PIMS Lectures

Over the Counter Markets

Lecture 2: Search and Bargaining

Preliminary Draft

Darrell Duffie
Stanford University

Pacific Institute for the Mathematical Sciences
Summer School

The Economics and Mathematics of Systemic Risk and
Financial Networks

University of British Columbia, Vancouver, July, 2014

Preface

These notes are not comprehensive; they cover some of the issues concerning over-the-counter markets that I plan to address in PIMS Summer School lectures. I am aiming for the interests of an audience whose core consists of doctoral-level students who may wish to obtain a sense of the key technical modeling approaches as well as some sense of the substantive issues. I assume a graduate-level knowledge of probability theory. Some of this content is based on my Princeton lectures, titled Dark Markets.

Rather than trading through a centralized mechanism such as an auction, specialist, or broadly accessible limit-order book, participants in an over-the-counter (OTC) market negotiate terms privately with other market participants, often pairwise. OTC investors, other than major dealers, may be largely unaware of prices that are currently available elsewhere in the market, or of recent transactions prices. In this sense, OTC markets are relatively opaque; investors are somewhat in the dark about the most attractive available terms and about who might offer them. I will focus attention on search and bargaining with counterparties.

The financial crisis of 2007-2009 brought significant concerns and regulatory action regarding the role of over-the-counter markets, particularly from the viewpoint of financial instability. OTC markets for derivatives, collateralized debt obligations, and repurchase agreements played particularly important roles in the crisis and in subsequent legislation. The modeling of OTC markets, however, is still undeveloped by comparison with the available research on central market mechanisms.

Chapter 1 familiarizes readers with the basic techniques used to model search and random matching in economies with many agents. The exact law of large numbers for random matching, stated rigorously in Appendix A, is used to calculate the cross-sectional distribution of types of matches across the population. This is then extended to treat multi-period search in both discrete-time and continuous-time frameworks. The optimal search intensity of a given agent, given the cross-sectional distribution of types in the population, is formulated and characterized with Bellman's Principle. The chapter ends with a brief formulation of equilibrium search and a short review of the early history of the literature.

Chapter 2, from work by Duffie, Gârleanu, and Pedersen (2005) and Duffie, Gârleanu, and Pedersen (2007), presents a simple introduction to asset pricing in over-the-counter markets with symmetric information. Investors search for opportunities to trade and bargain with counterparties. Each of two negotiating investors is aware that a failure to complete a trade could lead to a potentially costly delay in order to search for a different counterparty. In equilibrium, the two investors agree to trade whenever there are potential gains from trade. The equilibrium asset price that they negotiate reflects the degree of search frictions, among other aspects of the market.

Appendix A provides needed results for dynamic random matching from Duffie and Sun (2007), Duffie and Sun (2012), and Duffie, Qiao, and Sun (2014b). Appendix B reviews the basics of counting processes with an intensity, such as Poisson processes. Appendix C covers the essentials of bargaining theory in settings related to OTC markets, with a focus on the alternating-offers bargaining protocol of Rubinstein (1982a).

Portions of these notes are updated from earlier notes prepared for the 2008 Nash Lecture, hosted by Steven Shreve, at Carnegie-Mellon University; for a doctoral course at the University of Lausanne in the summer of 2009; for the Distinguished Lecture Series at the Mathematics Department of Humboldt University hosted in Berlin by Ulrich Horst in June 2010; for the 2010 Tinbergen Lectures at the Duisenberg Institute, hosted in Amsterdam by André Lucas and Ton Vorst; and for the Minerva Foundation Lectures in the Mathematics Department at Columbia University in March, 2011, hosted by Ioannis Karatzas and Johannes Ruf. I am grateful for many discussions with students and faculty during my visits to present these lecture series.

I have a large debt is to many collaborators in this topic area: Adam

Ashcraft, Nicolae Gârleanu, Gaston Giroux, Semyon Malamud, Gustavo Manso, Lasse Heje Pedersen, Lei Qiao, Martin Scheicher, David Skeel, Bruno Strulovici, Yeneng Sun, Guillaume Vuillemy, Chaojun Wang, and Haoxiang Zhu. I am especially grateful to PIMS Director Alejandro Adem and to the Summer School organizers, René Carmona, Ivar Ekeland, and George Papanicolaou, for the opportunity to give these lectures.

This draft is for the use of those attending the PIMS Lectures, only, and is not otherwise for distribution.

Darrell Duffie

Stanford University
July, 2014

Chapter 1

Search for Counterparties

This chapter introduces the modeling of search and random matching in large over-the-counter markets. The objective is to build intuition and techniques for later chapters. After some mathematical prerequisites, the notion of random matching is defined. The law of large numbers is then invoked to calculate the cross-sectional distribution of types of matches. This is extended to multi-period search, first in discrete-time settings, and then in continuous time. The optimal search intensity of a given agent, given the cross-sectional distribution of types in the population, is characterized with Bellman's Principle. We then briefly take up the issue of equilibrium search efforts.

1.1 Preliminaries

We fix some mathematical preliminaries, beginning with a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The elements of Ω are the possible states of the world. The elements of $\mathcal{F}$ are events, sets of states to which we can assign a probability. The probability measure $\mathbb{P}$ assigns a probability in $[0, 1]$ to each event. We also fix a measure space $(G, \mathcal{G}, \gamma)$ of agents, so that $\gamma(B)$ is the quantity of agents in a measurable subset B of agents. The total quantity $\gamma(G)$ of agents is positive, but need not be 1.

We suppose that the measure γ is atomless, meaning that there is an infinite number of agents, none of which has a positive mass. The set of agents is therefore sometimes described as a "continuum." For example, agents could be uniformly distributed on the unit interval $G = [0, 1]$. Combining the continuum property with a notion of the independence of search across agents will lead in this chapter to an exact law of large numbers, by which the the cross-sectional distribution of search outcomes is deterministic (almost surely). For example, with two types of investors, A and B, we will see that independent random matching implies that the fraction of the population of type-A investors that are matched to type-B investors in a given period is almost surely equal to the probability that a given type-A investor is matched to some type-B investor.

Search delays are typical in over-the-counter markets, but also proxy for delays associated with reaching an awareness of trading opportunities, arranging financing and meeting suitable legal restrictions, negotiating trades, executing trades, and so on. As indicated in Chapter 1 and theoretically modeled in Chapters 4 and 5, these delays have important implications for optimal investment behavior, the dynamics of returns, and the distribution of information held across the population of investors.

1.2 Random Matching

In the simplest and most common models of random matching, a typical agent α, whom we shall call "Al," is randomly assigned to at most one other agent, and not to himself. If Al is assigned to a particular agent β, whom we shall call "Beth," then Beth is also assigned to Al. We suppose for now that the probability p of being matched to someone is the same for all agents, and that the probability that an agent is matched to some agent in a particular measurable subset B of agents is proportional to the quantity $\gamma(B)$ of agents in that subset. This is a natural implementation of the idea that all agents are "equally likely" to be Al's counterparties. Thus, the probability that Al gets matched to someone in the set B must be $p\gamma(B)/\gamma(G)$.

In order to later rely on a law of large numbers for the independent random matching of agents, we assume a notion of joint measurability of match assignments, as functions on $\Omega \times G$, that is stated in Appendix A. We will use the phrase "almost surely" to describe an event of probability one, and use "for almost every agent" to describe a relationship that applies to every agent in G, except those in some subset of measure zero.

We take the indicator random variable $1_{\alpha,\beta}$ to have the outcome 1 in the event that Al is matched to Beth, and zero otherwise. By adding up, the indicator of the event that Al is matched to someone in a measurable set B of agents is the random variable

$$1_{\alpha,B} = \int_{\beta \in B} 1_{\alpha,\beta} \, d\gamma(\beta).$$

The quantity of matches of agents in A to agents in B is then

$$\int_{\alpha \in A} 1_{\alpha,B} \, d\gamma(\alpha).$$

By interchanging expectation and summation over agents (joint measurability justifies this application of Fubini's Theorem), the expected quantity of these matches is

$$E\left[\int_{\alpha \in A} 1_{\alpha,B} \, d\gamma(\alpha)\right] = \int_{\alpha \in A} E(1_{\alpha,B}) \, d\gamma(\alpha) = \gamma(A) p \frac{\gamma(B)}{\gamma(G)}. \tag{1.1}$$

Similarly, the expected quantity of matches of agents in B to agents in A is $p\gamma(B)\gamma(A)/\gamma(G)$. Thus, if A and B are disjoint, the total expected quantity of matches between agents in B and agents in A is $2p\gamma(A)\gamma(B)/\gamma(G)$.

For now, we suppose there is a finite number K of types of agents. For a two-type example, suppose that 60% of the agents are potential buyers of an asset and the remaining 40% are potential sellers. The total quantity $\gamma(G)$ of agents is, say, 10. Only buyer-to-seller or seller-to-buyer matches result in trade. Suppose that the probability p that a given agent is matched to someone is 0.5. From (1.1), the expected quantity of buyer-to-buyer matches is $0.5 \times 6 \times 0.6 = 1.8$, the expected quantity of seller-to-seller matches is $0.5 \times 4 \times 0.4 = 0.8$, and the expected quantity of buyer-to-seller matches is $0.5 \times 6 \times 0.4 = 1.2$, which is equal to the quantity of seller-to-buyer matches. Thus, the total expected quantity of matches is $1.8 + 2 \times 1.2 + 0.8 = 5$, which is indeed the same as the total quantity of agents multiplied by the 50% probability that any agent is matched.

By "independent random matching," we mean that, for almost every agent α, the matching result for α (the event of remaining unmatched or the agent to whom α is matched) is independent of the matching result for β, for almost every agent β. An implication of independent random matching is that, for almost every agent α, the type of agent to whom α is matched is independent of the type to whom another agent β is matched, for almost every other agent β. This independence property will allow us to apply the exact law of large numbers for random matching, stated in Appendix A, to calculate the total quantity of matches of pairs of agents that are of given respective types.

The conventional law of large numbers is applied to a sequence $X_1, X_2, \ldots$ of independent random variables, all with the same probability distribution ν. By this law, for any measurable subset B of outcomes of X_i, the empirical fraction

$$\nu_n(B) = \frac{1}{n} \sum_{j=1}^{n} 1_{\{X_j \in B\}} \tag{1.2}$$

of outcomes in B converges with increasing sample size n to $\mathbb{P}(X_i \in B)$, almost surely. That is, the empirical distribution ν_n converges (almost surely) to the underlying probability distribution ν.

The *exact* law of large numbers treats a family $\{X_\alpha : \alpha \in G\}$ of random variables, one for each agent, satisfying measurability conditions provided in Appendix A. These random variables need not be identically distributed.

We are interested in characterizing the cross-sectional empirical distribution μ of these random variables, defined at some subset B of outcomes as

$$\mu(B) = \frac{1}{\gamma(G)} \int_G 1_{\{X_\alpha \in B\}} \, d\gamma(\alpha),$$

the fraction of the outcomes that are in B, which is analogous to (1.2). Absent independence assumptions, this fraction $\mu(B)$ is a non-trivial random variable. Under the independence and technical conditions of the exact law of large numbers of Sun (2006b) (see Appendix A), we have, almost surely,

$$\mu(B) = \frac{1}{\gamma(G)} \int_G \mathbb{P}(X_\alpha \in B) \, d\gamma(\alpha), \tag{1.3}$$

which is the cross-sectional average probability that X_α is in B. That is, with independence, the empirical distribution is almost surely the same as the average probability distribution.

When later modeling equilibrium investor behavior, our task is dramatically simplified if agents correctly assume that the empirical cross-sectional distribution of matches is not merely approximated by its probability distribution, but is actually *equal* to it. This is known as the exact law of large numbers for random matching, conditions for which are given by Duffie and Sun (2007) and Duffie and Sun (2012), and re-stated in Appendix A. By this exact law, letting A_i denote the subset of type-i agents, the quantity of matches of type-i agents to type-j agents is almost surely equal to the expected quantity, $p\gamma(A_i)\gamma(A_j)/\gamma(G)$.

Thus, in our previous example, with independent random matching, the quantity of buyer-to-seller matches is almost surely equal to the expected quantity, $0.5 \times 6 \times 0.4 = 1.2$.

1.3 Dynamic Search Models

In dynamic search models, random matching occurs period after period. In many applications, when agents meet, the matching activity could change the agent's types, perhaps randomly. For example, a prospective seller and a prospective buyer could meet at random and, if they successfully negotiate a trade, become a prospective buyer and prospective seller, respectively. This is the basis of the dynamics described in Chapter 2. Through their bids and offers, agents could also exchange information with each other when they meet, which changes their types with respect to posterior beliefs, as in Duffie, Malamud, and Manso (2010). Agents' types could also change for exogenous reasons, such as changes in preferences, exogenous investment opportunities, or new private information.

For now, we suppose that, at integer times, agents are randomly matched, as in the previous section, and that the probabilities of matching assignments and of exogenous type changes for each agent depend only on that agent's current type and on the cross-sectional distribution of types in the population. As before, we assume independence of these events across almost every pair of agents, as defined precisely in Appendix A.

For a warm-up illustrative calculation, suppose that there are two types of agents, buyers and sellers. Before entering the market in a given period, a change in preferences (or endowments, or trading constraints, for example) could, at random, cause a buyer to become a seller with probability 0.4, and a seller to become a buyer with probability 0.5. For almost every pair of traders, these exogenous changes are assumed to be independent. The exact law of large numbers implies that these exogenous changes, in the aggregate, cause 40% of the buyers to become sellers, and 50% of the sellers to become buyers, almost surely.

An agent is randomly matched to another with probability 0.2. Whenever a buyer and a seller are matched, they leave the market. They otherwise stay.

The initial quantities of buyers and sellers are b and s, respectively. In each period, after mutation and matching, the quantity of new buyers entering the market is assumed to be 10% of the quantity of buyers in the previous period plus 4. The quantity of new sellers entering is 20% of the previous-period quantity of sellers plus 2. These entries occur after trade. This example is contrived merely as an instructive numerical illustration.

An application of the exact law of large numbers implies that the new quantity of buyers is almost surely

$$b' = (0.6b + 0.5s) - 2 \times 0.2 \times (0.6b + 0.5s)\frac{0.4b + 0.5s}{b + s} + (0.1b + 4).$$

Similarly, the new quantity of sellers is almost surely

$$s' = (0.4b + 0.5s) - 2 \times 0.2 \times (0.4b + 0.5s)\frac{0.6b + 0.5s}{b + s} + (0.2s + 2).$$

The first-time reader should review each term in these expressions as a check on understanding.

One often simplifies with a "steady-state" model, in which the quantities b and s of buyers and sellers are stationary, that is, $b' = b$ and $s' = s$. In order for this to be the case, the net quantity $Q_b = -0.4b + 0.5s + 0.1b + 4$ of additional new buyers arising from exogenous type changes and fresh arrivals must be equal to the quantity of buyer departures caused by trade. Because this is also the case for sellers, and because the quantity of trades is of course the same for buyers and sellers, we have $Q_b = Q_s$. We therefore have the linear equation

$$-0.4b + 0.5s + 0.4b + 4 = 0.4b - 0.5s + 0.2s + 2,$$

which implies that

$$s = \frac{7}{8}b - \frac{5}{2}.$$

We can substitute this result for s into the steady-state equation

$$b = 0.6b + 0.5s - 2 \times 0.2 \times (0.6b + 0.5s)\frac{0.4b + 0.5s}{b + s} + 0.1b + 4,$$

and arrive at the quadratic equation

$$\frac{359}{2500}b^2 - \frac{46}{5}b + 12 = 0,$$

whose unique positive real solution is $b \simeq 62.73$. We then find that $s \simeq 52.39$. If the market starts with these quantities of buyers and sellers, then these quantities will persist at all future times, almost surely.

This notion of stationarity under dynamic random matching and mutation has been used for almost a century to model stability in population genetics, as discussed in Section 1.8, which provides a brief outline of the development of the literature. This approach became popular in economics, mainly in order to simplify modeling or study the effect of independence, in the latter half of the twentieth century.

1.4 Markov Chain for Type

We now model the evolution of the cross-sectional distribution of agents' types as a dynamic system, letting μ_{it} denote the fraction of the population that is of type i at period t.

For simplicity, we take the total quantity $\gamma(G)$ of agents to be 1, and assume no entry or exit. Simplifying from the previous example, we assume that at each period, mutually exclusively, a given agent is (i) matched, (ii) mutates type, or (iii) is unaffected.

Each period, agents of type i are matched with probability p_i, and are matched with equal likelihood to sets of agents of equal measure. Appendix A extends to the case of "directed search," by which the likelihood of a match with a type j agent is of the form $\theta_{int}\mu_{jt}$, for a per-capita matching rate θ_{ijt} that can vary, as opposed to the uniform case.

Under the uniform-matching-rate assumption, an agent of type i is therefore matched to an agent of type j with probability $p_i \mu_{jt}$. Immediately after such a match, we suppose that the type-i agent changes to type k with probability q_{ijk}. For instance, if agents are either owners ($i = 1$) or non-owners ($i = 2$) of an asset, and if non-owners and owners trade the asset whenever they meet (the non-owner becoming an owner, and vice versa), then $q_{122} = q_{211} = 1$, and all other q_{ijk} are 1 for $k = i$ and zero for $k \neq i$. The probability that a particular agent makes a one-period transition from type i to type k through matching is therefore $p_i \mu_{jt} q_{ijk}$.

A type-i agent becomes a type-k agent in the next period with an exogenous "mutation" probability Φ_{ik}. Thus, we have $p_i + \sum_k \Phi_{ik} = 1$. In order to apply the exact law of large numbers, we assume that these type changes are pairwise independent, as stated precisely in Appendix A. The parameters of the model are (p, q, Φ) and the initial cross-sectional distribution μ_0 of types.

A result stated in Appendix A implies that quantity of type-k agents satisfies (almost surely)

$$\mu_{k,t+1} = \sum_{i=1}^{K} \mu_{it} \left(\Phi_{ik} + p_i \sum_j \mu_{jt} q_{ijk} \right).$$

Letting $\mu_t = (\mu_{1t}, \ldots, \mu_{Kt})'$ denote the vector of fractions of each type of agent, we have (almost surely)

$$\mu_{t+1} = (\Phi + Q(p, q, \mu_t))\mu_t, \tag{1.4}$$

where $Q(p, q, \mu_t)$ is the $K \times K$ matrix whose (i, k)-element is $p_i \sum_j \mu_{jt} q_{ijk}$. Details are given by Duffie and Sun (2012).

We can similarly model the probability transitions of a particular agent's type. For any particular agent, the probability π_{it} that this agent is of type i at time t satisfies, as a vector,

$$\pi_{t+1} = (\Phi + Q(p, q, \mu_t))\pi_t. \tag{1.5}$$

This means that the agent's type is a Markov chain, with deterministic but time-varying transition probabilities that depend on the current cross-sectional type distribution μ_t.

The transition matrices of (1.4) and (1.5) are the same. Thus, if the agent's initial type is drawn at random from the initial cross-sectional type distribution μ_0, it follows that $\pi_t = \mu_t$ for all t (almost surely). That is, the probability distribution of the given agent's type at any time t is identical to the deterministic cross-sectional distribution of types at that time.

From (1.4), a stationary vector μ_∞ of quantities of agents satisfies the algebraic Riccati (linear-quadratic polynomial) equation

$$0 = (\Phi - I + Q(p, q, \mu_\infty))\mu_\infty. \tag{1.6}$$

Duffie and Sun (2012) show that a stationary equilibrium exists. The same equation characterizes a stationary probability distribution π_∞ of a given agent's type.

1.5 Continuous-Time Search and Matching

In many cases, calculations are simplified in a continuous-time setting. For this, we use the notion of an intensity process λ for the arrival of events of a particular type. The intensity λ_t of a given event is defined as the conditional mean arrival rate of the event given all of the information available up until time t. For example, an intensity of 2 means an expected arrival rate of 2 events per unit of time. The mathematical foundations are reviewed in Appendix B. A special case is a constant intensity λ, the Poisson-process model by which the times between arrivals are independent and exponentially distributed with mean $1/\lambda$.

Now, suppose that an agent of type i is randomly matched to other agents at a constant intensity λ_i. Taking our typical assumption that the selection of a counterparty is uniform across the population, the intensity of matches to agents of type j is $\lambda_i \mu_{jt}$. When such a match occurs, we suppose that the agent of type i becomes an agent of type k with probability q_{ijk}, as in the discrete-time model of the previous section. The type of an agent can also mutate from i to k for other reasons, at a fixed intensity of η_{ik}. For example, Chapter 4 discusses mutation over time of an investor's preferences for the asset or the investor's liquidity needs.

Assuming, for almost every pair of agents, that these type transitions are independent and that the exact law of large numbers applies, the quantity μ_{kt} of type-k agents satisfies (almost surely) the ordinary differential equation

$$\frac{d}{dt}\,\mu_{kt} = \sum_i \eta_{ik}\mu_{it} \, - \lambda_k \mu_{kt} \, + \sum_i \sum_j \lambda_i \mu_{it} \mu_{jt} q_{ijk}, \tag{1.7}$$

where we define $\eta_{kk} = -1$ to capture the expected rate of change of mutation out of type k. The dynamic equation for $\mu_t = (\mu_{1t}, \ldots, \mu_{Kt})'$ is thus

$$\frac{d}{dt}\,\mu_t = (\eta + Q(\lambda, q, \mu_t))\mu_t, \tag{1.8}$$

for the same matrix-valued function $Q(\,\cdot\,)$ defined in the previous section. This type of equation has long been relied upon, by assumption, in economics and physics. In particle physics, Boltzmann referred to this form of application by assumption of the continuous-time exact law of large numbers for random matching as the "Stosszahlansatz." A rigorous justification of (1.8) based on independent random matching is now available in Duffie, Qiao, and Sun (2014a). Ferland and Giroux (2008) have shown in some settings that the type distributions of discrete-time or finite-agent models converge to the solution of (1.8) as the number of agents converges to infinity and the length of a time period converges to zero.

The algebraic Riccati equation corresponding to the stationary quantities of each type is

$$0 = (\eta + Q(\lambda, q, \mu_\infty))\mu_\infty, \tag{1.9}$$

which is identical to the discrete-time equation (1.6), after replacing $\Phi - I$ and p with their continuous-time counterparts η and λ, respectively. The discrete-time model, however, was restricted by assuming that, at each time period, the events of being matched and of having an exogenous type change are mutually exclusive. Without such a restriction, the discrete-time model would be slightly more complicated than the corresponding continuous-time model.

1.6 Optimal Search

Continuing in this continuous-time framework, suppose that an agent of type i collects utility at the rate $u(i)$ whenever of type i, and generates an additional utility (or expected utility) of $w(i,j)$ when matched to an agent of type j. The dynamics of type changes are determined by the meeting intensities of agents and by the parameters η and q_{ijk} whose roles are explained in the previous section. As opposed to the previous section, however, each agent chooses some search intensity process λ. Search with intensity process λ generates costs at the rate $c(\lambda_t)$ at time t, for some continuous $c : [0,\infty) \to \mathbb{R}$. The search intensity process is bounded above by some constant $\overline{\lambda}$. The agent's search intensity process λ must be based only on information that the agent has available. More precisely, the intensity process is assumed to be predictable with respect to the agent's information filtration, as defined in Appendix B.

To simplify, suppose that the agent conjectures that the population cross-sectional type distribution μ is constant. The agent's type process associated with a chosen intensity process λ is denoted ϕ^λ.

For a discount rate $r > 0$, the agent's lifetime expected discounted utility is then

$$U(\lambda) = E\left(\int_0^\infty e^{-rt}[u(\phi_t^\lambda) - c(\lambda_t)]\,dt + \sum_{k=1}^\infty e^{-r\,T_k}\,w(\phi^\lambda(T_k-), \theta_k)\right),$$

where T_k is the time of the k-th match of that agent to some other agent, $\phi^\lambda(T_k-)$ is the type of the agent immediately before any type change occurs at time T_k, and θ_k is the type to whom the agent is matched at time T_k.

We are interested in solving the stochastic control problem

$$\sup_{\lambda} \; U(\lambda). \tag{1.10}$$

A search intensity process λ^* is optimal if it solves this problem, that is, if $U(\lambda^*) \geq U(\lambda)$ for all λ.

Letting $V = (V(1), \ldots, V(K))$ denote the supremum utilities associated with the respective types, the Hamilton-Jacobi-Bellman (HJB) equation for optimal choice of intensity is

$$0 = \sup_{\ell \in [\,0, \overline{\lambda}\,]} \; B(i, V, \ell), \quad i \in \{1, \ldots, K\}, \tag{1.11}$$

where

$$\begin{aligned} B(i, V, \ell) \;\; &= \;\; -rV(i) + u(i) - c(\ell) + \sum_{k=1}^{K} \eta_{ik}(V(k) - V(i)) \\ &\quad + \ell \; \sum_{j=1}^{K} \mu_j \left[w(i,j) + \sum_{k=1}^{K} q_{ijk}(V(k) - V(i)) \right]. \end{aligned}$$

There is a unique V solving the HJB equation (1.11). Fixing a solution V of the HJB equation, the continuity of $B(i, V, \ell)$ with respect to ℓ implies that the supremum defined by (1.11) is attained by some intensity level denoted by $\Lambda(i)$. We conjecture that optimality is achieved by an intensity process that has the outcome $\Lambda(i)$ whenever the agent is of type i. We let ϕ^* denote a type process for the agent with this property. The corresponding search intensity process λ^* is then defined by $\lambda^*_t = \Lambda(\phi^*_{t-})$. We thus conjecture that λ^* solves (1.10).

Proposition 1.1 *Problem (1.10) is solved by the search intensity process λ^*.*

This result follows from a standard verification argument, as follows. For an arbitrary search intensity process λ, let ϕ^λ be the associated type process and let

$$Y_t = e^{-rt}V(\phi_t^\lambda) + \int_0^t e^{-rs}[u(\phi_s^\lambda) - c(\lambda_s)]\,ds + \sum_{\{k\,:\,T_k \le t\}} e^{-rT_k}w(\phi_{T_k-}^\lambda, \theta_k).$$

A calculation shows that a martingale Z is defined by

$$Z_t = Y_t - \int_0^t e^{-rs}B(\phi_s^\lambda, V(\phi_s^\lambda), \lambda_s)\,ds. \tag{1.12}$$

To check that Z is indeed a martingale, we let N_t be the number of type changes the agent has experienced by time t. Proposition B.2 implies that a martingale $\hat{N}$ is defined by

$$\hat{N}_t = N_t - \int_0^t \left[\eta_{\varphi_s^\lambda,k} + \lambda_s \sum_j \mu_j q_{\varphi_s^\lambda,j,k}\right] ds.$$

The fact that Z is a martingale now follows from another application of Proposition B.2 and the fact that

$$dZ_t = dY_t - e^{-rt}B(\phi_t^\lambda, V(\phi_t^\lambda), \lambda_t)\,dt = H_t\,d\hat{N}_t,$$

where H is a bounded process that can be calculated.

From the HJB equation, $B(\phi_s^\lambda, V(\phi_s^\lambda), \lambda_s) \le 0$, so Y is a super-martingale. Thus, for an agent of initial type i, and for any time t,

$$V(i) = Y_0 \ge E(Y_t).$$

Because $e^{-rt}\max_i |V(i)|$ converges with t to zero, $E(Y_t) \to U(\lambda)$, and we have

$$V(\phi_0^\lambda) \ge U(\lambda).$$

For the particular case of $\lambda = \lambda^*$, the HJB equation (1.11) implies that $B(\phi_s^*, V(\phi_s^*), \lambda_s^*) = 0$, so Y is a martingale, and again taking a limit,

$$V(\phi_0^*) = U(\lambda^*).$$

Because $\phi_0^* = \phi_0^\lambda$ is an arbitrary initial type, $U(\lambda^*) \geq U(\lambda)$, proving optimality of λ^*, and confirming that $V(i)$ is the optimal utility $U(\lambda^*)$ of an agent of initial type i.

1.7 Equilibrium Search Intensities

Continuing in the setting of the previous section, an equilibrium is a cross-sectional distribution μ of types with the property that, when μ is taken as given by each agent, the optimal search intensities of agents are in aggregate consistent with μ. In order to formulate this precisely, suppose that for each cross-sectional type distribution μ, the dependence of an optimal search intensity policy function $\Lambda(\,\cdot\,)$, characterized in the last section, on the assumed cross-sectional distribution μ is indicated by writing $\Lambda^\mu = \Lambda$.

So, an equilibrium can be viewed as a solution μ (in the set Δ^{K-1} of non-negative vectors in $\mathbb{R}^K$ that sum to one) of the equilibrium equation

$$0 = (\eta + Q(\Lambda^\mu, q, \mu))\mu. \tag{1.13}$$

It would be enough for the existence of an equilibrium to have continuity of the map from a conjectured distribution $\nu \in \Delta^{K-1}$ to the corresponding solution μ of the stationary-measure equation

$$0 = (\eta + Q(\Lambda^\nu, q, \nu))\mu. \tag{1.14}$$

Because Δ^{K-1} is compact and convex, Schauder's Theorem would then imply at least one equilibrium.

The purpose of this section is merely to explain the notion of equilibrium search intensities. We do not apply this notion here. An example is the model of equilibrium search intensities of Duffie, Malamud, and Manso (2009).

Our reliance on the exact law of large numbers for random matching is evident in the formulation of an agent's conjectures about the equilibrium market environment. In the proposed equilibrium, the agent correctly conjectures a deterministic distribution μ of types in the population. If μ is only a limiting approximation of the distribution of types as the number of agents gets larger and larger, then it could be substantially more difficult to characterize the agent's optimal search policy in a particular finite-agent setting. Moreover, it would not be assured that as the actual cross-sectional distribution of types converges to the limit distribution μ, the agent's optimal policy converges to the optimal policy associated with the limit distribution. For both of these reasons, the exact law of large numbers drastically simplifies our modeling. This tractability, however, is achieved at a cost in realism. Real market environments are much more complicated than our simple model suggests.

1.8 Development of the Search Literature

Historically,[1] reliance on the exact law of large numbers for independent random matching dates back at least to 1908, when G.H. Hardy and W. Weinberg[2] independently proposed that random mating over time in a large population leads to constant and easily calculated fractions of each allele in the population. Hardy wrote: "suppose that the numbers are fairly large, so that the mating may be regarded as random," and then used, in effect, an exact law of large numbers for random matching to deduce his results. Consider, for illustration, a continuum population of gametes consisting of two alleles, A and B, in initial proportions p and $q = 1 - p$. Then, following the Hardy-Weinberg approach, the new population would have a fraction p^2

[1]These historical remarks are based in part on Duffie and Sun (2007).

[2]See Hardy (1908) and Cavalli-Sforza and Bodmer (1971).

whose parents are both of type A, a fraction q^2 whose parents are both of type B, and a fraction $2pq$ whose parents are of mixed type (heterozygotes). These genotypic proportions asserted by Hardy and Weinberg are already, implicitly, based on the exact law of large numbers for independent random matching in a large population.

In the field of economics, Hellwig (1976) is the first, to my knowledge, to have relied on the effect of the exact law of large numbers for random pairwise matching in a market, in a 1976 study of a monetary exchange economy. (Diamond (1971) had earlier assumed random matching of a large population with finitely many employers, but not pairwise matching among a continuum of agents.)

Since the 1970s, a large economics literature has routinely relied on an exact law of large numbers for independent random matching in a continuum population. This implicit use of this result occurs in general equilibrium theory (e.g. Gale (1986a), Gale (1986b), McLennan and Sonnenschein (1991), Wolinsky (1990)), game theory (e.g. Binmore and Samuelson (1999), Burdzy, Frankel, and Pauzner (2001), Dekel and Scotchmer (1999), Fudenberg and Levine (1993), Harrington (1998)), monetary theory (e.g. Diamond and Yellin (1990), Green and Zhou (2002), Hellwig (1976), Kiyotaki and Wright (1993)), labor economics (e.g. Diamond (1982), Hosios (1990), Mortensen (1982), Mortensen and Pissarides (1994)), and financial market theory, (e.g. Duffie, Gârleanu, and Pedersen (2007), Krainer and LeRoy (2002)).

In almost all of this literature, dynamics are crucial. For example, in the monetary and finance literature cited above, each agent in the economy solves a dynamic programming problem that is based on the conjectured dynamics of the cross-sectional distribution of agent types. An equilibrium has the property that the combined effect of individually optimal dynamic behavior is consistent with the conjectured population dynamics. In order to simplify the analysis, much of the literature relies on equilibria with a stationary distribution of agent types, as in the previous section.

Chapter 2

A Simple OTC Pricing Model

This chapter, based entirely on Duffie, Gârleanu, and Pedersen (2005) and Duffie, Gârleanu, and Pedersen (2007), presents a simple introduction to asset pricing in over-the-counter markets. Investors search for opportunities to trade and bargain with counterparties, each counterparty being aware that failure to conduct a trade could lead to a costly new search for a counterparty. In equilibrium, whenever there is gain from trade, the opportunity to search for a new counterparty is dominated by trading at the equilibrium asset price. The asset price reflects the degree of search frictions.

Under conditions, illiquidity premia are higher when counterparties are harder to find, when sellers have less bargaining power, when the fraction of qualified owners is smaller, and when risk aversion, volatility, or hedging demand is larger. Supply shocks cause prices to jump, and then "recover" over time, with a pattern that depends on the degree of search frictions.

We show how the equilibrium bargaining powers of the counterparties are determined by search opportunities, using the approach of Rubinstein and Wolinsky (1985). This approach has an axiomatic foundation based on Nash bargaining, as shown by Binmore, Rubinstein, and Wolinsky (1986), as discussed in Appendix C.

Here, traders have the same information. The case of OTC trading with asymmetric information is considered by Duffie, Malamud, and Manso

(2014).

2.1 Basic OTC Pricing Model

This section introduces a simple model of asset pricing in an over-the-counter market, with risk-neutral investors. The effects of risk aversion is considered in Duffie, Gârleanu, and Pedersen (2007).

We fix a non-atomic measure space of investors. Each investor is infinitely lived, with a constant time-preference rate $\beta > 0$ for consumption of a single non-storable numéraire good. A probability space and a common information filtration are also fixed. A cumulative consumption process C is one that can be represented as the difference between an increasing adapted process and a decreasing adapted process, with C_t denoting the total amount of consumption that has occurred through time t. The agent is restricted to a consumption process C whose utility

$$U(C) = E\left(\int_0^\infty e^{-\beta t}\, dC_t\right)$$

is well defined. This allows for positive or negative consumption, "smoothly" over time, or in sudden "lumps."

An agent can invest at any time in a liquid security with a risk-free interest rate of r. As a natural form of credit constraint, the agent must enforce some lower bound on the liquid wealth process W. (Otherwise, the agent could borrow without limit and get unbounded utility.) We take $r = \beta$ in this baseline model.

Agents may trade a long-lived asset in an over-the-counter (OTC) market in which trade may be negotiated bilaterally whenever two counterparties are matched. We begin for simplicity by taking the traded asset to be a consol, a bond that continually pays one unit of consumption per unit of time. We later allow random dividend processes in order to examine the effects of risk aversion.

An agent has an intrinsic preference for asset ownership that is "high" or "low." A low-type agent, when owning the asset, has an asset holding cost of δ per time unit. A high-type agent has no such holding cost. We could imagine this holding cost to be a shadow price for ownership due, for example, to a pressing need for cash or a relatively low personal use for the asset, as may happen for certain durable consumption goods. When we later allow for risk aversion, the low-type agent will be one whose endowments are adversely correlated with the asset dividends.

The agent's intrinsic type is a Markov chain, switching from low to high with intensity λ_u, and back to low with intensity λ_d. The intrinsic-type processes of almost every pair of agents are independent. These occasional preference shocks will generate incentives to trade because, in equilibrium, low-type owners want to sell and high-type non-owners want to buy.

The per-capita supply s of the asset is initially endowed to a subset of the agents. As a simplification, investors can hold at most one unit of the asset and cannot shortsell. This restriction is relaxed by Gârleanu (2009) and by Lagos and Rocheteau (2009). Because agents have linear utility, it is without much loss of generality that we restrict attention to equilibria in which, at any given time and state of the world, an agent holds either 0 or 1 unit of the asset. The set of $K = 4$ agent types is then $\mathcal{T} = \{ho, hn, lo, ln\}$, with the letters "$h$" and "$l$" designating the agent's current intrinsic preference state as high or low, respectively, and with "o" or "n" indicating whether the agent currently owns the asset or not, respectively.

We next consider a continuous-time search and bargaining framework adapted from Trejos and Wright (1995). We let $\mu_\sigma(t)$ denote the fraction at time t of agents of type $\sigma \in \mathcal{T}$, so that

$$1 = \mu_{ho}(t) + \mu_{hn}(t) + \mu_{lo}(t) + \mu_{ln}(t). \tag{2.1}$$

Equating the per-capita supply s with the fraction of owners gives

$$s = \mu_{ho}(t) + \mu_{lo}(t). \tag{2.2}$$

Any agent is matched to some counterparty with a constant intensity of λ, a parameter reflecting the efficiency of the market technology, and perhaps also reflecting individual inattention to trading. We assume that the counterparty found is randomly selected from the pool of other agents, so that the probability that the counterparty is of type σ is $\mu_\sigma(t)$. Thus, the total intensity of being matched to a type-σ investor at time t is $\lambda\mu_\sigma(t)$. Based on the Stosszollansatz outlined in Chapter 1, hn investors thus meet lo investors at an aggregate (almost sure) rate of $\lambda\mu_{lo}(t)\mu_{hn}(t)$.

In keeping with the modeling convention used in other chapters, we are departing here from the notion of contact intensity of Duffie, Gârleanu, and Pedersen (2005) and Duffie, Gârleanu, and Pedersen (2007), which measures the intensity with which an agent contacts other agents (in a transitive-verb sense) separately from the intensity with which other agents contact the agent in question. The total intensity of being matched is the sum of these two. Thus, the intensity parameter used by Duffie, Gârleanu, and Pedersen (2005) and Duffie, Gârleanu, and Pedersen (2007) is half of that used here.

To solve the model, we proceed in two steps. First, we exploit the fact that the only form of encounter that provides gains from trade is one in which low-type owners meet high-type non-owners. In any equilibrium of the bargaining game that is played at each such encounter, trade occurs immediately. We can therefore determine the asset allocations without reference to prices. Given the time-dynamics of the cross-sectional type distribution $\mu(t)$, we then consider equilibrium asset pricing.

In equilibrium, the rates of change of the fractions of the respective investor types satisfy the special case of (1.8) given by

$$\begin{aligned}
\dot\mu_{lo}(t) &= -\lambda\mu_{hn}(t)\mu_{lo}(t) - \lambda_u\mu_{lo}(t) + \lambda_d\mu_{ho}(t) \\
\dot\mu_{hn}(t) &= -\lambda\mu_{hn}(t)\mu_{lo}(t) - \lambda_d\mu_{hn}(t) + \lambda_u\mu_{ln}(t) \\
\dot\mu_{ho}(t) &= \lambda\mu_{hn}(t)\mu_{lo}(t) - \lambda_d\mu_{ho}(t) + \lambda_u\mu_{lo}(t) \\
\dot\mu_{ln}(t) &= \lambda\mu_{hn}(t)\mu_{lo}(t) - \lambda_u\mu_{ln}(t) + \lambda_d\mu_{hn}(t),
\end{aligned} \tag{2.3}$$

where $\dot{\mu}(t)$ denotes the time derivative of $\mu(t)$.

The intuition for, say, the first equation in (2.3) is straightforward: Whenever an *lo* agent meets an *hn* investor, he sells his asset and is no longer an *lo* agent. This explains the first term on the right-hand side of (2.3). The second term is due to intrinsic type changes in which *lo* investors become *ho* investors, and the last term is due to intrinsic type changes from *ho* to *lo*.

Duffie, Gârleanu, and Pedersen (2005) show that there is a unique stable stationary solution for $\{\mu(t) : t \geq 0\}$, that is, a constant solution defined by $\dot{\mu}(t) = 0$. The steady state is computed by using (2.1)-(2.2) and the fact that $\mu_{lo} + \mu_{ln} = \lambda_d/(\lambda_u + \lambda_d)$ in order to write the first equation in (2.3) as a quadratic equation in μ_{lo}.

Having determined the stationary fractions of investor types, we compute the investors' equilibrium intensities of finding counterparties of each type and, hence, their utilities for remaining lifetime consumption, as well as the bargained price P. The utility of a particular agent depends on the agent's current type, $\sigma(t) \in \mathcal{T}$, and the wealth W_t held in the liquid "bank-account" asset. Specifically, an agent's continuation utility is $W_t + V_{\sigma(t)}$, where, for each investor type σ in $\mathcal{T}$, V_σ is a constant to be determined.

In steady state, the Bellman principle implies that the rate of growth of any agent's expected indirect utility must be the discount rate r, which yields the steady-state equations

$$\begin{aligned} 0 &= rV_{lo} - \lambda_u(V_{ho} - V_{lo}) - \lambda\mu_{hn}(P - V_{lo} + V_{ln}) - (1 - \delta) \\ 0 &= rV_{ln} - \lambda_u(V_{hn} - V_{ln}) \\ 0 &= rV_{ho} + \lambda_d(V_{ho} - V_{lo}) - 1 \\ 0 &= rV_{hn} + \lambda_d(V_{hn} - V_{ln}) - \lambda\mu_{lo}(V_{ho} - V_{hn} - P). \end{aligned} \tag{2.4}$$

2.2 Bargaining over the Price

The asset price is determined through bilateral bargaining. A high-type nonowner pays at most his reservation value $\Delta V_h = V_{ho} - V_{hn}$ for obtaining the asset, while a low-type owner requires a price of at least $\Delta V_l = V_{lo} - V_{ln}$. In any equilibrium of the bargaining game, trade must occur at an in-between price of the form

$$P = \Delta V_l (1-q) + \Delta V_h \, q \,, \tag{2.5}$$

where $q \in [0,1]$ is called the "bargaining power" of the seller. Because we are characterizing stationary equilibrium, we take the bargaining power q to be constant.

While a Nash equilibrium in the bargaining game is consistent with any exogenously assumed bargaining power, we can use the device of Rubinstein and Wolinsky (1985) to calculate the unique bargaining power that represents the limiting price of a sequence of economies in which, once a pair of counterparties meets to negotiate, one of the pair is selected at random to make an offer to the other, at each of a sequence of offer times separated by intervals that shrink to zero.

Specifically, suppose that when an owner who wishes to sell and a non-owner who wishes to buy find each other, one of them is chosen randomly, the seller with probability $\hat{q}$ and the buyer with probability $1 - \hat{q}$, to suggest a trading price. The counterparty either rejects or accepts the offer, immediately. If the offer is rejected, the owner receives the dividend from the asset during the current period. At the next period, T later, one of the two agents is chosen at random, independently, to make a new offer. The bargaining may, however, break down before a counteroffer is made. A breakdown may occur because, during the interim, at least one of the agents may change his intrinsic valuation type, or one of the agents may meet yet another agent and leave his or her current trading partner. (The opportunity to continue to search for alternative counterparties while engaged in negotiation will also be considered below.)

This bargaining setting is a slight extension of our basic model, in that once a pair of agents meet, they are given the opportunity to interact at discretely separated moments in time, T apart. Later, we return to our original continuous-time framework by letting T go to zero, and adopt the limiting behavior of their bargaining game as $T \to 0$.

We consider first the case in which agents can search for alternative counterparties during their bargaining encounter. We assume that, given contact with an alternative partner, they leave the present partner in order to negotiate with the newly found one. The offerer suggests the price that leaves the other agent indifferent between accepting and rejecting it. In the unique subgame perfect equilibrium, the offer is accepted immediately, as shown by Rubinstein (1982b). The value of rejecting is that associated with the assumption by agents that the equilibrium strategies are to be played from then onwards. Letting P_σ be the price suggested by the agent of type $\sigma \in \{lo, hn\}$, letting $\bar{P} = \hat{q}P_{lo} + (1-\hat{q})P_{hn}$, and making use of the dynamic equations governing V_{lo} and V_{hn}, we have

$$P_{hn} - \Delta V_l = e^{-(r+\lambda_d+\lambda_u+\lambda\mu_{lo}+\lambda\mu_{hn})T}(\bar{P} - \Delta V_l) + O(T^2) \tag{2.6}$$

$$-P_{lo} + \Delta V_h = e^{-(r+\lambda_d+\lambda_u+\lambda\mu_{lo}+\lambda\mu_{hn})T}(-\bar{P} + \Delta V_h) + O(T^2)\,. \tag{2.7}$$

These prices, P_{hn} and P_{lo}, have the same limit $P = \lim_{T\to 0} P_{hn} = \lim_{T\to 0} P_{lo}$. The limit price P and a limit type-dependent value V_σ satisfy

$$P = \Delta V_l\,(1-q) + \Delta V_h\,q, \tag{2.8}$$

with

$$q = \hat{q}. \tag{2.9}$$

Thus, the limiting bargaining power $q = \hat{q}$ does not depend on the model parameters, beyond the likelihood that the seller is chosen to make an offer. In particular, an agent's intensity of meeting other trading partners does not influence q. This is because one's own ability to meet an alternative trading partner makes oneself more impatient, and also increases the partner's risk of breakdown. These two effects happen to cancel each other.

Other bargaining procedures lead to other outcomes. For instance, if agents are unable to search for alternative trading partners during negotiation, then, as shown by Duffie, Gârleanu, and Pedersen (2005),

$$q = \frac{\hat{q}(r + \lambda_u + \lambda_d + \lambda\mu_{lo})}{\hat{q}(r + \lambda_u + \lambda_d + \lambda\mu_{lo}) + (1 - \hat{q})(r + \lambda_u + \lambda_d + \lambda\mu_{hn})}. \tag{2.10}$$

The linear system of equations defined by (2.4)-(2.5) has a unique solution, with

$$P = \frac{1}{r} - \frac{\delta}{r} \frac{r(1-q) + \lambda_d + \lambda\mu_{lo}(1-q)}{r + \lambda_d + \lambda\mu_{lo}(1-q) + \lambda_u + \lambda\mu_{hn}q}. \tag{2.11}$$

This price (2.11) is the present value $1/r$ of dividends, reduced by an illiquidity discount. The discount is larger (other effects held constant) if the distressed owner has less hope of switching type (lower λ_u), if the quantity μ_{hn} of other buyers to be found is smaller, if the buyer may more suddenly need liquidity himself (higher λ_d), if it is easier for the buyer to find other sellers (higher μ_{lo}), or if the seller has less bargaining power (lower q).

These intuitive results are based on partial derivatives of the right-hand side of (2.11). In other words, they hold when a parameter changes without influencing any of the others. It is the case, however, that the steady-state type fractions μ themselves depend on λ_d, λ_u, and λ, an equilibrium effect that must also be considered. The following proposition offers a characterization of the equilibrium steady-state effect of changing each parameter.

Proposition 2.1 *The steady-state equilibrium price P is decreasing in δ, s, and λ_d, and is increasing in λ_u and q. Further, if $s < \lambda_u/(\lambda_u + \lambda_d)$, then $P \to 1/r$ as $\lambda \to \infty$, and P is increasing in λ for all $\lambda \geq \bar{\lambda}$, for a constant $\bar{\lambda}$ depending on the other parameters of the model.*

The condition that $s < \lambda_u/(\lambda_u + \lambda_d)$ means that, in steady state, there is less than one unit of asset per agent of high intrinsic type. Under this condition, the Walrasian frictionless price is equal to the present value of dividends $1/r$ because the marginal owner is always a high-type agent who incurs no holding costs. Naturally, as the search intensity increases towards infinity and frictions vanish, the OTC price approaches the Walrasian price (that is, the liquidity discount vanishes). The proposition also states that the price decreases with the ratio s of assets to qualified owners, with reductions in the mean arrival rate λ_d of a liquidity shock, and with increases in the speed at which agents can "recover" by becoming of high type again. It can easily be seen that if agents can easily recover (that is, as $\lambda_u \to \infty$), the price also approaches the Walrasian price.

While the proposition above captures the intuitively anticipated increase in market value with increasing search intensity λ, the alternative is also possible. With $s > \lambda_u/(\lambda_u+\lambda_d)$, the marginal investor in perfect markets has the relatively lower reservation value, and search frictions lead to a "scarcity value." For example, a high-type investor in an illiquid OTC market could pay more than the Walrasian price for the asset because it is hard to find, and given no opportunity to exploit the effect of immediate competition among many sellers. This scarcity value could, for example, contribute to the widely studied on-the-run premium for Treasuries, or to the elevation of prices of bonds that are difficult to find for physical settlement of credit derivatives or futures contracts. Absent search delays, it is difficult to explain these pricing phenomena.

Appendix A

Foundations for Random Matching

This appendix summarizes the results of Duffie, Qiao, and Sun (2014b) providing for an exact law of large numbers for random matching of a "continuum" of investors in a static setting. The results generalize those suggested in Chapter 1, based on Duffie and Sun (2007), by allowing for directed search.

A.1 Mathematical Preliminaries

We fix an atomless probability space $(I, \mathcal{I}, \lambda)$ representing the space of agents and a sample probability space $(\Omega, \mathcal{F}, P)$ representing the states of the world, and we let $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes P)$ be a Fubini extension[1] of the usual product probability space. This Fubini extension includes a sufficiently rich collection

[1]A formal definition of Fubini extension was introduced by (Sun 2006a). A probability space $(I \times \Omega, \mathcal{W}, Q)$ extending the usual product space $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$ is said to be a *Fubini extension* of $(I \times \Omega, \mathcal{I} \otimes \mathcal{F}, \lambda \otimes P)$ if for any real-valued Q-integrable function g on $(I \times \Omega, \mathcal{W})$, the functions $g_i = g(i, \cdot)$ and $g_\omega = g(\cdot, \omega)$ are integrable respectively on $(\Omega, \mathcal{F}, P)$ for λ-almost all $i \in I$ and on $(I, \mathcal{I}, \lambda)$ for P-almost all $\omega \in \Omega$; and if, moreover, $\int_\Omega g_i \, dP$ and $\int_I g_\omega \, d\lambda$ are integrable, respectively, on $(I, \mathcal{I}, \lambda)$ and on $(\Omega, \mathcal{F}, P)$, with $\int_{I \times \Omega} g \, dQ = \int_I \left(\int_\Omega g_i \, dP \right) d\lambda = \int_\Omega \left(\int_I g_\omega \, d\lambda \right) dP$. To reflect the fact that the probability space $(I \times \Omega, \mathcal{W}, Q)$ has $(I, \mathcal{I}, \lambda)$ and $(\Omega, \mathcal{F}, P)$ as its marginal spaces, as required by the Fubini property, it is denoted by $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes P)$.

of measurable sets to allow applications of the exact law of large numbers that we shall need.

We begin with a static model of directed random matching, and then a dynamic model that incorporates random changes over time in agents' types that are caused by matching and mutation.

A.2 The static model

Let S be a finite or countably infinite agent type space and $\alpha : I \to S$ be a measurable type function, mapping individual agents to their types. For any k in S, we let $p_k = \lambda(\{i : \alpha(i) = k\}$ denote the fraction of agents that are of type k. We can view $(p_k)_{k \in S}$ as an element of the space Δ of probability measures on S.

A function $\theta : S \times S \to \mathbb{R}_+$ is a matching rate function for the type distribution p if $\theta_{kl} = \theta_{lk}$ for any k and l in S, and if $\sum_{l \in S} p_l \theta_{kl} \leq 1$ for each $k \in S$. The matching rate θ_{kl} specifies the "per-capita" rate of matching of agents of type k with agents of type l, in the sense that $q_{kl} = p_l \theta_{kl}$ is the probability that a given agent of type k is matched to an agent of type l. Thus, $q_k = 1 - \sum_{l \in S} p_l \theta_{kl}$ is the associated probability of no matching for an agent of type k.

A mapping π from $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes P)$ to I is defined to be a random matching if it satisfies the two conditions:

(i) For each $\omega \in \Omega$, π_ω is a bijection between I and itself. Letting $B^\omega = \{i \in I : \pi_\omega(i) = i\}$ denote the agents that not matched by π_ω to a distinct agent, the restriction of π_ω to $I \backslash B^\omega$ is one-to-one and satisfies $\pi_\omega(\pi_\omega(i)) = i$.

(ii) Letting J denote the event of no matching, an $\mathcal{I} \boxtimes \mathcal{F}$-measurable type assignment function g for π is defined by

$$g(i,\omega) = \begin{cases} \alpha(\pi(i,\omega)), & i \notin B^\omega \\ J, & i \in B^\omega. \end{cases}$$

We say that a random matching π with type assignment function g has parameters (p, θ) if, for λ-almost every agent $i \in I$ of type k, we have $P(g_i = J) = q_k$ and $P(g_i = l) = q_{kl}$, where g_i denotes the random variable $g(i, \cdot)$.

The following is a direct application of the exact law of large numbers. We say that π is pairwise independent in types if its type assignment function g is essentially pairwise independent.[2]

Proposition 1 *Let π be a random matching with type assignment function g and parameters (p, θ). If π is pairwise independent in types then, for P-almost every $\omega \in \Omega$:*

(i) $\lambda(\{i \in I : \alpha(i) = k, g_\omega(i) = J\}) = p_k q_k$.

(ii) For any $(k,l) \in S^2$, $\lambda(\{i : \alpha(i) = k, g_\omega(i) = l\}) = p_k q_{kl} = p_k \theta_{kl} p_l$.

Proposition 2 *For any type distribution p on S and any matching rate function θ for p, there exists a random matching π with parameters (p, θ) that is essentially pairwise independent in types.*

A.3 Dynamic directed random matching

In this section, we consider a dynamical system with random mutation, random matching with directed probabilities, and match-induced random type

[2] An $\mathcal{I} \boxtimes \mathcal{F}$-measurable process f from $I \times \Omega$ to a complete separable metric space X is said to be essentially pairwise independent if for λ-almost all $i \in I$, the random variables f_i and f_j are independent for λ-almost all $j \in I$. Two random variables ϕ and ψ from $(\Omega, \mathcal{F}, P)$ to X are said to be independent if the σ-algebras $\sigma(\phi)$ and $\sigma(\psi)$ generated respectively by ϕ and ψ are independent.

changes. We also allow for time-dependent parameters. We first define such a dynamical system. Then we formulate the key property of being Markovian and conditional independence in types. We then state a result providing for the existence and an exact law of large numbers for such a dynamical system. For time-independent parameters and with finitely many types, we also characterize stationarity.

A.3.1 Definitions for dynamic random matching

Here, we define a discrete-time random process for agent types with the property that at each integer time period $n \geq 1$, agents first experience a random mutation and then a random matching with directed probability. Finally, any pair of matched agents are randomly assigned new types whose probabilities depend on the prior types of the two agents in a manner to be defined.

A random type function is a measurable mapping from $(I \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, \lambda \boxtimes P)$ to S. The initial random type function α^0 is assumed to be essentially pairwise independent with a cross-sectional type distribution p^0 defined by $p^0_k = \lambda(\{i : \alpha^0_i = k\})$.

We will characterize a dynamical system with, at each period n, a random type function h^n, assigning to agent i the type h^n_i after mutation but before matching. At period n after matching, the types of agents are likewise specified by a random type function α^n. A key objective is to show the existence and properties of a random type process $(h, \alpha) = \{(h^1, \alpha^1), (h^2, \alpha^2), \ldots\}$ that respects specified properties and parameters for mutation, directed random matching, and match-induced random type change.

At period n, before random matching, any agent of type k mutates so as to become an agent of type l with some specified probability b^n_{kl}, where for each k, $(b^n_{k1}, b^n_{k2}, \ldots)$ is in Δ. We thus require that for λ almost-every agent i,

$$P(h^n_i = l \,|\, \alpha^{n-1}_i = k) = b^n_{kl}. \tag{A.1}$$

For $n \geq 1$ and for each $(k,l) \in S^2$, let θ^n_{kl} be a continuous function on Δ into $\mathbb{R}_+$ with the property that, for all k and all p in Δ,

$$\sum_{l \in S} \theta^n_{kl}(p) p_l \leq 1.$$

An agent of type k is matched at period n to an agent with type l at the per-capita matching rate $\theta^n_{kl}(\hat{p}^n)$, where $\hat{p}^n$ is the type distribution of h^n, defined by $\hat{p}^n_k = \lambda(\{i : h^n_i = k\})$.

When an agent of type k is matched at time n to an agent of type l, the agent of type k becomes an agent of type r with a specified probability $\nu^n_{kl}(r)$.

The parameters of the model are (p^0, b, θ, ν).

At each period n, agents are to be matched according to a random matching π^n with a type assignment function g^n that respects the property that for every type k and λ-almost every agent i,

$$P(g^n_i = l \,|\, h^n_i = k) = \hat{q}^n_{kl} \triangleq \theta^n_{kl}(\hat{p}^n)\hat{p}^n_l, \tag{A.2}$$

and

$$P(g^n_i = J \,|\, h^n_i = k) = \hat{q}^n_k = 1 - \sum_{l=1}^{\infty} \hat{q}^n_{kl}. \tag{A.3}$$

We also require that the type function α^n_i after match-induced type changes satisfies, for λ-almost all agent $i \in I$,

$$P(\alpha^n_i = r \,|\, h^n_i = k, g^n_i = J) = \delta_k(r) \tag{A.4}$$

and

$$P(\alpha^n_i = r \,|\, h^n_i = k, g^n_i = l) = \nu^n_{kl}(r), \tag{A.5}$$

where $\delta_k(r)$ is one if $r = k$, and zero otherwise.

For any given model parameters (p^0, b, θ, ν), by induction in the period n, we will rely on Duffie, Qiao, and Sun (2014b) for the existence of (h, α, π, g), determining agent types with random mutation, random matching with directed probability, and match-induced type changing, respecting the definitional transition probabilities (A.1)-(A.4)-(A.5) and matching probabilities (A.2). In this case, we say that $\mathbb{D} = (h, \alpha, \pi, g)$ is a dynamical system with parameters (p^0, b, θ, ν). Under the assumption that random mutation, matching, and match-induced type changes are essentially pairwise independent, an application of the exact law of large numbers implies that the cross sectional type distributions are almost surely deterministic, a property that is frequently used in applications.

A.3.2 Markov conditional independence in types

In this section we define Markovian and cross-sectional independence properties for a dynamical system $\mathbb{D} = (h, \alpha, \pi, g)$. The idea of the property is that each agent's type process is Markovian and, moreover, that the random mutation, random matching, and match-induced type changes that occur in any period n are probabilistically independent across almost all agents.

We say that $\mathbb{D}$ has random mutation that is Markovian and conditionally independent in types if, for λ-almost all $i \in I$ and λ-almost all $j \in I$,

$$P(h_i^n = k, h_j^n = l \mid \alpha_i^0, \ldots, \alpha_i^{n-1}; \alpha_j^0, \ldots, \alpha_j^{n-1}) = P(h_i^n = k \mid \alpha_i^{n-1}) P(h_j^n = l \mid \alpha_l^{n-1}),$$

for every period n and for all types k and l in S.

We say that $\mathbb{D}$ has random matching that is Markovian and conditionally independent in types if, for λ-almost all $i \in I$ and λ-almost all $j \in I$,

$$P(g_i^n = c, g_j^n = d \mid \alpha_i^0, \ldots, \alpha_i^{n-1}, h_i^n; \alpha_j^0, \ldots, \alpha_j^{n-1}, h_j^n) = P(g_i^n = c \mid h_i^n) P(g_j^n = d \mid h_j^n)$$

for every period n and for all c and d in $S \cup \{J\}$.

We say that $\mathbb{D}$ has match-induced random type change that is Markovian and conditionally independent in types if for λ-almost all $i \in I$, and λ-almost all $j \in I$,

$$P(\alpha_i^n = c, \alpha_j^n = d \,|\, \alpha_i^0, \dots, \alpha_i^{n-1}, h_i^n, g_i^n; \alpha_j^0, \dots, \alpha_j^{n-1}, h_j^n, g_j^n) = P(\alpha_i^n = c \,|\, h_i^n, g_i^n) P(\alpha_j^n = d \,|\, h_j^n, g_j^n)$$

for every period n and for all k and l in S.

Finally, we say that $\mathbb{D}$ is Markovian and conditionally independent in types if its random mutation, random matching, and match-induced type change is Markovian and independent in types.

Our main result, from Duffie, Qiao, and Sun (2014b), is the following.

Proposition 3 *For any parameters (p^0, b, θ, ν), there exists a dynamical system $\mathbb{D} = (h, \alpha, \pi, g)$ with these parameters that is Markovian and conditionally independent in types.*

A.3.3 Exact law of large numbers and stationarity

We now define a sequence Γ^n of mappings from Δ to Δ such that, for each $p = (p_1, \dots, p_k, \dots)$ in Δ,

$$\Gamma_r^n(p_1, \dots, p_k, \dots) = \tilde{q}_r^n \sum_{l=1}^{\infty} p_l b_{lr}^n + \sum_{k,l=1}^{\infty} \tilde{q}_{kl}^n \nu_{kl}^n(r) \tilde{p}_k,$$

where $\tilde{p}_k = \sum_{k=1}^{\infty} p_l b_{lk}^n$, $\tilde{q}_{kl}^n = \theta_{kl}^n(\tilde{p})\tilde{p}_l$ and $\tilde{q}_k^n = 1 - \sum_{l=1}^{\infty} \tilde{q}_{kl}^n$.

The following proposition from Duffie, Qiao, and Sun (2014b) provides an exact law of large numbers for agent type processes allowing for random mutation, random matching with directed probability, and match-induced random type changing that is Markov conditionally independent in types. The proposition also gives a recursive calculation of the the cross-sectional type distribution p^n.

Proposition 4 *If $\mathbb{D} = (h, \alpha, \pi, g)$ is a dynamical system with parameters (p^0, b, θ, ν) that is Markovian and conditionally independent in types, then:*

(1) *For each time* $n \geq 1$*, the expectation* $\bar{p}^n = E(p^n)$ *of the cross-sectional type distribution is given by*

$$\bar{p}^n_r = \Gamma^n_r(\bar{p}^{n-1}) = \tilde{q}^n_r \sum_{l=1}^{\infty} \bar{p}^{n-1}_l b^n_{lr} + \sum_{k,l=1}^{\infty} \tilde{q}^n_{kl} \nu^n_{kl}(r) \tilde{p}^n_k,$$

where $\tilde{p}^n_k = \sum_{l=1}^{\infty} b^n_{lk} \bar{p}^{n-1}_l$, $\tilde{q}^n_{kl} = \theta^n_{kl}(\tilde{p}^n) \tilde{p}^n_l$ *and* $\tilde{q}^n_k = 1 - \sum_{l=1}^{K} \tilde{q}^n_{kl}$.

(2) *For* λ*-almost all* $i \in I$, $\{\alpha^n_i\}_{n=0}^{\infty}$ *is a Markov chain with transition matrix* z^n *at time* $n-1$ *defined by*

$$z^n_{kl} = q^n_l b^n_{kl} + \sum_{r,j=1}^{\infty} \nu^n_{rj}(l) b^n_{kr} q^n_{rj}.$$

(3) *For* λ*-almost all* $i \in I$ *and* λ*-almost all* $j \in I$*, the Markov chains* $\{\alpha^n_i\}_{n=0}^{\infty}$ *and* $\{\alpha^n_j\}_{n=0}^{\infty}$ *are independent.*

(4) *For* P*-almost all* $\omega \in \Omega$*, at each time period* $n \geq 1$*, the realized cross-sectional type distribution after random mutation* $\lambda(h^n_\omega)^{-1}$ *is* $\tilde{p}^n$ *and the realized cross-sectional type distribution at the end of the period* n*,* $p^n(\omega) = \lambda(\alpha^n_\omega)^{-1}$*, is equal to its expectation* $\bar{p}^n$*.*

Appendix B

Counting Processes

This appendix reviews intensity-based models of counting processes. Brémaud (1981) is a standard source.

All properties below are with respect to a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a given filtration $\{\mathcal{F}_t : t \geq 0\}$ satisfying the usual conditions unless otherwise indicated. We say that some $X : \Omega \times [0, \infty) \to \mathbb{R}$ is adapted if, for each time t, the function $X(\,\cdot\,, t) : \Omega \to \mathbb{R}$, also denoted X_t or $X(t)$, is $\mathcal{F}_t$-measurable. For market applications, $\mathcal{F}_t$ corresponds to the information held by a given set of agents at time t. To say that a process X is adapted can be interpreted as a statement that X_t is observable at time t, or could be chosen by agents at time t, on the basis of the information represented by $\mathcal{F}_t$.

A process Y is predictable if $Y : \Omega \times [0, \infty) \to \mathbb{R}$ is measurable with respect to the σ-algebra on $\Omega \times [0, \infty)$ generated by the set of all left-continuous adapted processes. The idea is that one can "foretell" Y_t based on all of the information available up to, but not including, time t. Of course, any left-continuous adapted process is predictable, as is, in particular, any continuous process.

A counting process N is defined via an increasing sequence $\{T_0, T_1, \ldots\}$ of random variables valued in $[0, \infty]$, with $T_0 = 0$ and with $T_n < T_{n+1}$ whenever

$T_n < \infty$, according to

$$N_t = n, \quad t \in [T_n, T_{n+1}), \tag{B.1}$$

where we define $N_t = +\infty$ if $t \geq \lim_n T_n$. We may treat T_n as the n-th jump time of N, and N_t as the number of jumps that have occurred up to and including time t. The counting process is nonexplosive if $\lim T_n = +\infty$ almost surely.

Definitions of "intensity" vary slightly from place to place. One may refer to Section II.3 of Brémaud (1981), in particular Theorems T8 and T9, to compare other definitions of intensity with the following. Let λ be a nonnegative predictable process such that, for all t, we have $\int_0^t \lambda_s \, ds < \infty$ almost surely. Then a nonexplosive adapted counting process N has λ as its intensity if $\{N_t - \int_0^t \lambda_s \, ds : t \geq 0\}$ is a local martingale.

From Brémaud's Theorem T12, without an important loss of generality for our purposes, we can require an intensity to be predictable, as above, and we can treat an intensity as essentially unique, in that: If λ and $\tilde{\lambda}$ are both intensities for N, as defined above, then

$$\int_0^\infty |\lambda_s - \tilde{\lambda}_s| \lambda_s \, ds = 0 \quad \text{a.s.} \tag{B.2}$$

We note that if λ is strictly positive, then (B.2) implies that $\lambda = \tilde{\lambda}$ almost everywhere.

We can get rid of the annoying "localness" of the above local-martingale characterization of intensity under the following technical condition, which can be verified from Theorems T8 and T9 of Brémaud (1981).

Proposition B.1 *Suppose that N is an adapted counting process and λ is a nonnegative predictable process such that, for all t, $E(\int_0^t \lambda_s \, ds) < \infty$. Then the following are equivalent:*

(i) N is nonexplosive and λ is the intensity of N.

(ii) $\{N_t - \int_0^t \lambda_s \, ds : t \geq 0\}$ *is a martingale.*

Proposition B.2 *Suppose that* N *is a nonexplosive adapted counting process with intensity* λ*, with* $\int_0^t \lambda_s \, ds < \infty$ *almost surely for all* t*. Let* M *be defined by* $M_t = N_t - \int_0^t \lambda_s \, ds$*. Then, for any predictable process* H *such that* $\int_0^t |H_s| \lambda_s \, ds$ *is finite almost surely for all* t*, a local martingale* Y *is well defined by*

$$Y_t = \int_0^t H_s \, dM_s = \int_0^t H_s \, dN_s - \int_0^t H_s \lambda_s \, ds.$$

If, moreover, $E\left[\int_0^t |H_s| \lambda_s \, ds\right] < \infty$ *for all* t*, then* Y *is a martingale.*

In order to define a Poisson process, we first recall that a random variable K with outcomes $\{0, 1, 2, \ldots\}$ has the Poisson distribution with parameter β if

$$\mathbb{P}(K = k) = e^{-\beta} \frac{\beta^k}{k!},$$

noting that $0! = 1$. A Poisson process is an adapted nonexplosive counting process N with deterministic intensity λ such that $\int_0^t \lambda_s \, ds$ is finite almost surely for all t, with the property that, for all t and $s > t$, conditional on $\mathcal{F}_t$, the random variable $N_s - N_t$ has the Poisson distribution with parameter $\int_t^s \lambda_u \, du$. (See Brémaud (1981), page 22.)

Appendix C

Essentials of Bargaining Theory

This appendix explains some basic concepts of bargaining games due to Nash (1950), Rubinstein (1982a), and Binmore, Rubinstein, and Wolinsky (1986).

Two players, 1 and 2, will bargain over some outcome in

$$X = \{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 + x_2 \leq 1\}.$$

Player i has a utility $u_i(x)$ for outcome x, where $u_i : X \to \mathbb{R}$ is continuous, concave, and strictly increasing in x_i and does not depend on x_j for $j \neq i$. Failure to agree, a "breakdown," results in a specified outcome b in X. We suppose that there exists some x in X with $u_i(x) > u_i(b)$ for both $i = 1$ and $i = 2$. That is, there is a feasible outcome that both players prefer to a breakdown. The objective is to model how an outcome in X is determined through bargaining between the two players. This setup is adaptable, via changes of variables and minor modifications, to bargaining over other outcomes than the choice depicted here, which amounts to "splitting a pie."

Rubinstein (1982a) suggested the following dynamic bargaining protocol, a particular alternating-offers extensive form game that is played in steps. Bargaining begins with Step 1. For each positive integer $k > 1$, there is a potential to reach bargaining Step k. If and when Step $k \geq 1$ is reached, the following sequence occurs.

1. An offer $x_k \in X$ is made. If k is odd, Player 1 makes the offer. If k is even, Player 2 makes the offer. Random offers are permitted.[1]

2. The counterparty to the offer accepts or rejects the offer. If the counterparty accepts, x_k is the outcome of the game and play ceases. The agreement decision is also permitted to be random.[2]

3. If the counterparty does not accept, the agents observe an independent Bernoulli trial Z_k with outcomes B (for "breakdown"), which occurs with probability $\eta > 0$, and A (for "advance"), which occurs with probability $1 - \eta$.

4. In the event of breakdown (that is, $Z_k = B$), bargaining ceases and the outcome of the game is b. In the event that $Z_k = A$, Step $k + 1$ is reached.

Without loss of generality, the offer x_k at Step k is determined by some measurable $f_k : X^{k-1} \times [0, 1] \to \mathbb{R}$, evaluated at the prior rejected offers $x_1, \ldots, x_{k-1}$ and (for randomization purposes) an independent random variable W_k that is uniformly distributed on $[0, 1]$. That is, $x_k = f_k(x_1, \ldots, x_{k-1}, W_k)$. The decision to accept or not at Step k is similarly determined by some measurable $g_k : X^k \times [0, 1] \to \{A, B\}$, evaluated at the prior and current offers $x_1, \ldots, x_k$ and (for randomization purposes) an independent random variable Y_k that is uniform $[0, 1]$.

A strategy σ_1 for Player 1 is a sequence $(f_1, g_2, f_3, g_4, \ldots)$ of such functions. A strategy σ_2 for Player 2 is likewise a sequence of such functions of the form $(g_1, f_2, g_3, f_4, \ldots)$.

Because there is a strictly positive probability η of a breakdown after each rejection of an offer, the game is completed in finitely many steps almost

[1] For this purpose, the offer can depend (measurably) on W_k, where W_k is revealed at the beginning of Step k, and where $W_1, W_2, \ldots$ is *iid* and uniformly distributed on $[0, 1]$.

[2] Again, the randomness can be based similarly on some *iid* sequence of uniformly distributed $Y_1, Y_2, \ldots$, independent of the sequence $W_1, W_2, \ldots$.

surely. For a given pair $\sigma = (\sigma_1, \sigma_2)$, of strategies, at the completion of the game the players obtain the random outcome $J(\sigma)$ determined by these strategies and the random draws $(W_1, Y_1, Z_1), (W_2, Y_2, Z_2), \ldots$, in the manner described above.

A Nash equilibrium is a pair $\sigma^* = (\sigma_1^*, \sigma_2^*)$ of strategies with the property that for each i, the strategy σ_i^* solves

$$\max_{\sigma_i} \; E[u_i(J(\sigma_i, \sigma_{-i}^*))],$$

where $(\sigma_i, \sigma_{-i}^*)$ refers to the strategy pair consisting of σ_i for Player i and σ_{-i}^* for the counterparty.

There can in general be multiple Nash equilibrium, some of which are not intuitively natural. In order to isolate a natural equilibrium, we restrict attention to a *perfect equilibrium*, which is a Nash equilibrium whose strategies are optimal for each player given the other player's strategy, not only at time 0 but also at each Step k, given the information that the players have received by that step. Specifically, a Nash equilibrium (σ_1^*, σ_2^*) is a perfect equilibrium if there does not exist some step number k, some player i, some strategy σ_i for Player i, and some event C_k of strictly positive probability that is measurable with respect to the information available[3] to player i at Step k with the property that

$$E[u_i(J(\sigma_i, \sigma_{-i}^*)) \,|\, C_k] > E[u_i(J(\sigma_i^*, \sigma_{-i}^*)) \,|\, C_k].$$

Given the strict monotonicity of $u_i(x)$ with respect to x_i, an outcome x is Pareto efficient if $x_1 + x_2 = 1$.

Proposition C.1 [Rubinstein (1982)] *There exist unique Pareto efficient*

[3]That is, if Player i is making an offer at Step k, then C_k is in the σ-algebra generated by $\{(f_1, g_1, Z_1), \ldots, (f_{k-1}, g_{k-1}, Z_{k-1})\}$, and if Player i is receiving an offer at Step k, then C_k is in the σ-algebra generated by $\{(f_1, g_1, Z_1), \ldots, (f_{k-1}, g_{k-1}, Z_{k-1}), f_k\}$. If the event C_k includes a prior breakdown, the strict inequality shown could never be achieved given that the associated outcome is b for any strategy.

outcomes x^* and y^* with the property that

$$u_1(y^*) = (1-\eta)u_1(x^*) + \eta u_1(b)$$

and

$$u_2(x^*) = (1-\eta)u_2(y^*) + \eta u_2(b).$$

There exists a unique perfect equilibrium. For this equilibrium:

- *At any odd step, Player 1 offers x^* and Player 2 rejects any offer strictly inferior to x^*.*
- *At any even step, Player 2 offers y^* and Player 1 rejects any offer strictly inferior to y^*.*

Thus, the game is completed at Step 1 when Player 1 offers x^* and Player 2 accepts this offer.

This perfect equilibrium of the bargaining game has been shown by Binmore, Rubinstein, and Wolinsky (1986) to be closely related to the axiomatic solution of Nash (1950), who proposed axioms under which the outcome x_N is the solution of the problem

$$\max_{x \in X} \quad [u_1(x) - u_1(b)][u_2(x) - u_2(b)].$$

Proposition C.2 [Binmore, Rubinstein, and Wolinsky (1986)] *As the probability η of a single-step breakdown converges to zero, the unique perfect-equilibrium outcomes x^*_η and y^*_η proposed by Players 1 and 2, respectively, both converge to the Nash solution x_N.*

For large η, Player 1 has a bargaining advantage, in that if Player 2 rejects the offer of Player 1, the lower-utility breakdown outcome is a serious threat. The fact that both x^*_η and y^*_η converge to the Nash solution x_N as the breakdown probability becomes small implies that the relative advantage of being Player 1 (the first to offer) is increasingly unimportant.

For completeness, we re-state the axioms of Nash (1950), as follows. A *bargaining problem* consists of a pair (U, d), where

$$U = \{(u_1(x), u_2(x)) : x \in X\}$$

is the set of feasible utilities and $d = (d_1, d_2) = (u_1(b), u_2(b))$ is the pair of breakdown utilities. Without loss of generality, we suppose that $b = 0$. We can relax our earlier utility assumptions a bit, and require only that U is compact and convex. We maintain the assumption that there exists v in U with $v_i > d_i$.

The set of bargaining problems is denoted $\mathcal{B}$. A *bargaining solution* is a function $F : \mathcal{B} \to U$. We denote $F_i(U, d) = [F(U, d)]_i$. Consider the following axioms for a bargaining solution F.

1. **Pareto efficiency.** For any (U, d), there is no v in U with $v_1 > F_1(U, d)$ and $v_2 > F_2(U, d)$.

2. **Symmetry.** Suppose that (U, d) has $d_1 = d_2$ and satisfies $(v_2, v_1) \in U$ whenever $(v_1, v_2) \in U$. Then $F_1(U, d) = F_2(U, d)$.

3. **Invariance to equivalent utility representations.** For a given (U, d), let (U', d') be given, for some scalars $\alpha_i > 0$ and β_i, by the transformations $d'_i = \alpha_i d_i + \beta_i$ and

$$U' = \{(\alpha_1 u_1 + \beta_1, \alpha_2 u_2 + \beta_2) : u \in U\}.$$

   Then $F_i(U', d') = \alpha_i F(U, d) + \beta_i$.

4. **Independence of irrelevant alternatives.** Given a bargaining game (U, d) and some $U' \subset U$, suppose that $F(U, d) \in U'$. Then $F(U', d) = F(U, d)$.

Proposition C.3 [Nash (1950)] *There exists a unique bargaining solution F satisfying Axioms 1 through 4. For any bargaining problem (U, d), the outcome $F(U, d)$ of this bargaining solution is the Nash solution, that solving*

$$\max_{v \in U} \quad [v_1 - d_1][v_2 - d_2].$$

Various alternatives or re-characterizations of the Nash axioms have been proposed, for example by Lensberg (1988). Network-based characterizations of the connection between alternating-offers extensive-form bargaining games and axiomatic solutions of bargaining problems have been developed by Navarro and Perea (2013), Stole and Zwiebel (1996), Duffie and Wang (2014), and de Fontenay and Gans (2013).

Bibliography

Binmore, K., A. Rubinstein, and A. Wolinsky, 1986, "The Nash Bargaining Solution in Economic Modelling," *The Rand Journal of Economics*, 17(2), pp. 176–188.

Binmore, K., and L. Samuelson, 1999, "Evolutionary Drift and Equilibrium Selection," *Review of Economic Studies*, 66, 363–393.

Brémaud, P., 1981, *Point Processes and Queues: Martingale Dynamics*. New York: Springer-Verlag.

Burdzy, K., D. M. Frankel, and A. Pauzner, 2001, "Fast Equilibrium Selection by Rational Players Living in a Changing World," *Econometrica*, 69, 163–189.

Cavalli-Sforza, L. L., and W. F. Bodmer, 1971, *The Genetics of Human Population*. Freeman, San Francisco.

de Fontenay, C., and J. S. Gans, 2013, "Bilateral Bargaining with Externalities," working paper, Melbourne Business School Discussion Paper No. 2004-32.

Dekel, E., and S. Scotchmer, 1999, "On the Evolution of Attitudes towards Risk in Winner-Take-All Games," *Journal of Economic Theory*, 87, 125–143.

Diamond, P., 1971, "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 156–168.

——— , 1982, "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90, 881–894.

Diamond, P., and J. Yellin, 1990, "Inventories and Money Holdings in a Search Economy," *Econometrica*, 58, 929–950.

Duffie, D., N. Gârleanu, and L. H. Pedersen, 2005, "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847.

——— , 2007, "Valuation in Over-the-Counter Markets," *Review of Financial Studies*, 20, 1865–1900.

Duffie, D., S. Malamud, and G. Manso, 2009, "Information Percolation with Equilibrium Search Dynamics," *Econometrica*, 77, 1513–1574.

——— , 2010, "The Relative Contributions of Private Information Sharing and Public Information Releases to Information Aggregation," *Journal of Economic Theory*, 145, 1574–1601.

——— , 2014, "Information Percolation in Segmented Markets," *Journal of Economic Theory*, xx, yy.

Duffie, D., L. Qiao, and Y. Sun, 2014a, "Continuous-Time Random Matching (In preparation)," Working Paper, Graduate School of Business, Stanford University.

——— , 2014b, "Dynamic Directed Random Matching," Working Paper, Graduate School of Business, Stanford University.

Duffie, D., and Y. Sun, 2007, "Existence of Independent Random Matching," *Annals of Applied Probability*, 17, 386–419.

——— , 2012, "The Exact Law of Large Numbers for Independent Random Matching," *Journal of Economic Theory*, 147, 1105–1139.

Duffie, D., and C. Wang, 2014, "Efficient Contracting in Network Markets," Working Paper, Graduate School of Business, Stanford University.

Ferland, R., and G. Giroux, 2008, "Law of Large Numbers for Dynamic Bargaining Markets," *Journal of Applied Probability*, 45, 45–54.

Fudenberg, D., and D. Levine, 1993, "Steady-State Learning and Nash Equilibrium," *Econometrica*, 61, 547–573.

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

——— , 1986b, "Bargaining and Competition, Part II: Existence," *Econometrica*, 54, 807–818.

Gârleanu, N., 2009, "Pricing and Portfolio Choice in Imperfect Markets," *Journal of Economic Theory*, 144, 532–564.

Green, E. J., and R. Zhou, 2002, "Dynamic Monetary Equilibrium in a Random Matching Economy," *Econometrica*, 70, 929–970.

Hardy, G. H., 1908, "Mendelian Proportions in a Mixed Population," *Science*, 28, 49–50.

Harrington, J. E., 1998, "The Social Selection of Flexible and Rigid Agents," *American Economic Review*, 88, 63–82.

Hellwig, M., 1976, "A Model of Monetary Exchange," Econometric Research Program, Research Memorandum Number 202, Princeton University.

Hosios, A., 1990, "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, 57, 279–298.

Kiyotaki, N., and R. Wright, 1993, "A Search-Theoretic Approach to Monetary Economics," *American Economic Review*, 83, 63–77.

Krainer, J., and S. LeRoy, 2002, "Equilibrium Valuation of Illiquid Assets," *Economic Theory*, 19, 223–242.

Lagos, R., and G. Rocheteau, 2009, "Liquidity in Asset Markets with Search Frictions," *Econometrica*, 77, 403–426.

Lensberg, T., 1988, "Stability and the Nash solution," *Journal of Economic Theory*, 45(2), 330 – 341.

McLennan, A., and H. Sonnenschein, 1991, "Sequential Bargaining as a Noncooperative Foundation for Walrasian Equilibrium," *Econometrica*, 59, 1395–1424.

Mortensen, D., 1982, "Property Rights and Efficiency in Mating, Racing, and Related Games," *American Economic Review*, 72, 968–979.

Mortensen, D., and C. Pissarides, 1994, "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, 61, 397–415.

Nash, J., 1950, "The Bargaining Problem," *Econometrica*, 18, 155–162.

Navarro, N., and A. Perea, 2013, "A Simple Bargaining Procedure for the Myerson Value," *BEJTE*, 13(1), 131–150.

Rubinstein, A., 1982a, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, pp. 97–109.

——— , 1982b, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

Rubinstein, A., and A. Wolinsky, 1985, "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, 53, 1133–1150.

Stole, L., and J. Zwiebel, 1996, "Intra-Firm Bargaining under Non-Binding Contracts," *Review of Economic Studies*, 63(3), 375–410.

Sun, Y., 2006a, "The exact law of large numbers via Fubini extension and characterization of insurable risks," *Journal of Economic Theory*, 126, 31–69.

Sun, Y. N., 2006b, "The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks," *Journal of Economic Theory*, 126, 31–69.

Trejos, A., and R. Wright, 1995, "Search, Bargaining, Money, and Prices," *Journal of Political Economy*, 103, 118–140.

Wolinsky, A., 1990, "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58, 1–23.