that, for almost every pair of agents, the matching times and counterparties of one agent are independent of those of the other. We do not show the existence of such a random matching process.[3]

Suppose that whenever agents meet they communicate to each other their posterior probabilities, given all information to the point of that encounter, of the event that X is high. Duffie and Manso (2007) provide an example of a market setting in which this revelation of beliefs occurs through the observation of bids submitted by risk-neutral investors in an auction for a forward contract on an asset whose payoff is X. Using the same arguments as in Proposition 3 of Duffie and Manso (2007), we know that whenever an agent of type θ meets an agent with type ϕ and they communicate to each other their posterior distributions of X, they both attain the posterior type $\theta + \phi$. Moreover, whenever m agents of respective types $\phi_1, \ldots, \phi_m$ share their beliefs, they attain the common posterior type $\phi_1 + \cdots + \phi_m$.

We let μ_t denote the cross-sectional distribution of posterior types in the population at time t. That is, for any real interval (a, b), $\mu_t((a, b))$ (also denoted $\mu_t(a, b)$ for simplicity) is the fraction of the population whose type at time t is in (a, b). Because the total quantity $\gamma(G)$ of agents is 1, we can view μ_t as a probability distribution. The initial distribution μ_0 of types is that induced by some particular initial allocation of signals to agents. In the following analysis we assume that there is a positive mass of agents that has at least one informative signal. This implies that the first moment $m_1(\mu_0)$ is strictly positive if $X = H$, and that $m_1(\mu_0) < 0$ if $X = L$. We assume that the initial law μ_0 has a moment generating function, $z \mapsto \int e^{z\theta}\, \mu_0(d\theta)$, that is finite on a neighborhood of $z = 0$.

3 Two-Agent Meetings

We now calculate the explicit belief distribution in the population at any given time, and the rate of convergence of beliefs to a common posterior, in

[3]For the case of groups of size $m = 2$, Duffie and Sun (2007) show existence for the discrete-time analogue of this random matching model. For the case of a finite number of agent types, the associated exact law of large numbers for the cross-sectional distribution of the type processes is provided by Duffie and Sun (2005). Giroux (2005) proves convergence of the naturally associated finite-agent discrete-time model to the analogous continuous-time model matching model of Duffie, Gârleanu, and Pedersen (2005) as the number of agents grows large.

a setting with $m = 2$ agents at each meetings. This is the standard setting for search-based models of labor, money, and asset markets. In this setting, the cross-sectional distribution of types is determined by the evolution equation

$$\mu_t = \mu_0 + \lambda \int_0^t (\mu_s * \mu_s - \mu_s)\, ds, \tag{3}$$

where $*$ is the convolution operator. This is intuitively understood if μ_t has a density $f_t(\,\cdot\,)$, in which case the density $f_t(\theta)$ of agents of type θ is reduced at the rate $\lambda f_t(\theta)$ at which agents of type θ meet other agents and change type, and is increased at the aggregate rate $\lambda \int f_t(\theta - y) f_t(y)\, dy$ at which an agent of some type y meets an agent of type $\theta - y$, and therefore becomes an agent of type θ.

The following result provides an explicit solution for the cross-sectional type distribution, in the form of a Wild summation.[4]

Proposition 1 *The unique solution of (3) is*

$$\mu_t = \sum_{n \geq 1} e^{-\lambda t}(1 - e^{-\lambda t})^{n-1} \mu_0^{*n}, \tag{4}$$

where ρ^{*n} denotes the n-fold convolution of a measure ρ.

Proof As in Duffie and Manso (2007), we write the evolution equation (3) in terms of the Fourier transform $\varphi(\,\cdot\,, t)$ of μ_t, as

$$\frac{\partial \varphi(s,t)}{\partial t} = -\lambda \varphi(s,t) + \lambda \varphi^2(s,t), \tag{5}$$

with solution

$$\varphi(s,t) = \frac{\varphi(s,0)}{e^{\lambda t}(1 - \varphi(s,0)) + \varphi(s,0)}. \tag{6}$$

This solution can be expanded as

$$\varphi(s,t) = \sum_{n \geq 1} e^{-\lambda t}(1 - e^{-\lambda t})^{n-1} \varphi^n(s,0), \tag{7}$$

which is identical to the Fourier transform of the right-hand side of (4). ■

[4] See Wild (1951).

The Wild summation (4) shows that at each point in time the cross-sectional distribution of types is a mixture of convolutions of the initial distribution μ_0 of types. In the Wild summation (4), the term $e^{-\lambda t}(1-e^{-\lambda t})^{n-1}$ associated with the n-th convolution of μ_0 represents the fraction of agents that has been involved in $(n-1)$ direct or indirect meetings up to time t.[5]

This solution for the cross-sectional distribution of types is converted to an explicit distribution for the cross-sectional distribution π_t of posterior probabilities that $X = H$, using the fact that

$$\pi_t(0, b) = \mu_t\left(-\infty, \log\frac{b}{(1-b)} - \log\frac{\nu}{(1-\nu)}\right). \tag{8}$$

Like μ_t, the beliefs distribution π_t has an outcome that differs depending on whether $X = H$ or $X = L$.

We now provide explicit rates of convergence of the cross-sectional distribution of beliefs to a common posterior. In our setting, it turns out that all agents' beliefs converge to that of complete information, in that any agent's posterior probability of the event $\{X = H\}$ converges to 1 on this event and to zero otherwise. In general, we say that π_t converges to a common posterior distribution π_∞ if, almost surely, π_t converges in distribution to π_∞, and we say that convergence is exponential at the rate $\alpha > 0$ if there are constants κ_0 and κ_1 such that, for any b in $(0, 1)$,

$$e^{-\alpha t}\kappa_0 \leq |\pi_t(0, b) - \pi_\infty(0, b)| \leq e^{-\alpha t}\kappa_1.$$

Thus, if there is a rate of convergence, it is unique.

Proposition 2 *Convergence of the cross-sectional distribution of beliefs to that of complete information is exponential at the rate* λ.

Proof Because of (8), the rate of convergence of π_t is the same as the rate of convergence to zero or 1, for any a, of $\mu_t(-\infty, a)$. We will provide the rate of convergence to zero of $\mu_t(-\infty, a)$ on the event $X = H$. A like argument gives the same rate of convergence to 1 on the event $X = L$.

From (4),

$$\mu_t(-\infty, a) \geq e^{-\lambda t}\mu_0(-\infty, a). \tag{9}$$

[5] Agent A is involved in an indirect meeting with agent C if agent A is involved in a direct meeting with agent B after agent B has been involved in a (direct or indirect) meeting with agent C.

We fix n_0 such that $m_1(\mu_0) > a/n$ for $n > n_0$ and we let $\{Y_n\}_{n\geq 1}$ be independent random variables with distribution μ_0. Then,

$$\begin{aligned}\mu_t(-\infty, a) &= \sum_{n=1}^{n_0} e^{-\lambda t}(1-e^{-\lambda t})^{n-1}\,\mathrm{P}\left[\sum_{i=1}^{n}\left(Y_i - \frac{a}{n}\right) \leq 0\right] \\ &\quad + \sum_{n=n_0+1}^{\infty} e^{-\lambda t}(1-e^{-\lambda t})^{n-1}\,\mathrm{P}\left[\sum_{i=1}^{n}\left(Y_i - \frac{a}{n}\right) \leq 0\right]. \quad (10)\end{aligned}$$

It is clear that there exists a constant β such that, for all t, the first term on the right-hand side of equation (10) is less than $\beta e^{-\lambda t}$. Therefore, we only need to worry about the second term on the right-hand side of equation (10). From a standard result in probability theory,[6] if Y is a random variable with a finite strictly positive mean and a moment generating function that is finite on $(-c, 0]$ for some $c > 0$, then $\mathrm{P}(Y \leq 0) \leq \inf_{-c<s<0} E[e^{sY}] < 1$. For $n > n_0$, for some fixed $c > 0$, we then have that

$$\begin{aligned}\mathrm{P}\left[\sum_{i=1}^{n}\left(Y_i - \frac{a}{n}\right) \leq 0\right] &\leq \inf_{-c<s<0} E\left[e^{s(\sum_{i=1}^{n}(Y_i - a/n))}\right] \\ &= \left(\inf_{-c<s<0} E\left[e^{s(Y_1 - a/n)}\right]\right)^n \\ &\leq e^{ac}\left(\inf_{-c<s<0} E\left[e^{sY_1}\right]\right)^n \\ &\leq e^{ac}\gamma^n, \quad (11)\end{aligned}$$

with $\gamma < 1$. The first inequality comes from the standard result in probability theory stated above and the last inequality comes from the fact that Y has positive mean and a finite moment generating function.

From (11), we conclude that the second term on the right-hand side of equation (10) is bounded by $e^{ac}\frac{\gamma}{1-\gamma}e^{-\lambda t}$. Therefore,

$$\mu_0(-\infty, a)e^{-\lambda t} \leq \mu_t(-\infty, a) \leq \left(\beta + e^{ac}\frac{\gamma}{1-\gamma}\right)e^{-\lambda t}, \quad (12)$$

and the proof is complete. ■

[6] See, for example, Rosenthal (2000), pp. 90-92.

4 Multi-Agent Meetings

For the case of m agents at each meeting, the evolution of the cross-sectional distribution of types is similarly given by:

$$\mu_t = \mu_0 + \lambda \int_0^t (\mu_s^{*m} - \mu_s)\, ds, \tag{13}$$

as explained in Duffie and Manso (2007). A solution for the cross-sectional distribution of beliefs at any time t is given explicitly by (8) and the following extension of the Wild summation formula for the type distribution.

Proposition 3 *The unique solution of (13) is*

$$\mu_t = \sum_{n \geq 1} a_{[(m-1)(n-1)+1]} e^{-\lambda t} (1 - e^{-(m-1)\lambda t})^{n-1} \mu_0^{*[(m-1)(n-1)+1]}, \tag{14}$$

where $a_1 = 1$ *and, for* $n > 1$,

$$a_{(m-1)(n-1)+1} = \frac{1}{m-1} \left(1 - \sum_{\left\{ \substack{i_1,\ldots,i_{(m-1)} < n \\ \sum i_k = n+m-2} \right\}} \prod_{k=1}^{m-1} a_{[(m-1)(i_k - 1)+1]} \right). \tag{15}$$

Proof From (13), the Fourier transform of μ_t satisfies

$$\frac{\partial \varphi(s,t)}{\partial t} = -\lambda \varphi(s,t) + \lambda \varphi^m(s,t), \tag{16}$$

whose solution satisfies

$$\varphi(s,t)^{m-1} = \frac{\varphi(s,0)^{m-1}}{e^{(m-1)\lambda t}(1 - \varphi^{m-1}(s,0)) + \varphi^{m-1}(s,0)}. \tag{17}$$

Following steps analogous to those of Proposition 1,

$$\mu_t^{*(m-1)} = \sum_{n \geq 1} e^{-(m-1)\lambda t} (1 - e^{-(m-1)\lambda t})^{n-1} \mu_0^{*(m-1)n}. \tag{18}$$

Let ν_t denote the right-hand side of (14). By recursively calculating the convolution,

$$\begin{aligned}
\nu_t^{*(m-1)} &= \left(\sum_{n\geq 1} a_{[(m-1)(n-1)+1]} e^{-\lambda t}(1-e^{-(m-1)\lambda t})^{n-1}\mu_0^{*[(m-1)(n-1)+1]}\right)^{*(m-1)} \\
&= \sum_{n\geq 1} \beta_n e^{-(m-1)\lambda t}(1-e^{-(m-1)\lambda t})^{n-1}\mu_0^{*(m-1)n} \qquad (19)\\
&= \sum_{n\geq 1} e^{-(m-1)\lambda t}(1-e^{-(m-1)\lambda t})^{n-1}\mu_0^{*(m-1)n}, \qquad (20)
\end{aligned}$$

where

$$\beta_n = \sum_{\left\{\substack{i_1,\ldots,i_{(m-1)} \\ \sum i_k = n+m-2}\right\}} \prod_{k=1}^{m-1} a_{[(m-1)(i_k-1)+1]},$$

and where the last equality follows from the definition of $a_{[(m-1)(n-1)+1]}$ for $n \geq 1$. Thus, $\nu_t^{*(m-1)} = \mu_t^{*(m-1)}$, and it remains to show that the distribution μ_t is uniquely characterized by its convolution of order $m-1$. This follows[7] from the fact that μ_0, and therefore μ_t^{*k} for any t and k, have a moment generating function in a neighborhood of zero and a non-zero first moment on the event $\{X = H\}$. ■

Proposition 4 *For any meeting group size m, convergence of the cross-sectional distribution of beliefs to that of complete information is exponential at the rate λ.*

Proof Again, it is enough to derive the rate of convergence of $\mu_t(-\infty, a)$ to zero on the event $\{X = H\}$. From (14),

$$\mu_t(-\infty, a) \geq e^{-\lambda t}\mu_0(-\infty, a). \qquad (21)$$

[7]Because, on $\{X = H\}$, the derivative of the moment generating function of μ_0 at zero is the first moment of μ_0, which is positive, the moment generating function is strictly less than 1 in an interval $(-\epsilon, 0]$, for a sufficiently small $\epsilon > 0$. This implies that there is an analogous explicit solution for the moment generating function of μ_t^{*n}, for any n and t, on a small negative interval. The $(m-1)$-st root of the moment generating function of $\mu_t^{*(m-1)}$, on such an interval, uniquely determines the associated measure μ_t. For additional details, see Billingsley (1986), p. 408.

Now, from (18) and our analysis in Section 3, we know that for some constant $\kappa > 0$,

$$\mu_t^{*(m-1)}(-\infty, (m-1)a) \leq \kappa e^{-(m-1)\lambda t}. \tag{22}$$

From the fact that

$$(\mu_t(-\infty, a))^{m-1} \leq \mu_t^{*(m-1)}(-\infty, (m-1)a), \tag{23}$$

we conclude that

$$\mu_t(-\infty, a) \leq \kappa^{\frac{1}{m-1}} e^{-\lambda t}. \tag{24}$$

From (21) and (24), it follows that the rate of convergence of $\mu_t(-\infty, a)$ to zero is λ, completing the proof. ■

Because the expected rate at which a particular individual enters meetings is λ per year, independence and a formal application of the law of large numbers implies that the total quantity of m-agent meetings per year is λ/m, almost surely. So the total annual attendance at meetings is almost surely λ per year, invariant[8] to m. Our results show that total attendance at meetings is what matters for information convergence rates.

A simple calculation using equation (18) shows that the average number of signals observed (directly or indirectly) by an agent grows exponentially at rate $(m-1)\lambda$. This stands in contrast with Proposition 3, which shows that convergence to a common posterior is exponential at the rate λ, which is independent of meeting group size. The proof of Proposition 3 sheds some light on this issue. From equation (21), one can see that the rate of convergence when m agents meet is at most λ due to the first term in the Wild summation (14), which is associated with agents that have never met other agents. In our model, after some time has passed, most of the agents will be very well-informed and meeting only one such well-informed agent is likely to be enough to move an agent's beliefs close to the truth. Therefore, it is the agents who have not been involved in any meetings that are responsible

[8]This is not about large numbers, or uncertainty. For example, suppose each member of a group $\{A, B, C, D\}$ of 4 agents holds one meeting with a different member of the group. For example, A meets with B, and C meets with D. Then there is a total of two meetings, and each individual attends one meeting. If the 4 agents meet all together, once, we have the same total attendance, and the same rate at which each individual attends a meeting.

for slowing down convergence. From the Wild summation representation, the fraction of agents that have not been involved in any meetings up to time t is equal to $e^{-\lambda t}$, which is independent of meeting group size.

We have not shown that our invariance result extends from the case of a constant group size to a model in which the group size varies at random from meeting to meeting, say with the same mean group size across meetings. This is in fact the case under technical conditions on the distribution of group sizes, as related to us by Semyon Malamud in a subsequent private communication.

5 Market Example

In this section, we use our model to study information transmission in a decentralized market setting similar to that studied in Duffie and Manso (2007).

In this market example, uninformed buyers hedge the uncertainty in X, which is assumed to be revealed at some time $T > 0$. A continuum of risk-neutral sellers are initially endowed with signals about X, so that the initial cross-sectional distribution of their types is μ_0.

When an uninformed buyer arrives at the market he contacts m informed sellers randomly selected from the population and conducts a second-price auction, to determine the price at which he purchases a financial asset that pays 1 at time T if the true state of nature is H and 0 otherwise. After the purchase, the uninformed buyer leaves the market. Each informed seller participates in a sequence of these auctions with Poisson arrival times and mean arrival rate (intensity) λ. All bids that occur in an auction are observed by only the buyer and by the sellers participating in that auction. The discount rate is normalized to 1.

These second-price common-value auctions are known as "wallet games" and have been studied by Klemperer (1998). In the unique symmetric equilibrium of each auction, the sellers bid their posterior probabilities that X is high. Since there is a one-to-one mapping between type and posterior, informed sellers learn the types of the other sellers participating in the auction. The dynamics of information transmission is thus as described in Section 4.

References

Arrow, Kenneth. 1974. *The Limits of Organization.* Norton, New York.

Banerjee, Abhijit, and Drew Fudenberg. 2004. "Word-of-Mouth Learning." *Games and Economic Behavior*, **46:** 1-22.

Billingsley, Patrick. 1986. *Probability and Measure, Second Edition*, Wiley, New York.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. 2005. "Over-the-Counter Markets." *Econometrica*, **73:** 1815-1847.

Duffie, Darrell, and Gustavo Manso. 2007. "Information Percolation in Large Markets." *American Economic Review Papers and Proceedings*, **97:** 203-209.

Duffie, Darrell, and Yeneng Sun. 2007. "Existence of Independent Random Matching." *Annals of Applied Probability*, **17:** 386-419.

Duffie, Darrell, and Yeneng Sun. 2005. "The Exact Law of Large Numbers for Independent Random Matching." Stanford University Working Paper.

Gale, Douglas. 1986a. "Bargaining and Competition Part I: Characterization." *Econometrica*, **54:** 785-806.

Gale, Douglas. 1986b. "Bargaining and Competition Part II: Existence." *Econometrica*, **54:** 807-818.

Giroux, Gaston. 2005. "Markets of a Large Number of Interacting Agents." Working Paper (available from the author upon request).

Hayek, Friedrich. 1945. "The Use of Knowledge in Society." *American Economic Review*, **35:** 51 9-530.

Klemperer, Paul. 1998. "Auctions with Almost Common Values: The "Wallet Game" and its Applications." *European Economic Review*, **42:** 757–769.

Malamud, Semyon. 2007. "Private Communication of November 19, 2007," Department of Mathematics, ETH, Zurich.

Milgrom, Paul. 1981. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica*, **50:** 1089-1122.

Pesendorfer, Wolfgang and Jeroen Swinkels. 1997. "The Loser's Curse and Information Aggregation in Common Value Auctions." *Econometrica*, **65:** 1247-1281.

Reny, Philip and Motty Perry. 2006. "Toward a Strategic Foundation for Rational Expectations Equilibrium." *Econometrica*, **74:** 1231-1269.

Rosenthal, Jeffrey. 2000. *A First Look at Rigorous Probability Theory*. World Scientific. River Edge, New Jersey.

Rubinstein, Ariel and Asher Wolinsky. 1985. "Equilibrium in a Market With Sequential Bargaining." *Econometrica*, **53:** 1133-1150.

Satterthwaite, Mark, and Artyom Shneyerov. 2007. "Dynamic Matching, Two-sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition." *Econometrica* **75:** 155-200.

Vives, Xavier. 1993. "How Fast do Rational Agents Learn." *Review of Economic Studies*, **60:** 329-347.

Wild, E. 1951. "On Boltzmann's Equation in the Kinetic Theory of Gases," *Proceedings of the Cambridge Philosophical Society*, **47:** 602-609.

Wilson, Robert. 1977. "Incentive Efficiency of Double Auctions," *The Review of Economic Studies*, **44:** 511-518.