

Strategic Interactions in Large Populations

by

Glenn David Ellison

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1992

© Massachusetts Institute of Technology 1992. All rights reserved.

Author
Department of Economics
May 11, 1992

Certified by
Drew Fudenberg
Professor of Economics
Thesis Supervisor

Accepted by
Richard Eckaus
Chairman, Departmental Committee on Graduate Students

ARCHIVES
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 10 1992

LIBRARIES

Strategic Interactions in Large Populations

by

Glenn David Ellison

Submitted to the Department of Economics
on May 11, 1992, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis consists of four essays in game theory. The first two chapters are concerned with games involving large numbers of players. The first chapter discusses the dynamic implications of learning in a large population coordination game, focusing on the structure of the matching process which describes how players meet. The combination of experimentation and the myopic attempts of players to coordinate with those around them creates evolutionary forces which lead players to coordinate on the risk-dominant equilibrium. To understand play in reasonable finite horizons it is necessary to analyze not only the limits of the dynamic systems but also the rates at which they converge. Populations in which players interact with small stable sets of neighbors are far more amenable to rapid change and hence more likely to reflect evolutionary forces than are populations with more uniform matching.

The second chapter considers the repeated prisoner's dilemma in a large population random matching setting where players are unable to recognize their opponents. Despite the informational restrictions cooperation is still a perfect equilibrium supported by "contagious" punishments. The equilibrium does not require excessive patience, and contrary to previous thought need not be extraordinarily fragile. It is robust to the introduction of small amounts of noise and remains nearly efficient. Extensions are discussed to models with heterogeneous rates of time preference and without public randomizations.

The final two chapters discuss empirical tests of theoretical models of collusion under uncertainty. The third chapter discusses the implications of the Green and Porter (1984) and Rotemberg and Saloner (1986) theories for the pattern of price wars in Joint Executive Committee, an 1880's railroad cartel. A switching regressions model with Markov transitions is used to analyze the causes of the price wars. The results provide some support for the predictions of the first theory.

The fourth chapter further explores the behavior of the Joint Executive Committee. The interpretation of switching regressions models is discussed, and it is argued that to some extent the price wars may reflect cheating by the firms rather than an optimal cartel design.

Thesis Supervisor: Drew Fudenberg

Title: Professor of Economics

To Sara

Acknowledgments

I must begin an acknowledgement of all the help I've received in completing this thesis by noting the contributions of my wife Sara. She convinced me to enroll at MIT three years ago, and has stood by me ever since. In addition, I owe her thanks specifically for her help on each of the essays which comprise this thesis.

As for the essays, they have benefitted greatly from the opportunity I've had to interact with a number of people while at MIT. I'd like to thank Chris Harris, Paul Joskow, Michi Kandori, Wally Mullin, Danny Quah, Julio Rotemberg, Lones Smith, Jean Tirole, Jeff Wooldridge, and Jeff Zwiebel, along with all the participants at the theory lunch and the theory and industrial organization seminars. Special thanks are due to the second reader of this thesis, Frank Fisher, and to Rob Porter, who inspired and provided the data for the final two chapters.

I am especially grateful to Drew Fudenberg, who has provided far more help than could be expected from a thesis supervisor. In countless hours of discussions, he has influenced everything from the broadest directions of my research to the details of its exposition. I also owe him thanks simply for his support and encouragement.

I thank the Sloan Foundation for financial support.

I would also like to thank everyone else who has helped me enjoy my years at MIT. I'd be afraid to list all my friends for fear of forgetting any of them, so I'll just mention them as a group, with special thanks to everyone who played basketball on Friday mornings.

Finally, I want to thank my parents for everything they've done for me over the years. I hope they enjoy having this thesis sitting on their bookshelf for years to come.

Contents

Abstract	2
Acknowledgements	5
Introduction	8
1 Learning, Local Interaction, and Coordination	12
1.1 Introduction	12
1.2 The Model	17
1.2.1 Coordination Games with Bounded Rationality	17
1.2.2 Local and Uniform Matching Rules	19
1.3 Model Dynamics	22
1.4 Limits and Rates of Convergence	26
1.5 Simulation Results	38
1.5.1 Matching Rules	38
1.5.2 Population Sizes	44
1.5.3 Heterogeneity	46
1.6 Conclusion	50
References	51
Appendix	52
2 Cooperation in the Prisoner's Dilemma with Anonymous Random Matching	59
2.1 Introduction	59

2.2	The Random Matching Model	63
2.3	Stability and Efficiency with Noise	72
2.4	Cooperation without Public Randomizations	78
2.5	Conclusion	85
	References	87
	Appendix	89
3	Markov Models of Price Wars in the Joint Executive Committee	96
3.1	Introduction	96
3.2	History	97
3.3	Theory	98
3.4	Econometric Model	102
3.5	Previous Results	106
3.6	Data	109
3.7	Results	111
	3.7.1 The “standard” model	111
	3.7.2 The Green-Porter and Rotemberg-Saloner models	114
	3.7.3 Sensitivity of the results to functional form assumptions	118
3.8	Conclusion	121
	References	124
4	What Does Not Seeing Something Look Like? Secret Price Cuts in the Joint Executive Committee	126
4.1	Introduction	126
4.2	Models of Unobserved Regimes	127
4.3	Secret Price Cuts in the JEC	131
	4.3.1 Looking for price cuts via hidden regimes	132
	4.3.2 Do the estimated regimes look like secret price cuts?	138
4.4	Conclusion	144
	References	145

Introduction

Game theory has been widely applied in economics to describe the behavior of interacting agents, for example, of competing firms. The game-theoretic approach is distinguished by the assumption that the optimal behavior of each agent is influenced by the behavior of the others. The essays which constitute this thesis focus on two particular issues within a broad field. The first two chapters address problems which arise in the attempt to apply game-theoretic models to large populations. The second two chapters discuss empirical tests of game-theoretic models with uncertainty.

The first two chapters of this thesis discuss models of strategic play in large populations. The idea of discussing play in large populations is not to simply restate theorems for N players instead of two, but rather to model formally a host of concerns which take on increased importance in large populations. For example, the first essay uses a network to model the structure of friendships and business contacts which govern association and asks how the network structure affects the behavior we observe. The second essay explores informational restrictions motivated by the fact that members of large societies will not always know who their trading partners are, how they have behaved in the past, etc. In each case I allow for heterogeneity among the players and for the uncertainty which inevitably arises from mistaken and misinterpreted actions.

The first essay addresses a very basic problem in game theory. How should we expect people to play in models which have multiple equilibria? The question is addressed in the simplest framework available, coordination games where the players have common interests. These games have several equilibria corresponding to coordination on each of the available actions, so standard Nash equilibrium analysis gives us

no guidance as to which equilibrium we should expect to see. Nonetheless, it is often argued that some outcomes are more likely to be observed than others. For example, we might expect that players should coordinate on the choice they most prefer.

The first essay examines the choice of equilibrium in a dynamic model which incorporates a notion of bounded rationality. As the players attempt to learn how their opponents will play the popularity of the various actions changes. The rather surprising conclusion of previous work is that for a variety of specifications, the learning process leads players to play one particular equilibrium. The principal result of the first chapter is that the pattern of play is also greatly influenced by the structure of the interactions in the population. When players care mostly about the actions of a few close friends or colleagues, “evolutionary forces” are powerful and lead players to coordinate on what is known as the “risk-dominant” equilibrium. On the other hand, when interactions are largely impersonal no equilibrium selection takes place within a reasonable time period and we should expect instead to see an historically determined outcome repeated over and over again.

The second essay explores another very basic problem in game theory, namely whether it is possible for players who interact repeatedly to achieve cooperative outcomes. When two players meet repeatedly to play the Prisoner’s dilemma, it is a standard result that cooperative equilibria exist. Despite a short-run incentive to cheat each other, the players cooperate in order to avoid being punished in the future. To apply similar reputation-based punishments in a large society, however, requires that somehow cheaters must be identified and honest players must be kept informed of everyone’s reputation.

The second chapter explores whether large populations may achieve cooperation despite the inherent difficulties in observing and communicating reputations. Specifically, the possibility of cooperation is discussed given an informational extreme – an anonymous random matching model where it is impossible for players to either recognize cheaters or to communicate with others. There are two main motivations for this study. First, the theorems are extensions of familiar Folk theorem results. In this context, it is interesting to note how “contagious” punishments can be used to over-

come the informational limitations which have been imposed. The impact of trembles and public randomizations may also be compared with previous studies. Second, the results can be viewed as a comment on recent work in the study of institutions. If we wish to view the development of institutions as a response to the problems of larger societies, we must understand the types of behavior which are possible without any institutions.

The final two chapters present an approach to the empirical testing of game-theoretic hypotheses. In particular, both essays discuss the application of theories of collusion under uncertainty to the experience of the Joint Executive Committee, a railroad cartel which controlled grain shipments between Chicago and the Eastern seaboard in the 1880's. When firms' actions are observable, a cartel can collude effectively by punishing any firm which deviates from the agreement. When demand is uncertain and firms may offer secret price cuts, however, the problem of detecting cheating is much more difficult. Nonetheless, some degree of collusion is possible and theory predicts that interesting patterns of behavior will emerge. The first such result, due to Green and Porter, held that even in the optimal equilibrium we should expect to see price wars. More controversially, Rotemberg and Saloner have discussed the consequences of cyclical demand and argued that price wars are more likely to occur in good times than in bad times.

The third chapter of this thesis begins with an informal discussion of the ways in which these theories must be adapted to conform to the historical situation of the Joint Executive Committee. In the case of the Green and Porter theory, it is argued that many nearly efficient cartel mechanisms are possible and empirical analysis should focus on detecting whether any of these are present. It is also argued that we should not necessarily expect to see price wars occurring during good times although there may be cyclical patterns in price wars within each year. To test these predictions, a Markov model is proposed. The idea is that we wish to estimate two things: the structure of supply and demand in the industry, and the strategies of the firms. Each state of the Markov process represents a week of operation of the cartel during which the firms' actions are fixed, and from which we estimate supply and demand. A model

of probabilistic transitions between states is used to infer the firms' strategies, *i.e.* to determine why firms decide to begin a price war.

The data clearly indicate that price wars were present in the Joint Executive Committee. Further, the results of the third chapter provide some support for the hypothesis that a Green-Porter style cartel design was used. However, clear evidence is not found for a mechanism strong enough to support collusion. Certainly this could reflect simply the limited data which is available. No evidence is found for the commonly discussed price wars during booms effect.

The fourth chapter tries to resolve some of this uncertainty in exploring further whether the Green-Porter theory adequately describes the structure of the Joint Executive Committee. In particular, the results of the third chapter might also be interpreted as indicating that the firms erred in their cartel design so that it indeed was in each firm's private interest to cheat. The chapter therefore investigates the extent to which firms may have deviated from their agreements. It begins with a discussion of the interpretation of models with hidden regimes. The estimation follows a two step approach where several models are first estimated to identify whether significant unobserved variables are omitted from the models of the previous chapter. Most of the discussion is then devoted to the attempt to determine whether secret price cuts might account for the results. While the analysis is largely speculative, it is argued that cheating may have been common.

Chapter 1

Learning, Local Interaction, and Coordination

1.1 Introduction

Even the simplest game theoretic models all too often have multiple equilibria. A typical example is the coordination game which arises when two players must work together in order to achieve a commonly desired outcome, but in which neither player will benefit from his efforts if his partner does not do his part. In this case, we regard the players working together as the “good” equilibrium and speak of coordination failure if it does not occur. In other coordination games, like that of two cars approaching each other on a highway, we may not care which equilibrium (both keeping to the right or both keeping to the left) occurs, but it remains very important that the players somehow have common beliefs so that an equilibrium is played. In trying to understand play in these games we are led to ask why we should expect players to coordinate on an equilibrium and whether there is any reason to believe that one equilibrium is more likely than the other.

Recently, models of learning have been used to explore these and other fundamental questions of game theory. These models typically investigate whether we can predict behavior in a game by examining the disequilibrium process by which players learn their opponents’ play and adjust their strategies over time. In their analysis of

coordination games in large populations, Kandori, Mailath, and Rob (KMR) (1991) derive surprisingly strong predictions from such an approach. They show that the simple combination of random experimentation and the myopic attempts of players to coordinate with those around them creates powerful dynamic forces which influence the evolution of play over time. In analyzing the long-run limit of this dynamic process, they are able to show not only that players will achieve coordination on an equilibrium, but that one particular equilibrium, the “risk-dominant equilibrium” will be selected.

In this paper, I adopt a similar framework to examine the play of coordination games in large populations. In each period of a dynamic model the players are randomly matched and each pair plays a 2×2 coordination game. Players do not know who their opponents will be until after they have chosen their actions. The behavioral assumptions incorporating noise and myopic responses by boundedly rational players are also borrowed from KMR.

The model is intended not only as an abstract demonstration of how coordination might arise, but also as a practical model which may help us understand why some populations are more susceptible to coordination failure. For this reason, I discuss not only the limit of each dynamic process, but also whether the limit reflects forces which would be felt within an economically reasonable time frame. If, as will be the case for one model, a dynamic system takes 10^{100} periods to approach its limit, we cannot say that the limit is a good prediction for what we will see when the game is repeated a few hundred times. While it is very hard to draw a dividing line and say exactly how fast a system must converge for the limit to be relevant, the models of this paper exhibit such extreme contrasts that meaningful conclusions are possible. When a system adjusts very slowly, I shall conclude that whatever historical factors determine the initial play will continue to determine play long into the future so that, for example, even Pareto-superior alternatives will not be adopted. On the other hand, when a system approaches its limit quickly, I shall conclude that dynamic forces do lead us to expect to see the limiting behavior.

The model itself departs from that of KMR in that it allows for different matching

processes within the population.¹ The conclusion of this paper is that the nature of the matching process is crucial to a determination of whether historical factors or risk dominance will determine play. I focus on two extremes among the possible matching rules, which I shall describe as uniform and local.

The uniform matching rule is that used in KMR. Players are equally likely to be matched with any member of the large population and therefore place only small weight on coordinating with any particular individual. I argue that in such populations any evolution in play is unlikely to occur within a reasonable period of time, and hence, that historical factors will determine play. In contrast, I shall describe as local a matching rule in which players interact with a small group of close friends, neighbors, or colleagues. In such situations, I argue that rapid changes in play are indeed possible and we may expect to see coordination on the risk-dominant equilibrium.

To provide a clearer picture of the types of questions that can be addressed within this framework, I now begin with a lengthy discussion of several examples. Consider first the choice of location of a trade fair. In medieval England, a large portion of trade took place at annual trade fairs, of which the largest was the Sturbridge Fair.² The Sturbridge Fair was chartered in 1211, and by the 14th century it had grown to a sufficient size that people traveled hundreds of miles to attend and buy a year's provisions. When Daniel Defoe visited Sturbridge in 1723, he described a half mile square fairground with such a tremendous variety of commerce that he was convinced (probably without any point of comparison) the Fair was "... not only the greatest in the whole nation, but in the world;"³ On one street, for example, he notes that

Scarce any trades are omitted, goldsmiths, toyshops, braziers, turners, milliners, haberdashers, hatters, mercers, drapers, pewterers, china-warehouses, and in a word all the trades that can be named in London; with coffee-houses, taverns, brandy-shops, and eating-houses, innumerable, ...⁴

¹In independent work, Blume (1991) discusses a related continuous time model in which players are spatially distributed and interact with a finite set of neighbors.

²The description below is based on that of Walford (1883).

³Defoe (1986), p. 102

⁴*Ibid.* p. 102

The fair contained not merely retail businesses but also extensive wholesale trade with, for example, woolen manufactures in tents, “as vast ware-houses piled up with goods to the top.”⁵

The question I address here is why this trade was centered in Sturbridge of all places for about eight centuries. Sturbridge was not the most efficient location. To the contrary, the Sturbridge Fair (located just outside Cambridge) had two clear drawbacks. First, the town was too small to support such an event. Second, and perhaps more importantly, the interior location necessitated the rather cumbersome transport of goods down the Ouse and Cam rivers from the port of King’s Lynn. Observing the volume of hops traded, this glaring inefficiency puzzled Defoe, and led him to ask

... why this fair should be thus, of all other places in England, the centre of that trade; and so great a quantity of so bulky a commodity be carried thither so far:⁶

In this paper, I argue that we can understand the location of the Sturbridge Fair by modeling the location choice as the outcome of a coordination game played by the attending merchants. Each year, the merchants independently choose actions in deciding to travel to one of several possible locations. Payoffs are determined by a random matching process insofar as we can view each pair as having a mutually advantageous trade with some probability. The matching process may be approximated as uniform to the extent that the merchants do not know each other or have no way of knowing in advance who will have the goods they need. The payoffs are those of a coordination game because each pair of merchants meets and has the opportunity to realize gains from trade only if both have chosen to attend the same location. The outcome of this coordination game appears to have been determined by historical circumstances relevant only at the start of the game.

The basic framework of this paper is applicable to a variety of modern settings as well. I will briefly sketch two such examples in which players tend to interact with

⁵ *Ibid.* p. 103

⁶ *Ibid.* p.104

friends or colleagues and hence in which local matching rules may be appropriate. First, consider the attendance of a periodic convention or reunion. Here, the players are the potential attendees who must decide whether to go. If players attend, their payoffs reflect the useful information they exchange, the contacts they make, or simply pleasurable conversations they have with other players. The game is a coordination game with two equilibria if the players would like to attend if everyone else were to do so, but would prefer to stay home if the convention will be poorly attended. In the case of conventions, a local matching rule may describe players' expectations over who is likely to have the information they desire.

For a more economically significant example, consider the adoption of new technologies when network externalities are present. While the new technology may be superior, network externalities will initially favor the old established standard. Whether the new technology is adopted depends on the outcome of a game played by the firms or individuals who must choose a technological standard for each new project they undertake. For example, we might imagine a group of electronics manufacturers deciding which of several possible standards a new product will adhere to, or applied economists deciding which of several possible software packages to use for a new project. In the latter example, network externalities arise when an economist knows that at some point during the project he is likely to want to borrow programs from a colleague, which will only be easy to do only if the programs are written in the same language that he is using. When network externalities are strong enough to outweigh individual preferences, the payoffs are those of a coordination game. Whether a new superior technology is likely to become predominant or whether the network externalities will allow the incumbent technology to endure is an equilibrium selection problem in this coordination game.

The paper is structured as follows. The model is described in Section 2. Section 3 contains some simple examples of the dynamics of the learning model. Section 4 contains the main theoretical results on both limiting distributions of play and rates of convergence. Section 5 discusses the results of numerical simulations which illustrate both the speed of convergence for reasonable parameter values and the importance

of the various assumptions.

1.2 The Model

1.2.1 Coordination Games with Bounded Rationality

The model described here can be thought of as having two classes of assumptions: those concerning the nature of the game being played and those describing the particular behavioral rules which players follow.

The basic model is of a repeated game played in periods $t = 1, 2, 3, \dots$. There is a large population of N players (perhaps a few hundred for typical applications). In each period, player i chooses one of two possible actions $a_{it} \in \{A, B\}$. The payoff to player i is given by

$$u_i(a_{it}, a_{-it}) = \sum_{j \neq i} p_{ij} g(a_{it}, a_{jt})$$

where the payoffs g are those of the 2×2 coordination game pictured below. Formally, it is required that $a > d$ and $b > c$ so that (A, A) and (B, B) are both Nash equilibria. In addition, I assume that (A, A) is the “risk-dominant” equilibrium as defined in Harsanyi and Selten (1988). In a symmetric 2×2 game, (A, A) is risk dominant if and only if $(a - d) > (b - c)$. Note that when the strategies have equal security levels ($c = d$), (A, A) is also the Pareto optimum.

	<i>A</i>	<i>B</i>
<i>A</i>	a, a	c, d
<i>B</i>	d, c	b, b

In many applications, we can envision the players to be playing a random matching

game in which case the weights p_{ij} will represent the probability that players i and j are matched in period t , and $g(a_{it}, a_{jt})$ gives the payoff to player i when he is matched with player j .

Rather than assuming complete rationality, I instead simply specify behavioral rules which the players follow. The rules are intended to capture the intuitive notion of reactive players. The players will myopically maximize their short-run payoffs in each period, and in addition will be unsophisticated in not realizing how play will change over time. In particular, I assume that in period t player i chooses

$$a_{it} \in \arg \max_{a_i} u_i(a_i, a_{-it-1})$$

with probability $1 - 2\epsilon$, and with probability 2ϵ he chooses an action at random with 50-50 probability.

Several remarks are called for. First, note that the choice of strategy requires that player i observe the period $t - 1$ play of *all* other players with whom he may be matched in the future, not just the players with whom he was matched in period $t - 1$. The reasonableness of this assumption varies with the application, *e.g.*, in the case of an annual reunion it requires that all invitees be sent a list of those attending after the reunion is over. Second, the 2ϵ probability randomizations are meant to represent the cumulative effect of noise introduced into the system through deliberate experimentation, trembles in strategy choices, and the play of new players unfamiliar with the history of the game. Which of these is most reasonable again depends on the particular application intended. Finally, I should emphasize that this is best thought of as a specification of bounded rationality, *not* as a form of rationality with impatient players. In period t , each player is reacting to the play in period $t - 1$, not to a prediction of how his opponents will play in period t . There are several possible justifications for this assumption. Perhaps it is not possible for the players to extrapolate and predict the future play of their opponents because they either do not know the decision rules followed by their opponents or can not observe the play of their opponents' opponents. Alternatively, players might simply be stupid.

If we do wish to regard the players as rationally responding to incorrect beliefs, we must take the period $t - 1$ actions to be the period t beliefs. This approach leads to two problems. First, particularly in models with local matching, the beliefs may be excessively naive as there may be cycles in play that we would expect “rational” players to recognize. Second, once we introduce rational players we must justify the assumption that players choose myopic best responses. It is often argued that in large populations strategic play is unlikely. What is really important, however, is not the size of the population but rather how rapidly a player’s actions affect the distribution of play. In models which allow only very slow evolution, myopic responses are probably reasonable. When evolution is more rapid though, strategic play is possible even in large populations, and to ignore it requires the additional assumption that players are impatient. The presence of a few patient strategic players, however, would likely strengthen the conclusions of this paper as strategic players tend to spur evolution when evolution will be fast anyway and do not try to affect future play when there is great inertia.

1.2.2 Local and Uniform Matching Rules

Within the basic framework described above, I contrast two extreme specifications of the matching process. I term the two types of matching rules uniform and local. The *uniform* matching rule is given by

$$p_{ij} = \frac{1}{N - 1} \quad \forall j \neq i.$$

With this rule, player i will choose his period t strategy considering only the fraction of the population playing each strategy at time $t - 1$, not the identities of the players using each strategy.

Such a rule seems appropriate in the example of the medieval trade fair mentioned above. We can model the location decision as a process in which each trader decides each year whether to travel to Sturbridge at a cost of c or to travel to a more convenient port instead at a cost of 0 . When any pair of traders meet, assume that with

probability α they are able to complete a trade which increases each player's utility by one unit. If we assume that a player meets all other traders who attend the same fair we can write the payoffs as

$$u_i(a_i, a_{-i}) = (N - 1) \sum_{j \neq i} \frac{1}{N - 1} g(a_i, a_j),$$

where g is given by the matrix below. The assumption of uniform matching expresses the idea that each trader has no information whatsoever about with whom he is likely to be able to complete a trade until after he has arrived at the fair he has chosen to attend and can observe the goods offered for sale.

	Port	Sturbridge
Port	α, α	$0, -\frac{c}{N-1}$
Sturbridge	$-\frac{c}{N-1}, 0$	$\alpha - \frac{c}{N-1}, \alpha - \frac{c}{N-1}$

In contrast, I shall use the term *local* matching as an informal description of several matching rules in which each player is likely to be matched only with a small fixed subset of the population. For simplicity, I will usually envision the players as being spatially distributed around a circle. In the most extreme local matching rule, each player is only ever matched with one of his two immediate neighbors, *i.e.*,

$$p_{ij} = \begin{cases} \frac{1}{2} & \text{if } i - j \equiv \pm 1 \pmod{N} \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, for any $k \geq 1$ we can define a rule where each player has $2k$ neighbors by

$$p_{ij} = \begin{cases} \frac{1}{2k} & \text{if } i - j \equiv \pm 1, \pm 2, \dots, \pm k \pmod{N} \\ 0 & \text{otherwise.} \end{cases}$$

For a less extreme example, we might assign positive probability to any pair of players being matched, but construct the probabilities so that most of the time players are matched with those nearby, *e.g.* for N even

$$p_{ij} = \begin{cases} \frac{1}{2^{k+1}} & \text{for } k = \min\{|i - j|, N - |i - j|\} \neq \frac{N}{2} \\ \frac{1}{2^{N/2}} & \text{for } |i - j| = \frac{N}{2}. \end{cases} \quad (1.1)$$

Local matching rules are appropriate to describe situations where players interact not with the population as a whole, but rather with a few close friends or colleagues. Consider, for example, the group of economists choosing software packages. While there are thousands of economists in the country, each will be able to identify, upon beginning a project, a small group of colleagues with whom he or she is likely to want to share programs or data. Hence, he or she will only consider the software choices of that small group when making a decision.

Before moving on, I would like to emphasize three essential features of the local matching rules described above. First, each player assigns a large weight to a small subset of the population. Second, the locations of the players are fixed over time so that each player's likely opponents remain constant over time. Third, there is considerable overlap in the groups of neighbors so that a player's neighbors' neighbors are likely to be his neighbors as well. It is the combination of all these features which allow for the existence of small clusters within the population each member of which is matched with another member with probability at least $\frac{1}{2}$. It is the possibility of a new strategy gaining a foothold within one of these clusters which allows the relatively rapid transition to the risk-dominant equilibrium in these populations.

1.3 Model Dynamics

For the remainder of this paper, I discuss the dynamic pattern of play in the model described above. The approach of the paper is as follows. I assume that at some point in the past, arbitrary historical factors determined the initial strategies of the players. The behavioral rules then generate a dynamic system which describes the evolution of players' strategy choices over time. I formally discuss both the limit of this system and the rates at which the limit is approached. First though, I describe the dynamic evolution of play in a few simple cases in order to motivate subsequent results.

The dynamics of the model with uniform matching are virtually identical to those described in KMR. Let q_i be the fraction of player i 's opponents who played A in period $t - 1$. Note that

$$\begin{aligned} u_i(A, a_{-it-1}) \geq u_i(B, a_{-it-1}) &\iff q_i a + (1 - q_i)c \geq q_i d + (1 - q_i)b \\ &\iff q_i \geq \frac{b - c}{(a - d) + (b - c)} \equiv q^*. \end{aligned} \quad (1.2)$$

Hence, player i will play A in period t if and only if $q_i \geq q^*$ (assuming player i chooses A when he is indifferent). The assumption that (A, A) is the risk-dominant equilibrium implies that $q^* < \frac{1}{2}$. I will frequently discuss the behavior of the model with payoffs $a = 2$, $b = 1$, and $c = d = 0$ so that $q^* = \frac{1}{3}$ and player i will play A if at least $\frac{1}{3}$ of his opponents did so.

At time t , we describe the state of the system by $s_t \in S^u = \{0, 1, \dots, N\}$, with s_t indicating the total number of players playing A . The dynamics of the model with no noise are straightforward:

If $s_t < \lceil q^*(N - 1) \rceil$, then $q_i < q^*$ for all i so all players play B in period $t+1$ and $s_{t+1} = 0$.

If $s_t > \lceil q^*(N - 1) \rceil$, then $q_i > q^*$ for all i so all players play A in period $t+1$ and $s_{t+1} = N$.

If $s_t = \lceil q^*(N - 1) \rceil$, we have a knife edge case where $q_i < q^*$ if player i played A in period t , and $q_i \geq q^*$ if player i played B in period t . The result is that $s_{t+1} = N - \lceil q^*(N - 1) \rceil$ as only those players who played B in period t play A in period $t+1$.

The last transition results from the assumption that each player reacts to the play of his possible opponents and that this group does not include himself. While unsightly, this transition does not play a significant role in the subsequent analysis so the reader should not be too troubled by it.

The important thing to note is that the model with uniform matching and no noise has two steady states 0 and N corresponding to the Nash equilibria where all players coordinate on one of the two possible strategies. Note that each steady state has a large basin of attraction. When play starts close to either steady state it immediately jumps to that equilibrium. Once noise is introduced, the transitions are governed by a Markov process which assigns positive probability to any transition. Nonetheless, when play is near one equilibrium it will likely remain near that equilibrium for a long period of time.

Suppose that most players played B in period t so that $s_t < \lceil q^*(N - 1) \rceil$. Each player will then play B in period $t + 1$ with probability $1 - \epsilon$. All of these randomizations are independent, so with large populations it is extremely likely that the fraction of players who play A will be close to ϵ . We will generally envision ϵ to be much smaller than q^* . In this case we will very likely have $s_{t+1} < \lceil q^*(N - 1) \rceil$ so that the same reasoning again applies to describe the period $t + 2$ play. The important thing to note is that no gradual change is possible. The only mechanism for a shift from one equilibrium to the other is the sudden jump which results when at least $\lceil q^*(N - 1) \rceil$ ϵ -probability events occur simultaneously. This coincidence occurs only very infrequently. Hence, the players' strategies will likely resemble those determined by the initial conditions for a long period of time.

In models of local interaction, we denote the possible states by N -tuples $(a_1, a_2, \dots, a_N) \in S^t = \{A, B\}^N$, in order to keep track of the locations of the players using each strat-

egy in addition to the aggregate frequencies. To illustrate the dynamics of such models, I discuss a typical setup where N players are arranged uniformly around a circle and each places equal weight on being matched with his eight closest neighbors. Let the payoffs be given by

	A	B
A	2,2	0,0
B	0,0	1,1

so that each player has A as his best response whenever at least three of his eight neighbors play A .

Once again, I begin by describing the dynamics in the model with no noise. Clearly there are at least two steady states, $\vec{A} = (A, A, \dots, A)$ and $\vec{B} = (B, B, \dots, B)$. Each of these steady states has a non-trivial basin of attraction. Suppose that all but one or two of the players are playing B at time t . Then, each player has at least six neighbors playing B so all will play B in period $t + 1$. We may write two such transitions as

$$\begin{aligned} (A, B, B, \dots, B) &\longrightarrow \vec{B} \\ (A, A, B, \dots, B) &\longrightarrow \vec{B}. \end{aligned}$$

Similarly, for period t states sufficiently close to \vec{A} we will get an immediate jump to \vec{A} in period $t + 1$.

An important feature of the dynamics is that the basin of attraction of \vec{B} is relatively small. In particular, the existence of a small cluster of players playing A is sufficient to ensure that the dynamic process will eventually lead all players to play A . Suppose the period t state is $(A, A, A, A, B, \dots, B)$ so that players 1 through 4 played A . It is easy to see that players 1 through 6 and players N and $N - 1$ all have

at least three neighbors playing A . Those eight players will play A in period $t + 1$. In period $t + 2$, players $N - 2$, $N - 3$, 7, and 8 will switch to playing A . The cluster of players playing A will grow until eventually the state \vec{A} is reached.

In contrast, any relatively small cluster of players playing B will disappear over time. It is easy to verify that

$$\begin{aligned} (B, B, B, B, A, A, A, \dots, A) &\longrightarrow \vec{A} \\ (B, B, B, B, B, B, B, A, \dots, A) &\longrightarrow (A, A, B, B, B, A, \dots, A) \longrightarrow \vec{A}. \end{aligned}$$

The basin of attraction of \vec{A} is much larger than that of \vec{B} .

It is the differing sizes of these basins of attraction which causes the relatively rapid convergence of play to a limit concentrated around \vec{A} once noise is introduced. From the dynamics above, it should be clear that we usually only need to wait for four well placed randomizations to create a cluster of players playing A and lead us away from an initial condition where everyone is playing B . When the number of players is large, seeing four adjacent randomizations is far more likely than seeing the $\lceil \frac{N-1}{3} \rceil$ simultaneous randomizations required to shift play in the model with uniform matching.

The extreme local matching rule in which each player has only two neighbors is neither an apt description of any of the examples I have given, nor does it have particularly compelling dynamic behavior. Nonetheless, it is the easiest model of local interaction to analyze, and hence will reappear throughout this paper. For this reason, I briefly discuss its dynamics here.

First, note that regardless of the payoffs, the assumption that (A, A) is the risk-dominant equilibrium entails that each player will have A as his best response whenever at least one of his two neighbors plays A . In a model with no noise, we again have two steady states, \vec{A} and \vec{B} . There is also one stable cycle when N is even,

$$(A, B, A, B, \dots, A, B) \longrightarrow (B, A, B, A, \dots, B, A) \longrightarrow (A, B, A, B, \dots, A, B),$$

whose existence is an unfortunate but not particularly significant byproduct of the assumption that players in period t myopically respond to their opponents play in period $t - 1$, not to a forecast of their period t play. I shall write $\vec{A}\vec{B}$ as a shorthand for the state $(A, B, A, B, \dots, A, B)$ and $\vec{B}\vec{A}$ for the state $(B, A, B, A, \dots, B, A)$. The most important aspect of the dynamics to note is that the steady state \vec{B} now has no other states in its basin of attraction. If at least one player plays A in period t , then at least two players (his neighbors) will play A in period $t + 1$. In contrast, any state which contains a cluster of two adjacent players playing A lies in the basin of attraction of \vec{A} . Once noise is introduced, we will see that this leads to rapid convergence to a steady state concentrated around \vec{A} .

1.4 Limits and Rates of Convergence

In this section, I discuss the principal theoretical results of the paper. As mentioned above, the motivation for the analysis here is the assumption that at some point the initial actions of the players were determined by historical factors and that for some subsequent period of time play has evolved according to the behavioral rules specified above. The fundamental problem is then to determine how historical and evolutionary forces combine to determine the play we observe. To this end, I first discuss the limiting behavior of these systems as the number of periods of evolution grows to infinity. Subsequently, I discuss the rates at which the limits are approached in order to assess whether the limits are meaningful given that the economic systems modeled involve only some reasonable finite repetition of play.

In this section, I contrast the behavior of the model under the extreme assumptions of uniform and two neighbor matching. I make this choice for analytic tractability and will use numerical simulations to discuss the behavior of models with alternate specifications of the matching process in the next section.

As noted above, we may view the time t strategy profiles as the states s_t of a Markov process. In the case of uniform matching, $s_t \in \{0, 1, \dots, N\}$ indicates the number of players playing A . At times, I shall write \vec{A} for state N and \vec{B} for the

state 0. We may represent the time t probability distribution over the states by an $(N+1) \times 1$ vector v_t . For example, $v_t = (0.5, 0.5, 0, \dots, 0)'$ represents a system which is equally likely to be in state 0 and state 1. The evolution of the process is governed by

$$v_{t+1} = P^u(\epsilon)v_t$$

where $P^u(\epsilon)$ is the transition matrix whose elements are given by

$$p_{ij}^u(\epsilon) = \text{Prob} \{s_{t+1} = i | s_t = j\}.$$

For example, we have

$$p_{ij}^u(\epsilon) = \left(\frac{N}{i}\right)\epsilon^i(1-\epsilon)^{N-i}$$

whenever $j < [q^*(N-1)]$ as each player's best response is B so state i can arise only when exactly i ϵ -probability randomizations occur. Note that $P^u(\epsilon)$ is strictly positive for $\epsilon > 0$ so by standard results on Markov processes there is an unique steady-state distribution $\mu^u(\epsilon)$ such that

$$\mu^u(\epsilon) = P^u(\epsilon)\mu^u(\epsilon).$$

For any initial probability distribution ρ , the distribution of period t play is given by $P^u(\epsilon)^t\rho$. It is also standard that positive finite-state Markov processes are ergodic, *i.e.*

$$P^u(\epsilon)^t\rho \rightarrow \mu^u(\epsilon)$$

as $t \rightarrow \infty$ so that the steady-state distribution represents the distribution of play after infinitely many periods of transitions. I shall write $\mu_s^u(\epsilon)$ for the probability assigned to state s by the steady-state distribution $\mu^u(\epsilon)$.

We may define $P^l(\epsilon)$ and $\mu^l(\epsilon)$ analogously for the model with two neighbor matching. It is harder, however, to visualize the structure of the transition matrix in this case as $P^l(\epsilon)$ now acts on 2^N dimensional vectors and there is no intuitively appealing ordering of the states.

The first result compares the steady state distributions of the uniform and two neighbor models. KMR have shown that the evolutionary forces in a model virtually identical to the uniform model yield a steady state limit in which the risk-dominant equilibrium (A, A) is played with very high probability. The statement that $\mu_A^u(\epsilon) \rightarrow 1$ and $\mu_A^l(\epsilon) \rightarrow 1$ verifies that the steady state also exhibits play concentrated around the risk-dominant equilibrium both in the uniform model I have defined and in the model with two neighbor matching. The second part of the theorem compares the steady state probabilities with which the entire population coordinates on the equilibrium (B, B) . For sufficiently small ϵ , the theorem tells us that this equilibrium is even less common in the model with two neighbor matching than it is under uniform matching (although it is extremely rare in both models).

Theorem 1 *Let $\mu^u(\epsilon)$ and $\mu^l(\epsilon)$ be the steady state distributions of the general model of Section 2 under the uniform and two neighbor matching rules, respectively. Let q^* be as defined in (1.2) with $\lceil q^*(N-1) \rceil < N/2$. Then,*

(a)

$$\lim_{\epsilon \rightarrow 0} \mu_A^u(\epsilon) = 1$$

$$\lim_{\epsilon \rightarrow 0} \mu_A^l(\epsilon) = 1$$

(b)

$$\begin{aligned} \mu_B^u(\epsilon) &= O(\epsilon^{N-2\lceil q^*(N-1) \rceil+1}) \\ \mu_B^l(\epsilon) &= \begin{cases} O(\epsilon^{N-2}) & \text{for } N \text{ even} \\ O(\epsilon^{N-1}) & \text{for } N \text{ odd} \end{cases} \end{aligned}$$

Proof

Both the statement and the proofs of the results for the uniform model are virtually identical to those given in KMR, and hence I omit the proofs. The proof for the two neighbor model relies on the following characterization of the steady state cited in KMR. The reader may refer to that paper or to Freidlin and Wentzel (1984) for an exposition of the background material. An x -tree t on S is a function $t : S \rightarrow S$

such that $t(x) = x$ and such that for all $s \neq x$ there exists m with $t^m(s) = x$. We may think of an x -tree as a set of arrows connecting elements of S in which every element has a unique successor and all paths eventually lead to x . The steady state distribution $\mu^\epsilon(\epsilon)$ can be characterized by

$$\mu_x^\epsilon(\epsilon) = c(\epsilon) \sum_{t \in H_x} \prod_{i \neq x} p_{t(i)i}^\epsilon(\epsilon)$$

where H_x is the set of x -trees on S^ϵ . Note that $p_{ji}^\epsilon(\epsilon)$ is a polynomial in ϵ whose constant term is non-zero if and only if the transition $i \rightarrow j$ occurs in the model with no noise ($\epsilon = 0$). For any state x , the expression above allows us to express the quantity $\mu_x^\epsilon(\epsilon)/\mu_x^\epsilon(\epsilon)$ as a ratio of polynomials in ϵ . The order of this ratio can then be discerned by looking at the smallest order x - and \vec{A} -trees. I go through the details of this for N even.

An \vec{A} -tree s with $\prod_{i \neq \vec{A}} p_{t(i)i}^\epsilon(\epsilon) = O(\epsilon^2)$ is given by

$$\begin{aligned} t(\vec{B}) &= (A, B, B, \dots, B) \\ t(\vec{A}\vec{B}) &= (B, A, A, A, B, A, \dots, B, A) \end{aligned}$$

with all other states mapped to their successors in the no noise model. Clearly there are no $O(\epsilon)$ \vec{A} -trees. To prove (a) we need only show that an x -tree for any other x is of strictly higher order so that $\mu_x^\epsilon(\epsilon)/\mu_x^\epsilon(\epsilon) \rightarrow 0$. We can prove this by considering three cases.

For $x \notin \{\vec{B}, \vec{A}\vec{B}, \vec{B}\vec{A}\}$ an x -tree is of order at least ϵ^3 as at least one ϵ -probability transition is necessary to break out of each steady state or cycle. For $x = \vec{B}$, a \vec{B} -tree t with $\prod_{i \neq \vec{B}} p_{t(i)i}^\epsilon(\epsilon) = O(\epsilon^N)$ is given by

$$\begin{aligned} t(\vec{A}) &= \vec{A}\vec{B} \\ t(\vec{A}\vec{B}) &= \vec{B} \end{aligned}$$

with all other states again mapped to their successors in the no noise model. Clearly there are no \vec{B} -trees of lower order as there must be a path from \vec{A} to \vec{B} in any

\vec{B} -tree. A transition on such a path requires at least k ϵ -probability events whenever the number of players playing A decreases by k as the number of players playing A never decreases in the model with no noise. Finally, for $x = \vec{A}\vec{B}$ or $\vec{B}\vec{A}$, a similar argument shows that the minimum order x -tree has order $\epsilon^{N/2+1}$. As $N > 2$, this concludes the proof of (a).

Combining the characterization of the \vec{A} -trees and \vec{B} -trees gives (b). **QED.**

The theorem above implies that if the coordination games we have described are repeated enough times, play will eventually become concentrated around the risk-dominant equilibrium (A, A) . It remains to be seen, however, whether this 'eventually' is relevant. Large annual trade fairs in medieval England continued for perhaps eight centuries. For modern economic applications, we shall want to discuss annual events repeated far fewer times. Even for weekly interactions over several decades we will be limited to a few thousand repetitions. To discuss play in these finite games, we must then ask whether such numbers of repetitions are sufficiently large for play late in the game to resemble the steady-state limit.

I now begin a discussion of this problem with some theoretical results on the rates of convergence of the uniform and two neighbor models. For very small probabilities ϵ of randomization, I find a striking contrast between the two models, with far slower convergence in the uniform model. This suggests that observed play is far less likely to resemble the steady state in the model with uniform matching than it is with local matching.

To discuss the problem formally, let Δ be the set of probability distributions on S . For any two distributions $\mu, \nu \in \Delta$ define

$$\|\mu - \nu\| \equiv \max_{s \in S} |\mu_s - \nu_s|.$$

If P is a Markov transition matrix with steady state distribution μ , $\|P^t \rho - \mu\|$ measures how far from the steady state the system is t periods after we begin with initial distribution ρ . When this distance is small, the steady state is a good predictor for period t play. When the distance is large, it is not. It is well known that finite state

Markov processes converge at an exponential rate. Informally, we may think of this as saying that $\|P(\epsilon)^t \rho - \mu(\epsilon)\| \sim c_0 r^t$ for some $r < 1$. While exponential convergence is usually thought of as a rapid theoretical standard, convergence may in fact be quite slow for practical applications. If, for example, $r = 0.9999999$, then $r^{1,000,000} \approx 0.9$ so that after one million periods only 10% of the distance between the initial distribution and the steady state will have been eliminated.

The following theorem characterizes the rates of convergence for the uniform and two neighbor models.

Theorem 2 *Let $P^u(\epsilon)$ and $P^\ell(\epsilon)$ be the transition matrices for the uniform and two neighbor models and let $\mu^u(\epsilon)$ and $\mu^\ell(\epsilon)$ be the associated steady state distributions. Assume $\lceil q^*(N-1) \rceil \leq N/2$. For Δ the set of probability distributions on S define*

$$r^u(\epsilon) \equiv \sup_{\rho \in \Delta^u} \limsup_{t \rightarrow \infty} \|P^u(\epsilon)^t \rho - \mu^u(\epsilon)\|^{1/t}$$

$$r^\ell(\epsilon) \equiv \sup_{\rho \in \Delta^\ell} \limsup_{t \rightarrow \infty} \|P^\ell(\epsilon)^t \rho - \mu^\ell(\epsilon)\|^{1/t}$$

Then,

$$1 - r^u(\epsilon) = O(\epsilon^{\lceil q^*(N-1) \rceil})$$

$$1 - r^\ell(\epsilon) = O(\epsilon)$$

Proof

To simplify the right hand side of the expressions defining $r^u(\epsilon)$ and $r^\ell(\epsilon)$ we make use of two results from the Frobenius Theory of positive matrices.⁷ Let P be any strictly positive transition matrix and let μ be its unique steady state. The first result is that if we order the eigenvalues of P so that $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_N|$, then $\lambda_1 = 1$ and $|\lambda_2| < 1$. From this follows a result directly applicable to our problem, namely that

$$\sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|P^t \rho - \mu\|^{1/t} = |\lambda_2|.$$

⁷See Karlin and Taylor (1975) pp. 542 - 551.

A formal proof is given in Karlin and Taylor (1975). In the special case of P diagonalizable, we can obtain the second result from the first by exploiting the diagonalization

$$P = \Phi \Lambda \Psi \quad \Psi = \Phi^{-1}$$

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_N \end{bmatrix}$$

Note that for any fixed ρ ,

$$P^t \rho = \Phi \Lambda^t \Psi \rho.$$

Hence,

$$P^t \rho - \Phi \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \Psi \rho = \Phi \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & \lambda_2^t & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_N^t \end{bmatrix} \Psi \rho.$$

The right hand side of this expression converges to zero. Writing E_{ij} for the matrix whose ij^{th} element is 1 and with all other elements equal to zero, we have from the uniqueness of the steady state that

$$\Phi E_{11} \Psi \rho = \mu$$

for all distributions ρ . Hence,

$$\begin{aligned} \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|P^t \rho - \mu\|^{1/t} &= \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|P^t \rho - \Phi E_{11} \Psi \rho\|^{1/t} \\ &= \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\lambda_2^t \Phi \left(\frac{\Lambda^t - E_{11}}{\lambda_2^t} \right) \Psi \rho\|^{1/t} \\ &= |\lambda_2| \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\Phi E_{22} \Psi \rho\|^{1/t} \\ &= |\lambda_2|. \end{aligned}$$

We also see from this calculation that the supremum is in fact achieved for any ρ such

that $\Phi E_{22} \Psi \rho \neq 0$.

Given this result, the problem of finding $r^u(\epsilon)$ and $r^l(\epsilon)$ is reduced to the problem of finding the second largest eigenvalues of the matrices $P^u(\epsilon)$ and $P^l(\epsilon)$. The remainder of the proof is the rather lengthy solution to this problem and can be found in the appendix. **QED.**

The force of Theorem 2 is that when ϵ is small, convergence will be much slower in the model with uniform matching than it is in the model with two neighbor matching. The simplest way to see this is to take a numerical example. Suppose there are 100 players with payoffs $a = 2, b = 1, c = d = 0$ so that $\lceil q^*(N - 1) \rceil = 33$. As a thought experiment, suppose we start with a fairly small randomization probability ϵ and consider the effect of reducing ϵ to $\epsilon/2$. In the two neighbor model, the result is that the model takes about twice as long to converge. The model with uniform matching, however, exhibits more extreme behavior. Theorem 2 tells us that $1 - r^u(\frac{\epsilon}{2}) \approx 2^{-33}(1 - r^u(\epsilon))$. If $1 - r^u(\epsilon)$ is small we have the first order approximation

$$r^u\left(\frac{\epsilon}{2}\right)^{2^{33}} \approx \left(1 - \frac{1 - r^u(\epsilon)}{2^{33}}\right)^{2^{33}} \approx r^u(\epsilon).$$

Hence, convergence to within a given tolerance in the uniform model will now take not twice as many periods but rather 2^{33} or over 8 billion times as many. Clearly, one does not have to divide ϵ by two very many times before the powers of 8 billion far outweigh any other factors. The theorem is then describing a vast difference in rates of convergence for ϵ small.

It is also not hard to understand why something like Theorem 2 must be true. As described in Section 3, the only mechanism for change in the uniform model is the simultaneous randomization by at least $\lceil q^*(N - 1) \rceil$ players. When N is large, ϵ small and $q^* - \epsilon$ non-trivial, the law of large numbers tells us that we are depending on an extremely unlikely event. In particular, as $\epsilon \rightarrow 0$ the probability of this event is $O(\epsilon^{\lceil q^*(N-1) \rceil})$. For very small values of ϵ the very slow convergence reflects our intuition that in large populations this simultaneous randomization is not a plausible

mechanism for change. In contrast, regardless of the population size, only two randomizations are necessary to begin a shift from \vec{B} to \vec{A} in the two neighbor model. This provides the relatively rapid convergence.

I now address two further questions in order to provide a more complete understanding of the behavior of the models. First, note that so far I have principally commented on the relative speeds of convergence of the two models. Presumably, it could still be the case that both models converge very slowly or that both converge very quickly (although the latter seems unlikely). Second, one should not overemphasize limiting behavior. In the examples discussed, it seems reasonable that the players do make mistakes or adopt strategies in what can be viewed as a random manner a non-trivial fraction of the time. In order to comment on the applications, we must be sure that the results given reflect the behavior of the model not just for infinitesimal ϵ but for values of ϵ like 0.1 as well.

To see that convergence in the uniform model is indeed very slow it will suffice to compute the rate of convergence $r^u(\epsilon)$. Table 1.1 gives the value of $1 - r^u(\epsilon)$ for various population sizes N and randomization probabilities ϵ . In each case the payoffs are fixed with $q^* = \frac{1}{3}$ and $r^u(\epsilon)$ was obtained by solving (1.4).

Table 1.1: Rates of Convergence for the Uniform Model

	$1 - r^u(\epsilon)$	
	$\epsilon = 0.05$	$\epsilon = 0.1$
$N = 10$	1.15×10^{-2}	7.02×10^{-2}
$N = 100$	1.23×10^{-18}	3.23×10^{-10}
$N = 1000$	5.11×10^{-174}	2.24×10^{-89}

It is clear from Table 1.1 that for the larger population sizes $1 - r^u(\epsilon)$ is both very close to zero and decreasing very rapidly as ϵ is reduced. If we choose an initial distribution ρ which differs from $\mu^u(\epsilon)$ in the direction of the eigenvector associated

with $r^u(\epsilon)$ (a vector largely adding weight on \vec{B} relative to \vec{A}) then

$$\|P^u(\epsilon)^t \rho - \mu^u(\epsilon)\| = r^u(\epsilon)^t \|\rho - \mu^u(\epsilon)\|$$

If, for example, we have $N = 1000$ and $\epsilon = 0.1$ then $r^u(\epsilon) = 1 - 2.24 \times 10^{89}$ and $r^u(\epsilon)^{10^{86}} \approx 0.998$. Hence, after 10^{86} periods of play the distance between the initial distribution of play and the steady state has only been reduced by less than one quarter of one percent. For a game with one hundred players, a similar comment applies to the cumulative evolutionary force of 10 million periods of play.

Kandori, Mailath and Rob discuss an alternate, and perhaps more intuitive, measure of the rate of convergence, the expected number of periods necessary for the system to first enter the basin of attraction of \vec{A} given that play starts with everyone playing B . From the limited number of values in their table, it is not clear how rapidly the wait increases as ϵ becomes small and N large. When ϵ is small, the waiting times are very closely approximated by $1/(1 - r^u(\epsilon))$. Table 1.2 gives the expected waiting time for two different sets of payoffs: the first $a = 2, b = 1, c = d = 0$ as before and the second a more extreme case $a = 5, b = 1, c = d = 0$ where the payoff to (A, A) is far greater than the payoff to (B, B) . The expected waiting time is computed as

$$\frac{(q_1 + q_2) + p_1(1 - (q_1 + q_2))}{(q_1 + q_2)p_2 + p_1q_2},$$

with $p_1, p_2, q_1,$ and q_2 as in (1.3).

For the population of only ten players, these waiting times are not unreasonable, and we thus might expect to see movement to the risk-dominant equilibrium. For a population of fifty or one hundred, though, we would expect to see evolution take place *only* in the case when the payoff to (A, A) is much greater than the payoff to (B, B) , and a large fraction of the population randomizes in each period. Even when the payoff to (A, A) is five times the payoff to (B, B) , it must be conceivable that one sixth of the entire population will switch to A randomly in the same period. For reasonably small randomization probabilities and large population sizes, this is very unlikely. Hence, the waiting times are too long for us to expect to see such a transition

Table 1.2: Expected Waiting Times with Uniform Matching

	Expected Wait ($a = 2$)		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
$N = 10$	609	87	14
$N = 50$	3.83×10^{14}	6.54×10^9	2.63×10^5
$N = 100$	1.30×10^{27}	8.13×10^{17}	3.09×10^9
$N = 1000$	1.09×10^{268}	1.96×10^{173}	4.46×10^{88}

	Expected Wait ($a = 5$)		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
$N = 10$	41	12	4
$N = 50$	2.65×10^5	1323	17
$N = 100$	1.86×10^9	1.06×10^5	49
$N = 1000$	1.61×10^{82}	1.82×10^{41}	2.16×10^{10}

in practice. If we believe that the model with uniform random matching describes the behavior of any large population, we will conclude that the dynamics will produce a pattern of play exhibiting great inertia. Over an economically reasonable time period we should expect players to continue to play whatever equilibrium first arises, regardless of considerations of Pareto optimality or risk dominance.

In comparison, let us now return to the model with two neighbor matching. While it is difficult to compute the eigenvalues or waiting times analytically, it is fairly easy to estimate the waiting times via numerical simulations. Table 1.3 reports Monte Carlo estimates of the expected waiting time until more than 75% of the players play \tilde{A} in the same time period. The simulations were designed so that each estimate has a standard error of 0.1 or less. The parameters N and ϵ reported are the same as in the table above. (The waiting time is now independent of q^* for any $q^* \in (0, \frac{1}{2})$).

Table 1.3: Expected Waiting Times with Two Neighbor Matching

	Randomization Probability ϵ		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
$N = 10$	14.5	9.0	6.2
$N = 50$	11.0	8.1	6.3
$N = 100$	11.1	8.2	6.4
$N = 1000$	11.0	8.1	6.2

The most striking result in Table 1.3 is that the waiting times are not only small relative to those of the uniform model but in fact are very short. For each of the parameter values, the expected wait is less than twelve periods. If we believed that the two neighbor model on a circle were an apt description of reality, we would conclude that powerful evolutionary forces will cause the risk-dominant equilibrium to arise fairly soon after the start of the game. We could easily imagine, however, that for more reasonable matching rules convergence will be significantly slower. The next section examines some aspects of this problem. For now, I will conclude only that with local matching the observed play need not reflect the initial conditions and

that the learning process may lead to play concentrated around the risk-dominant equilibrium.

1.5 Simulation Results

The theorems of the previous section show that the evolution of play will be very slow in the model with uniform matching. In contrast, under the extreme assumption of two neighbor matching, convergence will be very fast. In this section, I use numerical simulations to investigate the extent to which this result carries over to specifications of the model which can be more reasonably applied to the economic examples of local matching mentioned in the introduction. To this end, I analyze the importance of the various assumptions I have made and discuss extensions.

1.5.1 Matching Rules

The most obvious shortcoming of Section 4 for practical applications is the limitation of the discussion to the two neighbor matching rule. In practice, the interactions within a population will never be this simple. The question I address here is whether this is likely to affect the conclusions drawn from the model.

Before presenting any simulation results it is important to recall the pattern of evolution with local matching described in Section 3. The shift to the risk-dominant equilibrium will usually begin when a small cluster of adjacent players switch to this new strategy. This small community must be sufficiently inward looking for its members to be satisfied playing the new strategy regardless of the play of the rest of the population. I have earlier noted that the existence of such clusters depends on three features of the model. First, each player must place a large weight on matching with a few close neighbors. Second, there must be a large degree of overlap among the groups of neighbors. Finally, the structure of the matching rule must be stable over time. Let me now examine each of these assumptions in a little more detail.

In each of the examples of local matching I have discussed, it is clear that each player will usually be matched with one of a few close friends or colleagues. While an

economist, for example, would probably have more than two colleagues with whom he frequently shares software, the correct number is almost surely no larger than ten or twenty. As in this example, we will typically have only a rough idea of the matching rule involved. We would hope then that the predictions of the model would be fairly robust to the particular specification chosen. Recall that with two neighbor matching we have seen that play shifts very quickly to the risk-dominant equilibrium. Table 1.4 investigates the extent to which this remains true for less concentrated matching rules. The table gives Monte Carlo estimates of the expected waiting time until more than 75% of the population shifts from playing B to playing the risk-dominant A . In each case, a population of 100 players and payoffs with $q^* = 1/3$ (e.g. $a = 2, b = 1, c = 0, d = 0$) are assumed. The matching rules labelled k neighbors place equal weight on a player's k closest neighbors on a circle. The rule labelled $1/2^{|i-j|}$ is that given in (1.1).

Table 1.4: Expected Waiting Times with Various Matching Rules*

	Randomization Probability ϵ		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
2 Neighbors	11.0 (0.1)	8.2 (0.1)	6.4 (0.1)
4 Neighbors	41.9 (1.2)	22.7 (0.4)	12.3 (0.2)
8 Neighbors	98.2 (7.5)	25.4 (0.8)	10.7 (0.2)
12 Neighbors	460.4 (42.7)	46.7 (3.4)	11.3 (0.3)
$1/2^{ i-j }$	39.5 (1.2)	20.0 (0.5)	10.3 (0.2)

* Standard Errors in Parentheses

Note that the effect of the matching rule varies with the frequency of the randomizations. For the smallest value of ϵ shown, waiting times increase significantly when players have more neighbors. In this case, evolution is only likely to be seen when

the matching rule is concentrated on a very few neighbors. For the larger values of ϵ , the waiting times are shorter and less dependent on the particular matching rule. Of course, whether we consider something like fifty periods to be a short time or a long time depends on the time frame of the particular application. I would nonetheless conclude that for a variety of local matching rules dynamic forces may lead players to coordinate on the risk-dominant equilibrium within a reasonable period of time.

The matching rules which I have discussed so far are far from general. Note in particular that I have maintained the assumption that the players are arranged around a circle.⁸ Inherent in this assumption is a great overlap of the groups of neighbors so that a player's neighbors' neighbors are likely to be his neighbors as well. This clearly facilitates the formation of stable clusters playing the risk-dominant equilibrium. The true geometry of social interactions will surely be more complex with less overlap of the neighborhoods. In the economist example, a player's neighbors will likely consist of some colleagues within her own department, colleagues at other schools with similar research interests and perhaps a cadre of older friends from graduate school. While a friend from graduate school would likely list the cadre of graduate school friends among his potential co-workers as well, he would likely have very little interaction with the other two groups. We would, therefore, like for our results to be robust to changes in the structure of the matching rule.

Rather than trying to propose a reasonable geometry for the interactions within a particular population, I simply examine a few alternate specifications which entail less overlap between groups of neighbors. Instead of envisioning a population of 400 players as arranged around a circle, for example, we can imagine them as arranged at the vertices of a 20×20 lattice on the surface of a torus. Figure 1-1 pictures part of such a lattice and the sets of neighbors which would give four and eight neighbor

⁸Blume (1991) discusses the effects of the lattice used to represent the matching process in a similar model and shows that different lattices can lead to strikingly different theoretical properties. Most significantly, for certain matching rules the Markov process may be ergodic for a countable population arranged along a line, but nonergodic when the population is arranged on a higher dimensional lattice. In our finite population framework, the analog of nonergodic behavior is very long waiting times for large populations.

matching rules. Note that with the eight neighbor matching rule moving from one to two dimensions has reduced the overlap of groups of neighbors to the point that no two players have more than four neighbors in common. If we move to a four dimensional lattice, no two players have more than two neighbors in common.

Figure 1-1: Four and Eight Neighbor Matching in Two Dimensions




Table 1.5 again gives estimates of the expected waiting time until 75% of the population shifts to the risk-dominant equilibrium in a population of 400 players with payoffs such that $q^* = 1/3$. The table compares the waiting times when players are arranged on three different lattices: a circle (labelled 400×1), a 20×20 lattice on the surface of a torus, and a four dimensional $4 \times 4 \times 5 \times 5$ lattice. For most of the parameter values shown, the expected waiting times are not greatly increased as we reduce the overlap of the neighbor groups by moving to higher dimensional lattices. However, the increased waiting times in the eight neighbor model for $\epsilon = 0.025$ are a reminder that the structure of the matching rule has the potential to greatly affect the behavior of the model.

So far, I have considered two aspects of the matching rule, the number of neighbors or likely opponents each player has and the structure of the interconnections among the players. For most applications, we will have only a rough knowledge of these aspects of the model and hence would hope that the predictions of the model are similar over a range of reasonable parameter values. Fortunately, this does seem

Table 1.5: Expected Waiting Times for Different Geometries*

Lattice	Four Neighbor Matching		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
400×1	45.8 (0.8)	23.3 (0.3)	12.4 (0.1)
20×20	43.0 (1.2)	21.3 (0.4)	11.5 (0.1)

Lattice	Eight Neighbor Matching		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
400×1	69.5 (2.0)	27.5 (0.5)	11.1 (0.1)
20×20	121.7 (14.7)	21.0 (0.5)	8.9 (0.1)
$4 \times 4 \times 5 \times 5$	1739.5 (382.6)	31.9 (1.7)	8.4 (0.1)

* Standard Errors in Parentheses

to be the case. There are, however, many details in the model which do affect the rate at which play evolves toward the risk-dominant equilibrium. Particularly when randomizations are infrequent, there are reasonable models in which play converges fairly slowly. As a result, we cannot predict that the risk-dominant equilibrium will arise in these models as confidently as we were able to predict that historical factors are likely to determine play in the model with uniform matching.

At this point, I discuss one more aspect of the matching process, the assumption that the matching rule remains constant over time. For some populations, this is reasonable. For example, when someone is deciding whether to attend a college reunion he will want to base his decision on whether the people who were his close friends when he was in college are attending. This group should not change greatly as the years pass.

At the other extreme, we can envision populations where there is constant change in the group of likely opponents. One example might be a population of antique traders who meet at weekend antique fairs. We can envision them as playing a matching game in the context of our model as follows. At the start of each week, each trader acquires an antique for which he has no use, but which he may be able to sell to one of the five or ten other traders who collect such items. Like our medieval traders, he must then decide independently which of two possible fairs to attend. After arriving at the fair, he tries to sell the antique and in doing so receives a larger expected payoff when more of the interested collectors have chosen to attend the same fair. This situation is similar to our local matching rules in that once the antique has been acquired, a trader cares only about the locational choices of the few potential buyers he can identify. The crucial difference, however, is that this group of neighbors changes each period when a new antique is to be sold.

The stability of groups of neighbors over time is necessary for the types of evolutionary patterns described in Section 3. Look, for example, at the eight neighbor matching rule with all players initially playing *B*. With stable groups of neighbors a cluster of four adjacent players playing *A* will soon arise and expand until nearly

everyone plays A . In contrast, we can model constantly changing groups of neighbors like those in the antiques example by envisioning the players as being randomly relocated around the circle at the start of each period. In this case, small clusters of adjacent players no longer provide stable inward looking communities where playing A can gain a foothold. At the start of each period, those players playing A are randomly scattered. As a result, any clusters which have formed will usually be broken up and subsequently disappear.

Table 1.6 compares the behavior of our standard model with a fixed local matching rule to the behavior of a similar model in which the players' locations on the circle are chosen randomly at the start of each period. Each simulation was run for a model with eight neighbor matching and with payoffs which imply $q^* = 1/3$. The time until 75% of the population first plays A is recorded. For very small populations, *e.g.* 10 players, the assumption of fixed sets of neighbors is not very important as there is only limited room for clusters to be broken up. In larger populations, however, most clusters will be broken up and hence the assumption of fixed neighbors is very important, particularly when randomizations are rare so that new clusters do not often arise. At the extreme of constantly changing sets of neighbors, evolution may be very slow in large populations.

For most applications, the reasonable matching rule will fall somewhere between the two extremes. In a population of economists, for example, there is clearly some change in the sets of co-workers over time, but the truth is probably much closer to the assumption of stable relationships than it is to the other extreme. In this case, we would still expect to see fairly rapid convergence to the risk-dominant equilibrium.

1.5.2 Population Sizes

In each of the applications mentioned in the introduction, there are a large number of players. Whether this means fifty players or several hundred players, though, is not certain. If we intend for the theory to capture what we think are characteristics of play common to all large populations, we should hope again that the predictions of the model are also fairly consistent over a range of reasonable population sizes. This

Table 1.6: Effect of Changing Locations over Time*

10 Players			
Matching Rule	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
Fixed	726.6 (105.3)	97.3 (14.7)	16.0 (1.6)
Random Locations	721.9 (95.0)	87.4 (12.1)	15.7 (1.7)


50 Players			
Matching Rule	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
Fixed	150.1 (17.2)	32.2 (3.3)	10.6 (0.5)
Random Locations	5020.0 (567.5)	108.0 (13.2)	9.5 (0.6)

* Standard Errors in Parentheses

is the case for the models of local matching discussed above. Figure 1-2 illustrates the relationship between the size of the population and the expected waiting time until 75% of the population has switched from playing *B* to playing *A* for three different matching rules. The waiting times pictured are estimates for $\epsilon = 0.05$ and payoffs such that $q^* = 1/3$.

The behavior illustrated above is quite different from that in the uniform model. For very small populations, the waiting times may be fairly long. Waiting times are usually decreasing in population size, however, and appear to be remarkably constant for all population sizes above about 100. Intuitively, what is happening is that there is little change in play until the first sufficiently large cluster of neighbors playing *A* forms. When the population is larger, there are more possible locations and therefore more chances for such a cluster to occur randomly. Hence, the wait until the first cluster appears is shorter. Eventually though, further increases in the population size can no longer greatly speed evolution. Whether a typical individual will have switched to playing *A* within the first, say, 40 periods depends not on whether a cluster has formed anywhere in the population but rather on whether a cluster has

Figure 1-2: Estimated Waiting Times for Various Population Sizes



formed sufficiently close to him so that it can grow and reach him within 40 periods. The important conclusion is that the results on the rates of convergence of models with local matching are applicable for a wide range of large population sizes.

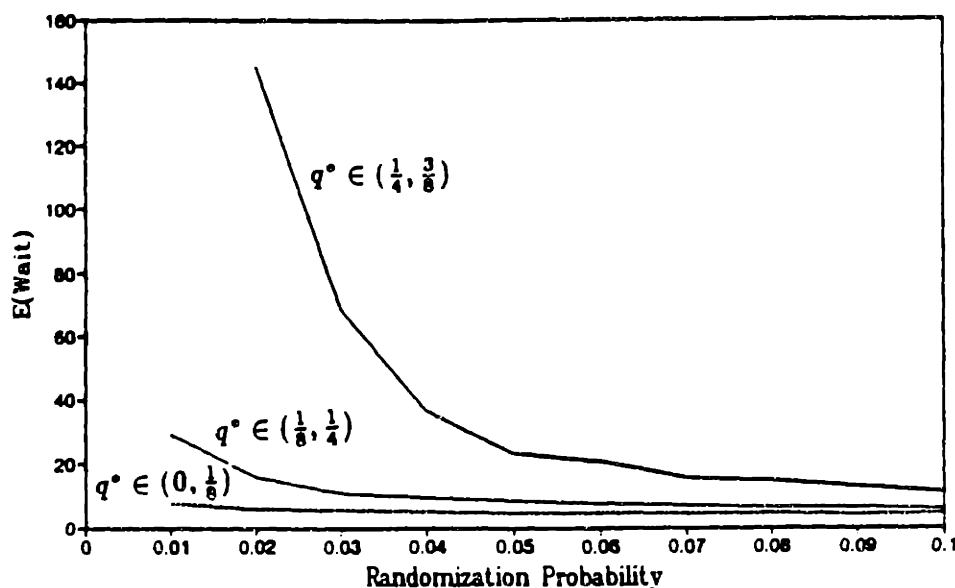
1.5.3 Heterogeneity

When modeling large populations it is unrealistic to assume that all players are identical. In the example of economists choosing computer packages clearly some will like each package more than others depending on their knowledge of the package, and on how well its features are suited to their work and their personal tastes. Obviously, if the players' preferences are sufficiently strong, players will completely ignore network externalities and no evolution will occur. What I show here, however, is that the addition of a small degree of heterogeneity actually speeds up convergence toward the risk-dominant equilibrium.

Before discussing heterogeneity, let us quickly look at the effect of simply increasing the payoff to (A, A) relative to that of (B, B) . Recall that in the model with eight neighbor matching on a circle and with our standard payoffs of $a = 2, b = 1, c = 0, d = 0, q^* = 1/3$ so that each player would like to play A if three of his eight

neighbors play A . The behavior of the model will be identical for any other payoffs with $q^* \in (1/4, 3/8)$. Suppose now that we either increase the payoff to (A, A) or decrease the payoff to (B, B) so that $q^* < 1/4$. Each player will now prefer to play A if only two of his neighbors play A . With no noise, a cluster of only two adjacent players playing A will be sufficient to guarantee that all players will eventually switch to A . The result is that convergence is much faster. Figure 1-3 graphs the estimated expected waiting time until 75% of the population has switched from B to A for three different ranges of q^* for various randomization probabilities.

Figure 1-3: Estimated Waiting Times for Various Payoffs



To model a limited degree of heterogeneity within the population, we might suppose that payoffs vary within the population. Perhaps the simplest assumption to make is that before the game begins each player is randomly endowed with preferences over the two equilibria. Rather than assuming that $u_i(A, A) = 2$ and $u_i(B, B) = 1$ for all i we can assume instead that for each player the two payoffs are determined by independent draws from separate distributions on $(0, \infty)$, the first with mean 2 and the second with mean 1. Player i 's behavior will then be determined by the fraction $q_i^* = u_i(A, A)/(u_i(A, A) + u_i(B, B))$ of his neighbors who must play A in order to make A his best response. In such a population, there will likely be a few players who will play A if only a few of their neighbors have done so. In the vicinity of these

players, smaller stable clusters of players playing A are possible and hence it will be far easier for stable clusters of players playing A to form. Once these clusters begin to expand, the fact that a few players prefer (B, B) will do little to slow their spread. The result is that we will see more rapid convergence.

Table 1.7 examines the effect of a small degree of heterogeneity on the expected waiting times. In each case, a population of 100 players is assumed. The first line gives the waiting time for any population in which all players have $q_i^* \in (1/4, 3/8)$. This includes our familiar assumption of a homogeneous population with $q^* = 1/3$. Rather than adding noise explicitly to the payoffs, I simply assume that somehow players have payoffs which result in the specified distribution of values of q_i^* .⁹ The second line records the behavior of a population where twenty players have $q_i^* \in (1/8, 1/4)$, sixty have $q_i^* \in (1/4, 3/8)$, and twenty have $q_i^* \in (3/8, 1/2)$. The third line adds more heterogeneity with five players having $q_i^* \in (0, 1/8)$ and five having $q_i^* \in (1/2, 5/8)$, so that there are players who in fact prefer to play B when both strategies are equally common among their neighbors. When ϵ is large so that evolution is rapid for a homogeneous population, heterogeneity has only a limited effect. What is more surprising, though, is that for ϵ small, the heterogeneity dramatically increases the rate at which play converges to the risk dominant A . This suggests that heterogeneity may play an important role in allowing rapid convergence for an even wider range of matching rules and parameter values than has already been identified.

⁹Virtually identical results can be obtained by assuming for example that the distribution of $u_i(A, A)$ and $u_i(B, B)$ is that of independent draws from lognormal distributions with $u_i(A, A) \stackrel{D}{\sim} 2.2u_i(B, B)$. (When $u_i(A, A) = 2.2$ and $u_i(B, B) = 1$, $q_i^* = 5/16$ is in the center of the interval $(1/4, 3/8)$.)

Table 1.7: Effect of Limited Heterogeneity*

Distribution of Payoffs	Randomization Probability		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
$q_i^* \in (\frac{1}{4}, \frac{3}{8})$	98.2 (7.5)	25.4 (0.8)	10.7 (0.2)
20 $q_i^* \in (\frac{1}{8}, \frac{1}{4})$ 60 $q_i^* \in (\frac{1}{4}, \frac{3}{8})$ 20 $q_i^* \in (\frac{3}{8}, \frac{1}{2})$	40.4 (2.0)	17.7 (0.5)	9.4 (0.2)
5 $q_i^* \in (0, \frac{1}{8})$ 20 $q_i^* \in (\frac{1}{8}, \frac{1}{4})$ 50 $q_i^* \in (\frac{1}{4}, \frac{3}{8})$ 20 $q_i^* \in (\frac{3}{8}, \frac{1}{2})$ 5 $q_i^* \in (\frac{1}{2}, \frac{5}{8})$	21.9 (0.6)	13.1 (0.3)	8.0 (0.2)

* Standard Errors in Parentheses

1.6 Conclusion

In this paper I have discussed a class of coordination games in order to examine the implications of a learning process among a large population of boundedly rational players. Kandori, Mailath and Rob (1991) introduced such a model and showed how the players' myopic adjustments create evolutionary forces which may select among the equilibria.

The analysis presented here yields two main conclusions. First, understanding the implications of dynamic models of learning requires that the rates of convergence be considered as well as the asymptotic distribution. Second, the nature of the matching rule which describes the interactions among the players can greatly affect the behavior we will observe. When each individual is equally likely to be matched with a great many opponents, play will exhibit great inertia. The play initially determined by arbitrary historical factors will persist for a long period of time regardless of whether a Pareto-superior or risk-dominant alternative is available. On the other hand, in communities in which players are only likely to be matched with a few close friends or colleagues, we have seen that the learning process not only leads to a shift in play toward the risk-dominant equilibrium, but also that the shift is rapid enough that it may very well be seen early in the game. This conclusion appears to be fairly robust to many aspects of the specification of the model.

References

- Blume, L., "The Statistical Mechanics of Strategic Interaction," Mimeo, 1991.
- Defoe, D., *A Tour Through the Whole Island of Great Britain*. London: Penguin, 1986. Original edition published London, 1724–6.
- Freidlin, M., and Wentzel, A., *Random Perturbations of Dynamical Systems*. New York: Springer Verlag, 1984.
- Harsanyi, J., and Selten, R., *A General Theory of Equilibrium in Games*. Cambridge: MIT Press, 1988.
- Kandori, M., Mailath, G., and Rob, R., "Learning, Mutation, and Long Run Equilibria in Games," Mimeo, 1991.
- Karlin, S. and Taylor, H., *A First Course in Stochastic Processes*. San Diego: Academic Press, 1975.
- Verlinden, C., "Markets and Fairs," *Cambridge Economic History of Europe*, Vol. III, 1963.
- Walford, C., *Fairs: Past and Present*. London, 1883. (Reprinted by A. M. Kelly: New York, 1968.)

Appendix

Proof of Theorem 2

I first consider $P^u(\epsilon)$. Note that $P_{ij}^u(\epsilon) = P_{i'j'}^u(\epsilon)$ if states j and j' have the same successor in the model with no noise. Hence, we may write

$$P^u(\epsilon) = R(\epsilon)Q$$

where

$$Q = \begin{bmatrix} 1 & 1 & \dots & \dots & 1 & 0 & 0 & \dots & \dots & 0 \\ 0 & 0 & \dots & \dots & 0 & 1 & 0 & \dots & \dots & 0 \\ 0 & 0 & \dots & \dots & 0 & 0 & 1 & \dots & \dots & 1 \end{bmatrix}$$

groups the states into three classes, the precursors of 0, $[q^*(N-1)]$, and the precursors of N , and $R(\epsilon)$ is an $N+1 \times 3$ matrix which gives the probability of each state arising given the class of the previous state. For example,

$$R_{i1}(\epsilon) = P_{ij}^u(\epsilon) \quad \forall j \in \{0, 1, \dots, [q^*(N-1)] - 1\}.$$

From this decomposition, it is clear that $P^u(\epsilon)$ has exactly three non-zero eigenvalues.

Let $C(\epsilon) = QR(\epsilon)$. $C(\epsilon)$ is a 3×3 matrix which can be regarded as giving the transitions between the three classes of states. Note that $P^u(\epsilon)^n = R(\epsilon)C(\epsilon)^{n-1}Q$. Let v be an eigenvector of $P^u(\epsilon)$ with eigenvalue $\lambda > 0$. We have $Qv \neq 0$ and $C(\epsilon)Qv = QR(\epsilon)Qv = \lambda Qv$ so Qv is an eigenvector of $C(\epsilon)$ with eigenvalue λ . As $C(\epsilon)$ has rank 3, this gives a one to one correspondence between the non-zero eigenvalues of $P^u(\epsilon)$ and the non-zero eigenvalues of $C(\epsilon)$.

We can write

$$C(\epsilon) = \begin{bmatrix} 1 - (p_1 + p_2) & q_1 & r_1 \\ p_1 & 1 - (q_1 + q_2) & r_2 \\ p_2 & q_2 & 1 - (r_1 + r_2) \end{bmatrix} \quad (1.3)$$

where, for example, p_1 gives the probability of a transition from any precursor of 0

to the state $[q^*(N-1)]$,

$$p_1 = \binom{N}{[q^*(N-1)]} \epsilon^{[q^*(N-1)]} (1-\epsilon)^{N-[q^*(N-1)]}.$$

Look at the characteristic polynomial of $C(\epsilon)$ as a polynomial in $z \equiv 1 - \lambda$. We find that

$$\text{Det}(C(\epsilon) - \lambda I) = z(z^2 - a_1 z + a_0) \quad (1.4)$$

with

$$a_1 = p_1 + p_2 + q_1 + q_2 + r_1 + r_2$$

$$a_0 = (p_1 + p_2)(q_1 + q_2) + (q_1 + q_2)(r_1 + r_2) + (r_1 + r_2)(p_1 + p_2) - q_2 r_2 - p_1 q_1 - p_2 r_1.$$

The root $z = 0$ corresponds to the eigenvalue $\lambda = 1$. Let $z_1(\epsilon) \geq z_2(\epsilon)$ be the other two roots of this equation. As $\epsilon \rightarrow 0$,

$$z_1(\epsilon) + z_2(\epsilon) \rightarrow 1$$

because $q_2(\epsilon) \rightarrow 1$ when $N - [q^*(N-1)] > [q^*(N-1)]$ and all other terms in the expression for a_1 converge to zero. We also have

$$z_1(\epsilon)z_2(\epsilon) = O(\epsilon^{[q^*(N-1)]})$$

as $q_2(\epsilon)p_1(\epsilon) = O(\epsilon^{[q^*(N-1)]})$ and all other terms in the expression for a_0 are of strictly higher order. Hence, we must have $z_1(\epsilon) \rightarrow 1$ and $z_2(\epsilon) = O(\epsilon^{[q^*(N-1)]})$. Clearly, $1 - z_2(\epsilon)$ is the second largest eigenvalue in absolute value for sufficiently small ϵ so the desired result for the uniform model follows from our characterization $r^u(\epsilon) = 1 - z_2(\epsilon)$.

Now consider the model with two neighbor matching. In the argument that follows I shall use S to denote the set of states $\{A, B\}^N$ and $s \in S$ to denote a member of that set. As before, I use \vec{A} , \vec{B} , $\vec{A}\vec{B}$, and $\vec{B}\vec{A}$ as shorthand for the states (A, A, \dots, A) , (B, B, \dots, B) , etc. Throughout the proof I assume that N is even; a

subset of this proof handles the case of N odd which is far simpler because the cycle $\vec{AB} \rightarrow \vec{BA} \rightarrow \vec{AB}$ does not exist. For v a probability distribution on S , v_s will denote the probability assigned to state s . It will often be convenient to write v as a 2^N -tuple ordered as

$$(v_{\vec{A}}, v_{\vec{AB}}, v_{\vec{BA}}, v_{\vec{B}}, \dots)$$

so that, for example, $(0, 0, 0, 1, 0, \dots, 0)$ represents a distribution which assigns probability 1 to all players playing B . (In writing this vector, some ordering is understood for the $2^N - 4$ states I have not specially named. Each of these states will be assigned probability zero in the 2^N -tuples which appear below.) Note that $P^\ell(\epsilon)$ can be considered to act on all $2^N \times 1$ vectors v , not just on probability distributions.

Let $CP_{P^\ell(\epsilon)}(x)$ be the characteristic polynomial of $P^\ell(\epsilon)$. $CP_{P^\ell(\epsilon)}$ and $CP_{P^\ell(0)}$ are polynomials of the same degree whose coefficients converge as $\epsilon \rightarrow 0$ so the set of roots of $CP_{P^\ell(\epsilon)}$ converges to the set of roots of $CP_{P^\ell(0)}$ (with multiplicity). If we write the eigenvalues of $P^\ell(\epsilon)$ with multiplicity as $1 = \lambda_1(\epsilon) > |\lambda_2(\epsilon)| \geq |\lambda_3(\epsilon)| \geq \dots \geq |\lambda_{2^N}(\epsilon)|$ this implies

$$|\lambda_i(\epsilon)| \rightarrow |\lambda_i(0)|.$$

$P^\ell(0)$ has a three dimensional space of eigenvectors of eigenvalue 1 spanned by

$$(1, 0, 0, 0, 0, \dots, 0), (0, \frac{1}{2}, \frac{1}{2}, 0, 0, \dots, 0), \text{ and } (0, 0, 0, 1, 0, \dots, 0).$$

The first and third eigenvectors correspond to the steady states where all players coordinate on A and B respectively. The second eigenvector corresponds to a 50-50 probability of being at each state of the $\vec{AB} \rightarrow \vec{BA} \rightarrow \vec{AB}$ cycle. $P^\ell(0)$ also has an eigenvector of eigenvalue -1 ,

$$(0, 1, -1, 0, 0, \dots, 0).$$

From this we know that

$$|\lambda_2(\epsilon)| \rightarrow 1, \quad |\lambda_3(\epsilon)| \rightarrow 1, \quad |\lambda_4(\epsilon)| \rightarrow 1,$$

and all other eigenvectors have absolute value bounded away from 1.

I will now show that $(1 - |\lambda_2(\epsilon)|)/\epsilon$ is bounded away from 0 and ∞ as $\epsilon \rightarrow 0$. If not, we can find a sequence $\{\epsilon_i\}$ converging to zero for which $(1 - |\lambda_2(\epsilon_i)|)/\epsilon_i$ is not bounded. Choosing an appropriate subsequence we may assume both that $\lambda_2(\epsilon_i)$ converges and that the corresponding eigenvectors $v(\epsilon_i)$ converge to a nonzero limit v .

Note that

$$P^\ell(0)v = \lim_{i \rightarrow \infty} P^\ell(\epsilon_i)v(\epsilon_i) = \lim_{i \rightarrow \infty} \lambda_2(\epsilon_i)v(\epsilon_i) = \pm v$$

so that v is an eigenvector of $P^\ell(0)$ with eigenvalue ± 1 . After a further normalization we must have either $v = (0, 1, -1, 0, 0, \dots, 0)$ or $v = (a, c, c, d, 0, \dots, 0)$ with $a + 2c + d = 0$. ($P^\ell(\epsilon_i)$ is a transition matrix so the sum of the elements of $P^\ell(\epsilon_i)v(\epsilon_i)$ is equal to the sum of the elements of $v(\epsilon_i)$. $\lambda_2(\epsilon_i) < 1$ then implies that the sum is zero.)

I now derive the contradiction that $(1 - |\lambda_2(\epsilon_i)|)/\epsilon_i$ is in fact bounded by considering two cases for v .

Case 1. $v = (a, c, c, d, 0, \dots, 0)$ $d \neq 0$

Here, we have $\lambda_2(\epsilon_i) \rightarrow 1$. We can write the eigenvalue as

$$\begin{aligned} \lambda_2(\epsilon_i) &= \frac{\sum_{s \in S} P_{\vec{B}s}^\ell(\epsilon_i)v_s(\epsilon_i)}{v_{\vec{B}}(\epsilon_i)} \\ &= P_{\vec{B}\vec{B}}^\ell(\epsilon_i) + \frac{\sum_{s \neq \vec{B}} P_{\vec{B}s}^\ell(\epsilon_i)v_s(\epsilon_i)}{v_{\vec{B}}(\epsilon_i)}. \end{aligned}$$

Note that

$$P_{\vec{B}\vec{B}}^\ell(\epsilon_i) = (1 - N\epsilon_i + O(\epsilon_i^2))$$

For $s \neq \vec{B}$, the successor of s in the model with no noise has at least two players playing A , so a transition from s to \vec{B} requires at least two ϵ -probability events. We then have that $P_{\vec{B}s}^\ell(\epsilon_i) = O(\epsilon_i^k)$ for some $k \geq 2$, and $v_s(\epsilon_i)$ is bounded. Hence, $P_{\vec{B}s}^\ell(\epsilon_i)v_s(\epsilon_i) = o(\epsilon_i)$ and

$$\lambda_2(\epsilon_i) = 1 - N\epsilon_i + o(\epsilon_i)$$

as desired.

Case 2. $v = (0, 1, -1, 0, 0, \dots, 0)$ or $v = (-2, 1, 1, 0, 0, \dots, 0)$

The proof for this case is somewhat longer. To make things clearer, let me begin by simply stating a pair of intermediate results I will prove later. The first is that we may assume without loss of generality that

$$v_{AB}(\epsilon_i) = v_{BA}(\epsilon_i) = 1$$

in the case $\lambda_2(\epsilon_i) \rightarrow 1$ and

$$v_{AB}(\epsilon_i) = -v_{BA}(\epsilon_i) = 1$$

in the case $\lambda_2(\epsilon_i) \rightarrow -1$. The second is that for \hat{S} the set of immediate predecessors of \vec{AB} (other than \vec{BA}) in the model with no noise (e.g. (B, B, B, A, B, A, ..., B, A)) and $s \in \hat{S}$

$$\lambda_2(\epsilon_i)v_s(\epsilon_i) = \begin{cases} \epsilon_i + o(\epsilon_i) & \text{if } s \text{ has } \frac{N}{2} + 1 \text{ players playing } B \\ o(\epsilon_i) & \text{otherwise} \end{cases} \quad (1.5)$$

Given these results, write

$$\begin{aligned} \lambda_2(\epsilon_i) &= \frac{\sum_{s \in \hat{S}} P_{ABs}^\ell(\epsilon_i)v_s(\epsilon_i)}{v_{AB}(\epsilon_i)} \\ &= P_{ABBA}^\ell(\epsilon_i) \frac{v_{BA}(\epsilon_i)}{v_{AB}(\epsilon_i)} + \sum_{s \in \hat{S}} P_{ABs}^\ell(\epsilon_i) \frac{v_s(\epsilon_i)}{v_{AB}(\epsilon_i)} + \sum_{s \notin \hat{S}} P_{ABs}^\ell(\epsilon_i) \frac{v_s(\epsilon_i)}{v_{AB}(\epsilon_i)}. \end{aligned}$$

The first term in this sum is equal to $\pm(1 - N\epsilon_i + o(\epsilon_i))$. There are $\frac{N}{2}$ states s in \hat{S} with $v_s(\epsilon_i) = \pm\epsilon_i + o(\epsilon_i)$ (for $N \geq 4$) and $P_{ABs}^\ell(\epsilon_i) = 1 - O(\epsilon_i)$ for all such states so the second term is equal to $\pm\frac{N}{2}\epsilon_i + o(\epsilon_i)$. Finally, the third term is $o(\epsilon_i)$ as for $s \in \{\vec{A}, \vec{AB}\}$, $P_{ABs}^\ell(\epsilon_i)$ is $O(\epsilon_i^{N/2})$ and $O(\epsilon_i^N)$ respectively, while for all other s , $P_{ABs}^\ell(\epsilon_i)$ is $O(\epsilon_i^k)$ for some $k \geq 1$ and $v_s(\epsilon_i)$ is $o(1)$.

Summing the three terms we find that

$$\lambda_2(\epsilon_i) = \pm(1 - \frac{N}{2}\epsilon_i + o(\epsilon_i)),$$

Thus, as in the previous case, we have the contradiction that $(1 - |\lambda_2(\epsilon_i)|)/\epsilon_i$ is bounded. (In fact, the existence of an eigenvalue with $\lambda(\epsilon) \rightarrow -1$ guarantees that $(1 - |\lambda_2(\epsilon)|)/\epsilon \rightarrow \frac{N}{2}$.)

To complete the proof, I now return to the two details I omitted earlier. First, let $v(\epsilon_i)$ be any eigenvector of $P^\ell(\epsilon_i)$ with eigenvalue $\lambda_2(\epsilon_i)$. Define $v(\epsilon_i)^r$ to be the vector with

$$v_s(\epsilon_i)^r \equiv v_{s^r}(\epsilon_i)$$

where s^r is the reverse of the state s , *i.e.*

$$(a_1, a_2, \dots, a_N)^r \equiv (a_N, a_{N-1}, \dots, a_1).$$

As $P^\ell(\epsilon_i)$ is symmetric, $v(\epsilon_i)^r$ is also an eigenvector with eigenvalue $\lambda_2(\epsilon_i)$ as are $v(\epsilon_i) + v(\epsilon_i)^r$ and $v(\epsilon_i) - v(\epsilon_i)^r$. Replacing $v(\epsilon_i)$ by $v(\epsilon_i) + v(\epsilon_i)^r$ in the $\lambda_2(\epsilon_i) \rightarrow 1$ case, and by $v(\epsilon_i) - v(\epsilon_i)^r$ in the $\lambda_2 \rightarrow -1$ case allows us to assume $v_{AB}(\epsilon_i) = v_{BA}(\epsilon_i)$ and $v_{AB}(\epsilon_i) = -v_{BA}(\epsilon_i)$ respectively.

For the second result, let R denote the set of all states in the basin of attraction of $\vec{AB} \rightarrow \vec{BA}$ cycle in the model with no noise, and let $T_{AB} \subset R$ be those states for which \vec{AB} is reached before \vec{BA} . Note that T_{AB} is naturally viewed as an \vec{AB} -tree with an arrow leading from each element to its immediate successor in the no noise model. For any $s \in T_{AB}$ we may define the height $h(s)$ to be the length of the path from s to \vec{AB} in this tree.

By induction, we first show that $v_s(\epsilon_i) = o(\epsilon_i)$ for $h(s) \geq 2$. Consider first all elements s with $h(s) = \max_{s' \in T_{AB}} h(s') = \frac{N}{2} - 1$ (if $\frac{N}{2} - 1 \geq 2$). We can write

$$\lambda_2(\epsilon_i)v_s(\epsilon_i) = \sum_{s' \in S} P_{ss'}^\ell(\epsilon_i)v_{s'}(\epsilon_i) \tag{1.6}$$

For $s' \notin \{\vec{A}, \vec{AB}, \vec{BA}\}$, $v_{s'}(\epsilon_i) = o(1)$ as $v(\epsilon_i) \rightarrow v$ and $P_{s,s'}^\ell(\epsilon_i) = O(\epsilon_i^k)$ for some $k \geq 1$ as s has no predecessors in the model with no noise. For $s' \in \{\vec{A}, \vec{AB}, \vec{BA}\}$, at least $\frac{N}{2} - 1$ ϵ -probability events are required for a transition from s' to s so $P_{s,s'}^\ell(\epsilon_i)v_{s'}(\epsilon_i)$ is $o(\epsilon_i)$ if $h(s) \geq 2$. Once the result has been established for all s with $h(s) = k > 2$, the result for $h(s) = k - 1$ again follows from the relation (1.6) with only the added calculation that for s' an immediate predecessor of s , $P_{s,s'}^\ell(\epsilon_i)v_{s'}(\epsilon_i)$ is again $o(\epsilon_i)$ as $v_{s'}(\epsilon_i)$ is $o(\epsilon_i)$ by the inductive hypothesis.

To get the desired result (1.5) take any $s \in T_{AB}$ with $h(s) = 1$ and use (1.6). For $s' = \vec{AB}$, $v_{AB}(\epsilon_i) = 1$ and

$$P_{s,AB}^\ell(\epsilon_i) = \begin{cases} \epsilon_i + o(\epsilon_i) & \text{if } s \text{ has } \frac{N}{2} + 1 \text{ players playing } B, \text{ and} \\ o(\epsilon_i) & \text{otherwise.} \end{cases}$$

For any $s' \in T_{AB}$ with $h(s') = 2$, $v_{s'}(\epsilon_i) = o(\epsilon_i)$. For $s' \in \{\vec{A}, \vec{BA}\}$, $P_{s,s'}^\ell(\epsilon_i)$ is at least $O(\epsilon_i^{\frac{N}{2}+1})$ and $O(\epsilon_i^{3N/4})$ respectively. Finally, for any other s' , $v_{s'}(\epsilon_i) = o(1)$ and $P_{s,s'}^\ell(\epsilon_i) = O(\epsilon_i^k)$ for some $k \geq 1$. Adding all the terms gives (1.5) as desired.

QED.

Chapter 2

Cooperation in the Prisoner's Dilemma with Anonymous Random Matching

2.1 Introduction

Ever since the earliest work on the Folk Theorem, it has been well known that when two players face each other in a repeated prisoner's dilemma the "cooperative" outcome can be sustained as a perfect equilibrium (Friedman, 1971; Aumann and Shapley, 1976). A variety of extensions are possible. Given additional assumptions, the Folk Theorem has been shown to apply to N player games, finite horizon games of incomplete information, and games with imperfect observations (Fudenberg and Maskin, 1986; Fudenberg, Levine, and Maskin, 1991). Recently, Kandori (1989) and Okuno and Postlewaite (1990) have discussed extensions of the Folk Theorem in a different direction, loosening informational requirements which may be unreasonable for matching games played by large populations of players. In this paper I continue in this direction, looking not at the Folk Theorem for general games, but simply at the prisoner's dilemma in a random matching setting under the most extreme informational assumptions—that players not only do not observe the outcomes of games in which they are not involved, but also are completely anonymous in that they cannot

recognize or communicate the identities of any of their past opponents. The fairly robust conclusion that players can still cooperate has implications for both applied and game theoretic uses of random matching models.

Recent papers by Greif (1989) and Milgrom et. al. (1990) have used random matching models to discuss characteristics of Medieval trade. Greif discusses the Maghribi traders, a group of North African Jews who conducted trade in many Mediterranean countries in the 11th century. Milgrom et. al. discuss trade in cities and fairs in Medieval Europe. In each case, the underlying model is one of a large number of traders who in each period are randomly paired with a trading partner. Each pair is presumed to play a game like the prisoner's dilemma with each party having both the opportunity and a private incentive to cheat the other by under-reporting sales on consignment, renegeing on promises to make future payments or deliveries, supplying goods of inferior quality, etc.

Standard Folk Theorem results imply that we can construct an equilibrium under full information where all traders will have an incentive to cooperate. The standard equilibrium involves a punishment phase directed at any player who cheats followed by a subsequent reward for players who carried out the punishment. In the random matching environment, this equilibrium makes considerable informational demands on the players. Not only must players be able to recognize a cheater whenever they meet him again, but somehow his identity must also be communicated to all the other players as well. Greif argues that the closeness of the Maghribi community allowed them to maintain cooperation. He cites evidence that many traders maintained ties to traders in other cities. Via this network of relationships they would quickly learn the identity of any cheaters, allowing the offending parties to be punished. Milgrom et. al. argue that such closeness no longer existed with the development of larger towns and trade fairs where keeping up with the reputation of every individual would be unreasonably demanding. They argue that this informational problem was resolved by the development of the Law Merchant, a private legal code whereby disputes could be tried before a judge who often lacked the power of enforcement. Nonetheless, this system could sustain cooperation if all traders consulted the judges' lists of cheaters

and refused any dealings with them. Maintaining this system is clearly costly, but Milgrom et. al. claim that the system seems well designed to keep costs as small as possible.

Random matching models have also been viewed as a possible method of reducing the multiplicity of equilibria in repeated games. Rosenthal (1979) discusses "rational Markovian hypotheses" in which all players react to steady state conjectures based only on their current opponent's play in the previous period, not on any further history. In the case of the prisoner's dilemma, both players cheating in every period is the only such equilibrium (except in one special case). Similarly, Milgrom et. al. (1990) note that with an infinite population and an extreme matching rule where no player can affect his future opponents' play in any way, cheating is the only Nash equilibrium outcome. While this is clearly an extreme case, it might be thought that in large finite populations individuals can have only a small impact on the future, and hence that there should be little difference.

In this paper, I discuss a random matching model similar to all those discussed above. There is a large finite population of players who are randomly paired to play the prisoner's dilemma in each period. The important informational assumption is that of anonymous interactions. In each period players only observe the outcome of the game in which they participate. They do not observe the identity of their opponent, nor do they observe the outcome of any of the games played by other pairs of players. The one informational convenience I do allow myself through most of the paper is to assume that some publicly observable random variable allows the players to coordinate their actions. For example, all traders at a fair might be able to observe the weather or hear official announcements. I later discuss how many of the results of the paper can be obtained without public randomizations.

Because of the assumption of completely anonymous interactions, the standard Folk Theorem does not apply. Nonetheless, Kandori (1989) shows that sometimes it may be possible to sustain cooperation. He defines "contagious" punishments where when one player cheats in period t , his period t opponent cheats from period $t+1$ on, infecting another player who cheats from period $t+2$ on, etc. For a fixed population

size, he shows that we can define payoffs for the prisoner's dilemma which allow cooperation in a perfect equilibrium. However, when the population is large the argument applies only to games with extreme payoffs.

In this paper, I adapt Kandori's arguments to study two main problems. First, for general payoffs in the prisoner's dilemma, does a cooperative equilibrium exist? I find that the answer is yes for sufficiently patient players. The basic idea of the equilibrium is to use contagious punishments which lead to a breakdown of cooperation after a single deviation. The public randomizations allow the severity of the punishments to be chosen so that the players fear a breakdown enough that they will not deviate first and destroy cooperation, but do not fear the breakdown so much that they are unwilling to contribute to its spread once it has begun. At several points I emphasize that this cooperation does not require unduly patient players. This result has implications for each of the applications of matching games mentioned above. In the models of trade, it suggests that while the institutions discussed may have facilitated cooperation in practice, they are by no means necessary or the least costly option in the abstract model. With regard to matching games as a device for reducing the equilibrium set, the results suggest that for the prisoner's dilemma it is the particular assumptions about the matching rule or the players' reactions which drive the results.

The second problem is a study of the stability and efficiency of the equilibrium in a world with noise. Kandori observes that in the equilibrium he constructs a single deviation causes a permanent end to cooperation and comments that this fragility may make the equilibrium inappropriate as a model for trade. His observation reflects two quite distinct concerns. The first is a modelling issue I shall refer to as "stability". If we intend the equilibrium to model cooperation in actual social settings and believe that in the real world punishments never last infinitely long we would like to construct an equilibrium with this property. Given public randomizations, this is not difficult. The second is a desire for a model which retains its efficiency in a world with noise. If we introduce noise by assuming that players either tremble and accidentally play the wrong strategy or misinterpret the actions of others, the equilibrium Kandori gives will be inefficient. Because cooperation eventually breaks down, the expected

payoff to very patient players will be near the non-cooperative level. In the standard repeated prisoner's dilemma with noise, the results of Fudenberg, Levine and Maskin (1991) imply that this inefficiency can be avoided. In the random matching model I am able to show that for sufficiently small probabilities of mistakes being made there is a perfect equilibrium in which players need not change their strategies in response to the presence of mistakes, and in which the inefficiency is small. We can conclude that the cooperative equilibrium with anonymous matching need not be as fragile as it has been portrayed.

The paper is organized as follows. Section 2 describes the model more precisely and exhibits a perfect equilibrium which sustains cooperation. Section 3 discusses the problem of stability and also shows that after introducing noise into the model we can still construct an equilibrium whose payoff approaches the efficient level as the amount of noise tends to zero, even for very patient players. Section 4 discusses the extension of the results to a model without public randomizations. While I believe that the public randomizations are appropriate for many applications, it is in the spirit of this paper to make due with as little information as possible. I show that we can in fact construct a cooperative equilibrium without public randomizations. Interestingly, this equilibrium remains nearly efficient in a model with noise even though it is no longer stable.

2.2 The Random Matching Model

For the remainder of this paper, I analyze the model described below. The game has M players indexed by $i \in \{1, 2, \dots, M\}$ where $M \geq 4$ is an even number. In each time period $t \in \{1, 2, 3, \dots\}$, the players are randomly matched into pairs with player i facing player $o_i(t)$. It is assumed that the pairings are independent over time and uniform so that

$$\text{Prob}\{o_i(t) = j | h_{t-1}\} = \frac{1}{M-1} \quad \forall j \neq i$$

for all possible histories h_{t-1} . At time t , each pair of players plays the prisoner's dilemma as shown below. The payoff g is taken to be positive with ℓ non-negative so

that each player has D as a dominant strategy in the stage game. All players have discount factor $\delta \in (0, 1)$ and their payoffs are the discounted sum of the payoffs in each stage game. At the end of period t , each player observes only the outcome of the prisoner's dilemma he and his opponent played. He does not observe the identity $\sigma_i(t)$ of his opponent and does not observe the outcome of any of the games played by other pairs of players.

	C	D
C	1, 1	$-\ell, 1+g$
D	$1+g, -\ell$	0, 0

In addition, I assume in this section and in the one which follows that before players choose their actions in period t , they observe a public random variable q_t which is drawn independently from a uniform distribution on $[0, 1]$. In some situations, it seems reasonable to assume that such a randomization is available. For example, all traders at a market may have access to the same newspaper or hear the same government announcements. In any case, the use of public randomizations simplifies the exposition below. I later discuss how many of the same results can be obtained without public randomizations.

The first thing to note about this model is that we can not implement the types of strategies usually used to prove the Folk Theorem. For example, when a player is the first to deviate, there is no way of identifying him, so it will be impossible to punish one player more severely and reward others for carrying out the punishment. Also, there is no obvious way to convey any information about the precise time of the deviation so that players could coordinate on something like T -period punishments.

Kandori (1989) shows that contagious punishments can be used to sustain collusion in some circumstances. Specifically, he shows that for any population size M , we

can choose the payoff ℓ so that cooperation is a perfect equilibrium for sufficiently patient players. The choice of ℓ is used to give players an incentive to carry out the punishment which follows a deviation. Unfortunately, the value of ℓ Kandori uses grows without bound as M increases and may be unreasonable for moderate values of M .

The main result of this section is that cooperation is indeed a perfect equilibrium of the random matching game for any payoffs g and ℓ . The equilibrium is supported by strategies like Kandori's which rely on contagious punishments. All subsequent results will rely on similar strategies. The following proposition gives the basic result.

Proposition 1 *Consider the random matching model with public randomizations described above where $M \geq 4$ players play the prisoner's dilemma with $g > 0$, $\ell \geq 0$. Then, $\exists \underline{\delta} < 1$ such that $\forall \delta \in [\underline{\delta}, 1)$ there is a perfect equilibrium $s^*(\delta)$ of the repeated game in which all players play C in every period.*

Before giving a formal proof, I first discuss the strategies which $s^*(\delta)$ which will support the equilibrium. The strategies described below employ a contagious process by which a deviation in period t will usually lead to two players playing D in period $t+1$, then four players playing D in period $t+2$, etc. The result is a breakdown of social cooperation which punishes all players after one deviates. Given a function $q(\delta)$ to be defined below, the strategies are as follows.

In period 1, all players begin play according to phase I.

Phase I. Play C in period t .

If (C, C) is the outcome for matched players i and j , both play according to phase I in period $t+1$.

If (C, D), (D, C), or (D, D) results in the game between players i and j , then at time $t+1$ both play according to phase II if $q_{t+1} \leq q(\delta)$ and according to phase I if $q_{t+1} > q(\delta)$.

Phase II. Play D in period t .

In period $t+1$ play according to phase I if $q_{t+1} > q(\delta)$ and according to phase II if $q_{t+1} \leq q(\delta)$.

The public randomizations are used to adjust the severity of the punishment phase so that it lasts $1/(1 - q(\delta))$ periods on average. The basic idea of the proof is this. In a perfect equilibrium the continuation payoffs of the players must satisfy two constraints. First, players must not want to deviate and play D in phase I. When punishments are of infinite duration (*i.e.* for $q(\delta) = 1$), sufficiently patient players will not want to cause a breakdown of cooperation in phase I so this constraint is satisfied. Second, we must recognize that in phase II players might deviate and play C in hopes of slowing the spread of the contagious punishment. When punishments never occur (*i.e.* for $q(\delta) = 0$) there is no possible gain to deviating in phase II so this constraint will be satisfied.

To prove the theorem, I show that there exists at least one value $q(\delta)$ which is both large enough to prevent deviations in phase I and small enough to prevent deviations in phase II. The intuitive reason why this can be done is simple. In either phase I or phase II player i gets the same short term gain of g from playing D when his opponent cooperates. However, starting a punishment by playing D in phase I causes a greater loss in continuation payoff than does spreading a punishment by playing D in phase II. Once play is in phase II, cooperation is breaking down anyway so one extra deviation has limited impact. Choosing an appropriate punishment severity, the loss from starting a punishment deters playing D in phase I, but the loss from spreading a punishment does not deter playing D in phase II.

To formalize this argument let k be the number of players who are playing according to phase II at the start of period t . Define $f(k, \delta, q)$ be player i 's (per period) continuation payoff from period t on when all players are playing the strategies above, and player i and $k - 1$ others are playing according to phase II. If player i deviates and plays D in phase I in period t , he gains g in period t but will have a lower continuation payoff from period $t+1$ on. To show that no deviation is profitable in phase I we must show that

$$(1) \quad (1 - \delta)g \leq \delta q(\delta)(1 - f(2, \delta, q(\delta))).$$

We can also derive a similar sufficient condition for there to be no profitable deviation in phase II. If player i deviates and plays C in phase II at time t we have one of two possibilities. First, he could be matched with someone else who is playing according to phase II. In this case, the result in period t is (C, D) instead of (D, D), and continuation payoffs are unaffected. Clearly, player i is not better off because $\ell \geq 0$. Second, player i might be matched with someone who is playing according to phase I. The period t outcome is then (C, C) instead of (D, C) so player i loses g in period t . In the continuation game, however, one fewer player will be playing according to phase II. The deviation is not profitable if

$$(1 - \delta)g \geq \delta q(\delta) E_j [(f(j, \delta, q(\delta)) - f(j + 1, \delta, q(\delta)))]$$

where the expectation reflects the fact that player i does not know how many other players will play according to phase II at time $t+1$. To show that this relation holds for all possible beliefs of player i , a sufficient condition is to show that it holds pointwise, *i.e.*

$$(2) \quad (1 - \delta)g \geq \delta q(\delta) (f(j, \delta, q(\delta)) - f(j + 1, \delta, q(\delta))) \quad \forall j \geq 3.$$

When (1) and (2) hold, we have a perfect equilibrium. In establishing these relations, both the result and the intermediate calculations of the following lemma will prove useful.

Lemma 1 $f(k, \delta, q)$ is convex in k for $k \geq 1$, *i.e.*

$$f(k, \delta, q) - f(k + 1, \delta, q) \geq f(k + s, \delta, q) - f(k + s + 1, \delta, q)$$

for all $s \geq 1$.

The lemma simply states that the loss in continuation payoff from having one extra player infected declines as the number of infected players grows. This is to be expected as when many players are infected, the one extra player not infected

in period t is likely to become infected in period $t+1$ anyway and thus never have a chance to affect player i 's payoff. The proof is straightforward once I introduce enough notation.

Proof of Lemma 1

Note that

$$f(k, \delta, q) = E_{\omega} g(k, \delta, q, \omega),$$

where ω is the random variable whose realization is a pairing of all the players in each period, and the function g gives player 1's continuation payoff for a given matching when he and players $2, \dots, k$ are playing according to phase II. For expositional convenience I define $h(k, \delta, q, \omega)$ to be player i 's continuation payoff when he and players $2, \dots, k$ and player M are playing according to phase II. Clearly

$$E_{\omega} g(k+1, \delta, q, \omega) = E_{\omega} h(k, \delta, q, \omega).$$

I show that

$$E_{\omega} [g(k, \delta, q, \omega) - h(k, \delta, q, \omega)] \geq E_{\omega} [g(k+s, \delta, q, \omega) - h(k+s, \delta, q, \omega)]$$

by showing that the inequality holds for every realization of ω .

Define the set $C(t, k, \omega)$ by

$$\begin{aligned} C(0, k, \omega) &= \{k+1, k+2, \dots, M\} \\ C(t+1, k, \omega) &= \{i \in C(t, k, \omega) \mid o_i(t, \omega) \in C(t, k, \omega)\}. \end{aligned}$$

$C(t, k, \omega)$ will be the set of players still playing according to phase I in period t when $q_s \leq q$ for all $s \leq t$ and players $1, 2, \dots, k$ begin in phase II in period 0.

Define the set $D(t, \omega)$ by

$$\begin{aligned} D(0, \omega) &= \{M\} \\ D(t+1, \omega) &= D(t, \omega) \cup \{i \mid o_i(t, \omega) \in D(t, \omega)\}. \end{aligned}$$

$D(t, \omega)$ gives the set of all players who will be playing according to Phase II in period t if player M begins in phase II in period 0.

Note that the payoff to player 1 in period t differs between the situations of $g(k, \delta, q, \omega)$ and $h(k, \delta, q, \omega)$ only if $q_s \leq q$ for all $s \leq t$ and only if his opponent $\sigma_1(t, \omega)$ plays C when players $1, 2, \dots, k$ start in phase II but plays D when players $1, 2, \dots, k$, and player M start in phase II. Thus,

$$(3) \quad g(k, \delta, q, \omega) - h(k, \delta, q, \omega) = \sum_{t=0}^{\infty} (1 - \delta) q^t \delta^t (1 + g) I\{\sigma_1(t, \omega) \in C(t, k, \omega) \cap D(t, \omega)\}.$$

The definition of C clearly implies that

$$C(t, k + s, \omega) \subset C(t, k, \omega)$$

so

$$C(t, k + s, \omega) \cap D(t, \omega) \subset C(t, k, \omega) \cap D(t, \omega)$$

and the expansion (3) gives the desired result. **QED**

We are now in a position to give

Proof of Proposition 1

Let $s^*(\delta)$ be the strategy profile given above. It suffices to demonstrate the existence of a $\underline{\delta} < 1$ such that (1) and (2) hold for all $\delta \in [\underline{\delta}, 1)$. To establish the relation (1), we will simply define $\underline{\delta}$ and $q(\delta)$ on $[\underline{\delta}, 1)$ so that (1) holds with equality. To see that this is possible, note that for $q(\delta) = 1$, punishments are infinite so

$$f(2, \delta, 1) = \sum_{t=0}^{\infty} \delta^t a_t$$

where a_t is the expected payoff in the t^{th} period after phase II play begins. With probability 1 all players will eventually be infected and start playing D so that $a_t \rightarrow 0$. We then have

$$\lim_{\delta \rightarrow 1} \frac{\delta}{1 - \delta} (1 - f(2, \delta, 1)) = \infty$$

$$\lim_{\delta \rightarrow 0} \frac{\delta}{1 - \delta} (1 - f(2, \delta, 1)) = 0.$$

By continuity we can choose $\underline{\delta} \in (0, 1)$ so that

$$\frac{\underline{\delta}}{1 - \underline{\delta}} (1 - f(2, \underline{\delta}, 1)) = g.$$

Note that when (1) holds with equality, a player in phase I is exactly indifferent between playing C and D. The payoff to a player who plays D in period 1 is $f(1, \delta, q(\delta))$. Thus, (1) holds with equality only if

$$(4) \quad \frac{\delta q(\delta)}{1 - \delta} (f(1, \delta, q(\delta)) - f(2, \delta, q(\delta))) = g.$$

The converse is also true. When (4) holds, a player in phase I is exactly indifferent between playing D in period t (and following the equilibrium strategies thereafter) and playing C in period t then deviating and playing D in period $t+1$. Applying the same indifference again, he is also indifferent between deviating in period t and playing C in periods t and $t+1$ and then deviating in period $t+2$. Repeating this process, he is indifferent between deviating in period t and cooperating in all future periods. This implies that (1) holds with equality.

From expansion (3) we have

$$(5) \quad \frac{\delta q}{1 - \delta} (f(k, \delta, q) - f(k + 1, \delta, q)) = \sum_{t=0}^{\infty} (\delta q)^{t+1} (1+g) \text{Prob}\{o_1(t, \omega) \in C(t, k, \omega) \cap D(t, \omega)\}.$$

As the right hand side depends only on the product δq , we simply define

$$q(\delta) \equiv \underline{\delta} / \delta.$$

Then, for any $\delta \in [\underline{\delta}, 1)$, $\delta q(\delta) = \underline{\delta}$ and

$$\frac{\delta q(\delta)}{1 - \delta} (f(1, \delta, q(\delta)) - f(2, \delta, q(\delta))) = g$$

as desired.

The result of the lemma now immediately implies (2) and hence completes the proof. **QED**

In a full information model, the "grim" strategies immediately punish a player who has cheated once. In contrast, the contagious punishment takes time to spread throughout the population so that a player may be able to cheat several opponents before he begins to suffer from the punishment phase he has brought on. This observation leads us to ask whether the equilibrium described in Proposition 1 requires undue patience on the part of the players.

Table 2.1 gives the minimum value of δ which can sustain cooperation for several population sizes M and for several values of the gain g to deviation. The discount factor for $M = 2$ is the discount factor necessary for the standard "grim" equilibrium in a two player game. While the contagious equilibrium requires more patient players as the population size grows, we certainly can not say that the increase in patience is so dramatic as to render the cooperative equilibrium impractical for large populations. If players discount the future at a rate of 5% per year, a discount factor of 0.996 is appropriate for monthly interactions. From the table, it is clear that in a model with frequent interactions like weekly or monthly trade at a market, contagious punishments will support cooperation for a broad range of payoffs and population sizes. Even if we wish to model annual interactions so that $\delta = 0.95$ is reasonable, cooperation is possible provided that the gain from cheating is not too large. While contagious punishments are less potent than the full information "grim" strategies, they are still sufficiently powerful to sustain cooperation with anonymous interactions for very reasonable discount factors.

To better appreciate the power of the contagious punishments in large populations, it is instructive to compare the discount factors of Table 2.1 to those necessary for another large population equilibrium. While we have focussed on completely anonymous matching, it should be noted that for many applications it may be reasonable to make the less stringent assumption that identities can be observed but not communicated,

perhaps because you recognize someone who cheated you but can not communicate an adequate verbal description of the person's appearance. In this model, we could sustain cooperation through personal retaliation with strategies where a deviation by player i in period t causes his period t opponent $o_i(t)$ to play D whenever they are matched in the future. Note that this equilibrium requires frequent *individual* interactions, and thus requires far more patient players than does the equilibrium with contagious punishments. If player i cheats in period t , he gains g in period t and loses 1 in each future period in which he is again matched with $o_i(t)$. This gives a cooperative equilibrium only if

$$\sum_{t=1}^{\infty} \delta^t \frac{1}{M-1} \geq g \quad \iff \quad \delta \geq \frac{g(M-1)}{1+g(M-1)}$$

For $g = 1$ and $M = 500$, for example, this requires $\delta = 0.998$, whereas $\delta = 0.92$ is sufficient for the equilibrium with contagious punishments.

2.3 Stability and Efficiency with Noise

The cooperative equilibrium described in section 2 exhibits the desirable property of global stability described by Kandori (1989). That is, after *any* finite history, the continuation payoffs of the players eventually return to the cooperative level (with probability 1). Obviously, this is a result of the introduction of public randomizations. The stability does suggest, though, that robustness in this sense is not a big problem for this model.

A more interesting question is whether we can still sustain a nearly efficient outcome in a model with noise. Suppose we really believed that the model of section 2 with its completely rational players and perfect observations were an accurate depiction of reality. Even if players follow the strategies of an equilibrium with infinite punishments, in equilibrium the punishment never begins, so we have no reason to care about the behavior of the continuation payoffs after a deviation. On the other hand, suppose that there is noise in the model, as players either act irrationally some

fraction of the time, or try to cooperate but make mistakes and play the wrong strategy or misinterpret their opponent's action. Again, I would argue that whether an equilibrium is stable is not the appropriate question to ask. If we have a globally stable equilibrium in which the continuation payoffs return to the cooperative level so slowly so that with noise the equilibrium has an expected payoff near zero, stability is not comforting. Suppose we have two different equilibria which have the same loss of efficiency after any deviation. Should we care if one equilibrium has all the inefficiency right away and then returns to cooperation while the other spreads out the same inefficiency over an infinite time period? The answer, I think, is that all that matters is the degree of efficiency the equilibrium attains in a model with the noise explicitly modeled.

In the two player repeated prisoner's dilemma complete efficiency can be attained in the limit $\delta \rightarrow 1$ (Fudenberg, Levine, and Maskin, 1991). I now introduce noise into the model of section 2 by assuming that all players are constrained to play D with probability at least $\epsilon > 0$ at every possible history. In the trade example, this could correspond to players trying to supply a high quality good but accidentally supplying one which proves defective. A similar result could be obtained if we assumed instead that there was only noise in observing opponents' actions. While the equilibrium of section 2 is not robust to this noise (because of the exact indifference during phase I play), the proposition below shows that for a slightly longer punishment length we do in fact have an equilibrium robust to this noise. While the existence of a fully efficient equilibrium is still an open question, the equilibrium described is approximately efficient in the sense that it approaches efficiency as $\epsilon \rightarrow 0$.

Proposition 2 *Under the assumptions of Proposition 1, there exists $\underline{\delta}' < 1$ and a set of strategy profiles $s^*(\delta)$ for $\delta \in [\underline{\delta}', 1)$ of the random matching game with the following three properties:*

1. *In the game with discount factor δ , $s^*(\delta)$ is a perfect equilibrium with all players playing C on the path in every period.*
2. *Define $s^*(\delta, \epsilon)$ to be the strategy which at each history assigns probability ϵ to D and probability $1 - \epsilon$ to the action given by $s^*(\delta)$. Then, there exists $\bar{\epsilon} > 0$ such that $\forall \epsilon < \bar{\epsilon}$*

$s^*(\delta, \epsilon)$ is a perfect equilibrium of a perturbed game where all players are required to play D with probability at least ϵ at each history.

3. For u_i defined to be player i 's expected per period payoff,

$$\lim_{\epsilon \rightarrow 0} \lim_{\delta \rightarrow 1} u_i(s^*(\delta, \epsilon)) = 1.$$

Outline of Proof

We will show that $s^*(\delta)$ can be taken to have the same form as the strategy profile in the proof of Proposition 1, but with a slightly larger probability $q'(\delta)$ of continuing in a punishment phase. The proof requires attention to some tedious details, so I only outline the proof here and leave the rest to the appendix.

To begin, I give a slight extension of Lemma 1, showing that the continuation payoff function f is strictly convex. The strict convexity allows us to choose a slightly larger $q'(\delta)$ so that the two inequalities which describe a player's loss from deviating in phase I or phase II of the model with no noise hold strictly. Formally, the appendix shows that we can choose $\eta > 0$ (independent of δ) for which

$$(6) \quad \frac{\delta q'(\delta)}{1 - \delta} (f(0, \delta, q'(\delta)) - f(2, \delta, q'(\delta))) > g + \eta$$

and

$$(7) \quad \frac{\delta q'(\delta)}{1 - \delta} (f(k, \delta, q'(\delta)) - f(k + 1, \delta, q'(\delta))) < g - \eta \quad \forall k \geq 2.$$

This immediately gives property 1.

To show that these strategies give an equilibrium for all sufficiently small ϵ requires two further steps. First, it must be shown that the left hand side of each equation is uniformly continuous in ϵ so that for small enough ϵ the inequalities above still hold but with η replaced by $\eta/2$. For $f(k, \delta, q, \epsilon)$ defined to be the continuation payoff of the strategies $s^*(\delta, \epsilon)$ the appendix demonstrates the existence of an $\bar{\epsilon}_1 > 0$ such that

for any $\epsilon < \bar{\epsilon}_1$

$$(8) \quad \frac{\delta q'(\delta)}{1 - \delta} (f(1, \delta, q'(\delta), \epsilon) - f(2, \delta, q'(\delta), \epsilon)) > g + \eta/2$$

and

$$(9) \quad \frac{\delta q'(\delta)}{1 - \delta} (f(k, \delta, q'(\delta), \epsilon) - f(k + 1, \delta, q'(\delta), \epsilon)) < g - \eta/2 \quad \forall k \geq 2.$$

Second, we have a new complication in that when a player is playing according to phase I, he can no longer believe with probability 1 that all other players are doing so. Again though, as $\epsilon \rightarrow 0$, this uncertainty also has an effect which vanishes so that the incentives to cooperate are maintained for sufficiently small ϵ . This completes the proof of 2.

Finally, the proof that we get efficiency in the limit is easy. The basic idea is that the punishment phases have a finite expected length bounded above by a constant independent of δ for δ close to 1. As $\epsilon \rightarrow 0$ a vanishing fraction of the periods is spent in a punishment phase, so the expected payoff tends to the efficient level. Again the details are in the appendix. **QED**

The results of Proposition 2 indicate that the perfect equilibrium I have described is far less fragile than it might appear at first. The same strategies yield an equilibrium for all sufficiently small amounts of noise, so players can cooperate even if they do not know the precise frequency with which other players make mistakes. Further, the strategies are truly supporting cooperation in the sense of having nearly efficient payoffs with noise.

In applying this conclusion to models of trade, however, some caution is called for. If we think that mistakes are common enough that say one player in the population makes a mistake in each period, the strategies given clearly will not support a nearly efficient equilibrium. The set up of Milgrom et. al. (1990) suggests one possible way of ameliorating the problem. Suppose we modify the stage game so that a player

who accidentally cheats has the opportunity to give back his excess payoff after a trial and avoid the start of a punishment phase. With this stage game, if we assume that mistakes are the result of independent trembles at each information set, a player would have to tremble twice in a row to start a breakdown. For a given probability of trembling, accidental punishments are far less common so we may have a far more efficient equilibrium. The Law Merchant could then serve a different but still very useful purpose in the equilibrium I have given by reducing the frequency with which breakdowns of cooperation occur.

The fact that each action in our equilibrium with contagious punishments is a strict best response also allows the further extension that follows. In a large population (like a group of medieval traders), we may want to allow for heterogeneity among the players. In particular, it is probably reasonable to assume that the players have different rates of time preference. In each of the first two propositions, the equilibrium strategy profile $s^*(\delta)$ is a function of the discount factor. For each discount factor δ , the equilibrium involves a different probability $q(\delta)$ of continuing within the punishment phase. Hence, the strategies are only appropriate for a population of players all of whom share a common discount factor. As long as all of the players are sufficiently patient, however, we can eliminate this restriction. The proposition below guarantees the existence of a perfect equilibrium strategy profile s^* which is not a function of δ . This profile will then sustain cooperation regardless of whether the population shares a common discount factor. For convenience, I shall discuss only a model without noise although the arguments clearly extend to the results of Proposition 2 as well. The proof is similar to that of Proposition 2, but is less involved.

Proposition 3 *Under the assumptions of Proposition 1, there exists a strategy profile s^* and a constant $\underline{\delta}'' < 1$ such that $\forall \delta \in [\underline{\delta}'', 1)$, s^* is a perfect equilibrium of the repeated matching game and all players play C in every period on the path of s^* .*

Proof

Once again, let s^* be a strategy profile like the one described in the proof of Proposition 1, but this time use the punishment probability $q'' \equiv \lim_{\delta \rightarrow 1} q'(\delta) = \underline{\delta}'$, a constant

independent of δ . An intuitive argument for the rest of the proof is that for δ close to 1, q'' is close to $q'(\delta)$ so the strategy profile s^* is very close to the $s^*(\delta)$ of Proposition 2. Because $s^*(\delta)$ is a strict equilibrium, this should suffice.

Completing the outline above involves some tedious calculations so I give a simpler constructive argument instead. Let $\underline{\delta}'' = \underline{\delta}/\underline{\delta}'$, where $\underline{\delta}$ is as defined in Proposition 1. For any $\delta \in [\underline{\delta}, 1)$, we have $\delta q'' < q'' = \underline{\delta}'$. Hence, from (5) and (7) we know that

$$\begin{aligned} \frac{\delta q''}{1 - \delta} (f(2, \delta, q'') - f(3, \delta, q'')) &< \frac{\underline{\delta}'}{1 - \underline{\delta}'} (f(2, \underline{\delta}', 1) - f(3, \underline{\delta}', 1)) \\ &< g - \eta. \end{aligned}$$

We also have $\delta q'' \geq \underline{\delta}$. Hence, (5) and the definition of $\underline{\delta}$ gives

$$\begin{aligned} \frac{\delta q''}{1 - \delta} (f(1, \delta, q'') - f(2, \delta, q'')) &\geq \frac{\underline{\delta}}{1 - \underline{\delta}} (f(1, \underline{\delta}, 1) - f(2, \underline{\delta}, 1)) \\ &= g. \end{aligned}$$

As in Proposition 1, these two conditions are sufficient to show that we have a perfect equilibrium. **QED**

A potentially disturbing aspect of the preceding proof is that because it involves another limit as $\delta \rightarrow 1$, the equilibrium with heterogeneous discount factors might require far more patient players than was previously necessary. From Table 2.1 we know that Propositions 1 and 2 do not require unreasonably patient players. Certainly, the equilibrium described in Proposition 3 will sometimes require more patient players. This is particularly true when the gain g from deviation is small so that it is hard to get players to carry out punishments. For example, for a population of 100 players, if we take g to be 0.01 the the equilibrium as constructed requires $\delta = 0.96$. Usually, though, we will think of g as being much larger. In the trade example, the payoff of 1 represents the profit or consumer surplus from an honest transaction. The gain of g from cheating might represent the additional cost savings from either failing to deliver the good or producing a good of inferior quality. These potential gains are

liable to be at least as large the profits from honest trade, so it is more appropriate to assume that g is near 1 than 0.01.

For such parameter values, equilibrium with heterogeneity does not require unduly patient players. In fact, it often requires that players be no more patient than in the model with homogeneous players. When the constraint that players be willing to carry out punishments is sufficiently far from binding, we can simply use infinite punishments for all $\delta \in [\underline{\delta}, 1)$ to get an equilibrium. Numerical calculations show this to be the case for each of the population sizes given in Table 2.1 for $g = 1$ or $g = 10$.

2.4 Cooperation without Public Randomizations

Throughout this paper, I have assumed that a public randomizing device is available. For many applications, including trade at a market, the assumption seems reasonable. Whenever all the players are present at the same physical location it seems likely that if the players looked hard enough they could find some random factor like the weather which everyone could observe and hence use to coordinate. Nonetheless, the focus of this paper is to describe how cooperation can be maintained with very little information available to the players. In this spirit then, I discuss what can be done without public randomizations.

In Fudenberg and Maskin's (1986) proof of the perfect Folk Theorem, public randomizations played a crucial role in allowing the adjustment of players continuation payoffs necessary for maintaining exact indifference. Fudenberg and Maskin (1991) show that public randomizations are, in fact, not necessary for this purpose. The crucial insight is that payoffs in the convex hull of the set of feasible payoffs can be obtained instead from a deterministic sequence of play.

In this paper, public randomizations are playing two quite distinct roles. First, they are used as a coordinating device so that all players can simultaneously return to cooperation at the end of a punishment phase. The simultaneity is important because all players only slightly prefer cooperating when all others are doing so. If the probability that everyone else returns to cooperation in period t is not very close

to one, no one will be willing to try returning to cooperative play. Coordination then allows the construction of a globally stable equilibrium. Whether global stability is possible without the public randomizations is unknown.

The second role of the public randomizations in this paper is to adjust the expected duration, and hence the severity of the punishments. This is the property which enabled us to construct strategies where punishments deter cheating, but are not so severe that individuals would be unwilling to carry them out. In the argument below, I show that for large enough discount factors it is possible to adjust the severity of the punishments in a completely different way – spreading out the punishments over time. This will allow us to establish the most important results of the paper even without the availability of public randomizations.

The ability to soften punishments by delaying them is at the heart of the following lemma. The lemma guarantees that any game which has a cooperative equilibrium for some interval of discount factors has a cooperative equilibrium for all discount factors near one as well. I hope that the very simple proof makes the lemma interesting in its own right.

Lemma 2 *Let $G(\delta)$ be any repeated game of complete information, and suppose that there is a non-empty interval (δ_0, δ_1) such that $G(\delta)$ has a perfect equilibrium $s^*(\delta)$ with outcome a for all $\delta \in (\delta_0, \delta_1)$. Then, there exists $\underline{\delta} < 1$ such that $\forall \delta \in [\underline{\delta}, 1)$ we can also define a strategy profile $s^{**}(\delta)$ which is a perfect equilibrium of $G(\delta)$ with outcome a .*

Proof

The key observation here is that for δ close enough to 1, we can simulate the situation of smaller discount factors by using slower responses.

Take $\underline{\delta} = \delta_0/\delta_1$. For any $\delta \in [\underline{\delta}, 1)$ there exists an integer $N(\delta)$ for which

$$\delta^{N(\delta)} \in (\delta_0, \delta_1).$$

When there is more than one such integer take $N(\delta)$ to be as large as possible. Now, have the players treat the game $G(\delta)$ as if it were $N(\delta)$ separate games, the first

taking place in periods

$$1, N(\delta) + 1, 2N(\delta) + 1, 3N(\delta) + 1, \dots,$$

the second in periods

$$2, N(\delta) + 2, 2N(\delta) + 2, 3N(\delta) + 2, \dots,$$

etc. Just as is the case in finding Markov equilibria, if for some set T all other players play strategies in period t which do not depend on the outcomes in all periods $t' \in T$, then the best response for player i can be taken to be independent of the outcomes in all periods $t' \in T$ as well. Hence, to show that we have an equilibrium $s^{**}(\delta)$ for $G(\delta)$ it suffices to show:

1. The strategies $s^{**}(\delta)$ give play in period $aN + b$ which does not depend on play in period $cN + d$ if $(b - d)$ is not a multiple of N .
2. Restricting consideration to each "component" game played in periods

$$b, N(\delta) + b, 2N(\delta) + b, 3N(\delta) + b, \dots,$$

the restriction of the strategy profile $s^{**}(\delta)$ gives a perfect equilibrium.

The obvious choice of $s^{**}(\delta)$ is to play the equilibrium $s^*(\delta^{N(\delta)})$ in each of the $N(\delta)$ component games described above. In our prisoner's dilemma example, this would mean that if player i or his opponent plays D in period $aN(\delta) + b$, player i plays D in periods

$$(a + 1)N(\delta) + b, (a + 2)N(\delta) + b, \dots,$$

but does not change his planned play in any other period. Within these component games, players have discount factor $\delta^{N(\delta)}$, so $s^*(\delta^{N(\delta)})$ satisfies the second condition. Clearly, we have a perfect equilibrium. **QED**

Note that when $q'(\delta) = 1$, the strategies described in the proof of Proposition 2 prescribe infinite punishments, and hence do not require public randomizations. In particular, $\underline{\delta}'$ was defined so that taking $q = 1$ gives a perfect equilibrium. In order to apply Lemma 2, we need only show that infinite punishments also yield a perfect equilibrium for a small interval of discount factors around $\underline{\delta}'$. This result is not hard. It is simply another application of the fact that each action is a strict best response. The resulting equilibrium of the game has a peculiar appearance with punishments being softened by being delayed into the future, spread among intervening periods of cooperation. In the trade example, this might mean that if a single deviation occurs on a Friday, eventually we will see all players cheating on every third Friday but cooperating on all other days. The punishments are of infinite duration so with noise, eventually all players will cheat in all periods. Despite this, the punishments are still no more severe than the punishments of the previous section. As players become more patient, the punishment periods become correspondingly further apart. The somewhat surprising result is that in the limit as the amount of noise vanishes, the equilibrium approaches efficiency. These results are summarized below.

Proposition 4 *The results of Proposition 2 still hold in a model where no public randomizations are available.*

Proof

In order to establish the first two results of Proposition 2, that there is a perfect equilibrium which remains an equilibrium for sufficiently small amounts ϵ of noise, it will suffice to show that for a fixed range of discount factors the standard strategies with $q = 1$ give a perfect equilibrium. Just as in Proposition 3, we apply continuity of the payoff functions to show that a strict equilibrium for one discount factor implies that nearby discount factors also give an equilibrium.

Recall that in the proof of Proposition 2, $\underline{\delta}'$ was defined so that the contagious strategies with parameter $q'(\underline{\delta}') = 1$ give a perfect equilibrium. An important intermediate step in the proof was to establish the existence of an $\bar{\epsilon}_1$ such that (8) and (9)

held for all $\epsilon < \bar{\epsilon}_1$ and all $\delta \in [\underline{\delta}', 1)$. Substituting $\underline{\delta}'$ into these equations gives

$$(10) \quad \frac{\underline{\delta}'}{1 - \underline{\delta}'}(f(1, \underline{\delta}', 1, \epsilon) - f(2, \underline{\delta}', 1, \epsilon)) > g + \eta/2$$

and

$$(11) \quad \frac{\underline{\delta}'}{1 - \underline{\delta}'}(f(k, \underline{\delta}', 1, \epsilon) - f(k + 1, \underline{\delta}', 1, \epsilon)) < g - \eta/2 \quad \forall k \geq 2.$$

Restricting attention to values $\delta \in [\underline{\delta}', \frac{1 + \underline{\delta}'}{2}]$, we once again can easily establish bounds on the derivatives of the left hand sides of the equations (10) and (11). For example, using expression (A3) from the proof of Proposition 2 we get

$$\begin{aligned} & \frac{\partial}{\partial \delta} \left(\frac{\delta}{1 - \delta} (f(k, \delta, 1, \epsilon) - f(k + 1, \delta, 1, \epsilon)) \right) \\ &= \sum_{t=0}^{\infty} (t + 1) \delta^t (1 + g) \text{Prob}\{\alpha_i(t) \in C(t, k) \cap D(t) \cap \overline{E(t)}\} \\ &\leq \frac{1 + g}{(1 - \delta)^2} \\ &\leq \frac{4(1 + g)}{(1 - \underline{\delta}')^2}. \end{aligned}$$

Hence, we can find a value δ_1 such that for all $\delta \in [\underline{\delta}', \delta_1]$ and all $\epsilon \leq \bar{\epsilon}_1$ we have

$$(12) \quad \frac{\delta}{1 - \delta}(f(1, \delta, 1, \epsilon) - f(2, \delta, 1, \epsilon)) > g + \eta/4$$

and

$$(13) \quad \frac{\delta}{1 - \delta}(f(k, \delta, 1, \epsilon) - f(k + 1, \delta, 1, \epsilon)) < g - \eta/4 \quad \forall k \geq 2.$$

From here, the same steps as in the proof of Proposition 2 but with $\eta/2$ in place of η show that for sufficiently small ϵ , the strategies with $q = 1$ give an equilibrium for all $\delta \in [\underline{\delta}', \delta_1]$. Now, the construction in Lemma 2 gives us an equilibrium without public randomizations for all $\delta \in [\underline{\delta}'/\delta_1, 1)$.

A further consequence of Lemma 2 is that the per period payoff to a player with discount factor δ of the no randomization equilibrium $s^{**}(\delta, \epsilon)$ is *exactly* equal to the per period payoff that the strategies with $q = 1$ give a player with discount factor $\delta^{N(\delta)}$, $f(0, \delta^{N(\delta)}, 1, \epsilon)$. The function f is continuous in its second argument and $\delta^{N(\delta)} \rightarrow \underline{\delta}'$ as $\delta \rightarrow 1$, so for u_i being player i 's expected utility in the game with discount factor δ ,

$$\lim_{\epsilon \rightarrow 0} \lim_{\delta \rightarrow 1} u_i(s^{**}(\delta, \epsilon)) = \lim_{\epsilon \rightarrow 0} f(0, \underline{\delta}', 1, \epsilon).$$

This, however, is merely the limit of the expected payoff for a *fixed* discount factor as $\epsilon \rightarrow 0$ so efficiency in the limit is easy. For any $\gamma > 0$, we can simply choose T so that $(1 - \underline{\delta}')(1 + \underline{\delta}' + \dots + \underline{\delta}'^T) > 1 - \gamma/2$ then pick ϵ small enough so that with very high probability there are no ϵ -probability events in the first T periods, hence giving an expected payoff of at least $1 - \gamma$ in the game with ϵ noise. **QED**

If we had not worried about noise in this section, we could have found a perfect equilibrium without public randomizations whenever $\delta \in [\underline{\delta}, \delta_1]$ where $\underline{\delta}$ is defined by

$$\underline{\delta}(1 - f(2, \underline{\delta}, 1)) = (1 - \underline{\delta})g$$

and δ_1 is defined either by

$$\delta_1(f(3, \delta_1, 1) - f(4, \delta_1, 1)) = (1 - \delta_1)g$$

or by $\delta_1 = 1$ if the equation above has no solution. Table 2.2 gives $\underline{\delta}$, δ_1 and $\underline{\delta}/\delta_1$ for a range of values of g and M . For $\delta \geq \underline{\delta}$ a cooperative perfect equilibrium exists with public randomizations, and for $\delta \geq \underline{\delta}/\delta_1$ a perfect equilibrium exists without them. Note that for many of the parameter values, δ_1 is in fact equal to one. In this case, eliminating public randomizations does not require any additional patience on the part of the players. When $g = 0.01$, the difficulty in getting the players to carry out punishments results in much more patient play being necessary to support the equilibrium I have given.

Table 2.1: Discount factor sufficient to maintain cooperation

	$g = 0.01$	$g = 1$	$g = 10$
$M = 2$	0.01	0.50	0.91
$M = 4$	0.03	0.68	0.95
$M = 10$	0.08	0.79	0.97
$M = 100$	0.35	0.89	0.985
$M = 500$	0.54	0.92	0.989

Table 2.2: Discount factors with and without public randomizations

		$g = 0.01$	$g = 1$	$g = 10$
$M = 4$	$\underline{\delta}$	0.03	0.68	0.95
	δ_1	0.03	1.00	1.00
	$\underline{\delta}/\delta_1$	0.96	0.68	0.95
$M = 10$	$\underline{\delta}$	0.08	0.79	0.97
	δ_1	0.08	1.00	1.00
	$\underline{\delta}/\delta_1$	0.96	0.79	0.97
$M = 100$	$\underline{\delta}$	0.35	0.89	0.985
	δ_1	0.36	1.00	1.00
	$\underline{\delta}/\delta_1$	0.96	0.89	0.985
$M = 500$	$\underline{\delta}$	0.52	0.92	0.989
	δ_1	0.54	1.00	1.00
	$\underline{\delta}/\delta_1$	0.96	0.92	0.989

2.5 Conclusion

In all of the results above, cooperation has been sustained in equilibrium by the use of "contagious" punishments which lead eventually to a breakdown of cooperation after a single deviation. The results illustrate the extent to which the convexity of the breakdown process can be exploited, and the interesting patterns of play which can arise in equilibrium. In addition, the contagious punishments are a fairly powerful tool for enforcing cooperation. Besides the basic result that cooperation can be sustained despite a large population with infrequent individual interactions, we have seen that cooperation is still possible with heterogeneity in time preferences or without public randomizations.

The great advantage of contagious punishments, though, is not their power but rather their informational simplicity. In models of trade in large populations, we can support cooperation even in the extreme case of completely anonymous interactions. I have probably overemphasized how reasonable the discount factors are, but it is important to make the point that in models of trade, for example, contagious punishments could support cooperation. While equilibria based on contagious punishments may not correspond to what was observed in particular historical situations, they at least were feasible, suggesting that reputation-based institutions were not the only possible solution.

I have also argued that these results can be made far more robust than Kandori's first example suggests. Global stability is not a problem if public randomizations are available. In a model in which players are explicitly assumed to make mistakes, we can still sustain nearly efficient outcomes in a perfect equilibrium even if players do not know the exact frequency with which mistakes are made. In this abstract sense, the equilibrium is quite robust. Whether or not full efficiency can be achieved for a small fixed amount of noise is still open. Of course, near efficiency is achieved only in the limit, and for an economically reasonable frequency of mistakes the equilibrium may not be particularly close to efficient. It is probably robustness in this sense, not in the limit, which is most relevant when considering the possibility of applying contagious

punishments to models of trade. Hence, I am unable to say whether concerns about robustness prevent that application.

Finally, I should note that I have also left one major question of game-theoretic interest unanswered. The results of this paper rely heavily on the fact that the prisoner's dilemma has a dominant strategy equilibrium. In light of Kandori's Folk Theorem for games with a more complex information structure, it would be interesting to know whether the results of this paper extend to a more general class of games. If so, we would have a much more general Folk Theorem. If not, we would have a sharper picture of the type of information transmission which is necessary to maintain cooperation.

References

Aumann, R., and Shapley, L., "Long Term Competition – A Game Theoretic Analysis," Mimeo (1976), Hebrew University.

Friedman, J., "A Non-cooperative Equilibrium for Supergames," *Review of Economic Studies* 38(1971), 1–12.

Fudenberg, D., and Maskin, E., "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* 54(1986), 533–54.

Fudenberg, D., and Maskin, E., "On the Dispensability of Public Randomizations in Discounted Repeated Games," *Journal of Economic Theory* 53(1991), 428–438.

Fudenberg, D., Levine, D., and Maskin, E., "The Folk Theorem with Imperfect Public Information," Mimeo(1991), MIT.

Greif, A., "Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders," *Journal of Economic History* 69(1989), 857–82.

Kandori, M., "Social Norms and Community Enforcement," CARESS Working Paper #89-14 (1989), University of Pennsylvania.

Milgrom, P., North, D., and Weingast, B., "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs," *Economics and Politics* 2(1990), 1–23.

Okuno-Fujiwara, M., and Postlewaite, A., "Social Norms and Random Matching Games," CARESS working paper #90-18 (1990), University of Pennsylvania.

Rosenthal, R., "Sequences of Games with Varying Opponents," *Econometrica* 47(1979), 1353-66.

Appendix

Proof of Proposition 2

I begin by establishing equations (6) and (7) which are analagous to equations (1) and (2) from the proof of Proposition 1. I shall write $f(k, \delta, q, \epsilon)$ for the per period continuation payoff of player i when at the start of period t , k players (including player i if $k > 0$) are playing according to phase II of the strategies described in the proof of Proposition 1. I wish to show that there exists $\underline{\delta}' < 1$, $\eta > 0$ and a function $q' : [\underline{\delta}', 1) \rightarrow [0, 1]$ such that (6) and (7) hold for all $\delta \in [\underline{\delta}', 1)$.

Note first that because I have not yet introduced noise, $f(0, \delta, q'(\delta), 0) = 1$. I begin by establishing a degree of strict convexity of f . From equation (3) in the proof of Lemma 1 we know that

$$\begin{aligned} & ((f(1, \delta, q, 0) - f(2, \delta, q, 0)) - (f(2, \delta, q, 0) - f(3, \delta, q, 0))) \\ &= E_{\omega} \left[\sum_{t=0}^{\infty} (1 - \delta) q^t \delta^t (1 + g) I(o_1(t, \omega) \in (C(t, 1, \omega) - C(t, 2, \omega)) \cap D(t, \omega)) \right]. \end{aligned}$$

The second term of this sum is

$$(A1) \quad (1 - \delta) q \delta (1 + g) \text{Prob} \{o_1(1, \omega) \in (C(1, 1, \omega) - C(1, 2, \omega)) \cap D(1, \omega)\}.$$

If player 2 is matched with player M in period 0 under ω we have

$$\begin{array}{ll} 2 \in C(1, 1, \omega) & M \in C(1, 1, \omega) \\ 2 \notin C(1, 2, \omega) & M \notin C(1, 2, \omega) \\ D(1, \omega) = \{2, M\} & \end{array}$$

Together, these imply

$$(C(1, 1, \omega) - C(1, 2, \omega)) \cap D(1, \omega) = \{2, M\}.$$

From this, we know that the probability term in (A1) is at least the probability that players 2 and M are matched in period 0 and that player 1 subsequently is matched against one of 2 or M in period 1. This probability is $2/(M - 1)^2$.

Hence, for $\underline{\delta}$ as defined in Proposition 1, we have for any $\delta > \underline{\delta}$

$$(A2) \quad \frac{\delta}{1-\delta} ((f(1, \delta, 1, 0) - f(2, \delta, 1, 0)) - (f(2, \delta, 1, 0) - f(3, \delta, 1, 0)))$$

$$\geq \frac{2\underline{\delta}^2(1+g)}{(M-1)^2}$$

$$\equiv \gamma$$

From equation (4) we know that

$$\frac{\underline{\delta}}{1-\underline{\delta}} (f(1, \underline{\delta}, 1, 0) - f(2, \underline{\delta}, 1, 0)) = g.$$

From expansion (5) in the proof of Proposition 1 it is immediate that

$$\frac{\partial}{\partial \delta} (f(1, \delta, 1, 0) - f(2, \delta, 1, 0)) \big|_{\underline{\delta}} > 0.$$

Thus for some $\eta < \gamma/2$ we can choose $\underline{\delta}' \in (\underline{\delta}, 1)$ so that

$$\frac{\underline{\delta}'}{1-\underline{\delta}'} (f(1, \underline{\delta}', 1, 0) - f(2, \underline{\delta}', 1, 0)) = g + \eta.$$

By (A2) we know

$$\frac{\underline{\delta}'}{1-\underline{\delta}'} (f(2, \underline{\delta}', 1, 0) - f(3, \underline{\delta}', 1, 0)) < g - \eta.$$

Now, we simply set

$$q'(\delta) \equiv \underline{\delta}'/\delta$$

and note from (5) that $\forall \delta \in [\underline{\delta}', 1)$

$$\frac{\delta q'(\delta)}{1-\delta} (f(k, \delta, q'(\delta), 0) - f(k+1, \delta, q'(\delta), 0)) = \frac{\underline{\delta}'}{1-\underline{\delta}'} (f(k, \underline{\delta}', 1, 0) - f(k+1, \underline{\delta}', 1, 0)).$$

As $q'(\delta) > q(\delta)$, players will not deviate in phase I of a model with no noise so

$$f(0, \delta, q'(\delta), 0) \geq f(1, \delta, q'(\delta), 0).$$

This establishes (6) and (7) as desired.

The next major step in the proof is to establish that the similar inequalities (8) and (9) hold for a model with sufficiently little noise. To do this, I extend expansion (3) to a model with noise. Note that

$$(A3) \quad \begin{aligned} & f(k, \delta, q, \epsilon) - f(k+1, \delta, q, \epsilon) \\ &= E_{\omega} \left[\sum_{t=0}^{\infty} (1-\delta)q^t \delta^t (1+g) I\{o_1(t, \omega) \in C(t, k, \omega) \cap D(t, \omega) \cap \overline{E(t, \omega)}\} \right] \end{aligned}$$

where a realization of ω now includes also the set of players who “tremble” and play D accidentally in each period and $E(t, \omega)$ is defined to be the set of players affected by an ϵ -probability tremble up to and including time t . If $T(t, \omega)$ is the set of players who tremble at time t for a realization of ω , $E(t, \omega)$ can be formally defined by

$$\begin{aligned} E(0, \omega) &= T(0, \omega) \\ E(t+1, \omega) &= E(t, \omega) \cup T(t+1, \omega) \cup \{i | o_i(t, \omega) \in E(t, \omega)\} \end{aligned}$$

Using the expansions (3) and (A3) we get

$$(A4) \quad \begin{aligned} & \frac{\delta q'(\delta)}{1-\delta} ((f(k, \delta, q'(\delta), 0) - f(k+1, \delta, q'(\delta), 0)) - (f(k, \delta, q'(\delta), \epsilon) - f(k+1, \delta, q'(\delta), \epsilon))) \\ &= E_{\omega} \left[\sum_{t=0}^{\infty} q'(\delta)^{t+1} \delta^{t+1} (1+g) I\{o_1(t, \omega) \in C(t, k, \omega) \cap D(t, \omega) \cap E(t, \omega)\} \right] \\ &\leq E_{\omega} \left[\sum_{t=0}^{\infty} \delta^{t+1} (1+g) I\{o_1(t, \omega) \in E(t, \omega)\} \right] \end{aligned}$$

Given $\eta > 0$ as defined above, we can choose T such that

$$\frac{\underline{\delta}^{T'}}{1 - \underline{\delta}'} < \frac{\eta}{4(1 + g)}.$$

Next, choose $\bar{\epsilon}_1$ sufficiently small such that

$$\text{Prob} \{E(T, \omega) \neq \emptyset\} < \frac{\eta}{4(1 + g)(1 + \dots + \underline{\delta}^{T'})}.$$

Now, for any $\delta \in [\underline{\delta}', 1)$ and any $\epsilon < \bar{\epsilon}_1$ the right hand side of equation (A4) is bounded above by $\eta/2$. This and equations (6) and (7) gives

$$\frac{\delta q'(\delta)}{1 - \delta} (f(1, \delta, q'(\delta), \epsilon) - f(2, \delta, q'(\delta), \epsilon)) > g + \eta/2$$

and

$$\frac{\delta q'(\delta)}{1 - \delta} (f(2, \delta, q'(\delta), \epsilon) - f(3, \delta, q'(\delta), \epsilon)) < g - \eta/2.$$

The first equation is (8). Using the expansion (A3) in place of (3) it is easy to see that the result for Lemma 1 carries over to the model with ϵ noise. This and the second equation above gives us (9).

Now that (8) and (9) have been established, I proceed to show that there are no profitable deviations from either phase I or phase II play in the ϵ -constrained game. The phase II case is easier so I'll start with that. Note that we can rewrite (9) to give $\forall \epsilon < \bar{\epsilon}_1, \delta \in [\underline{\delta}', 1)$, and $k \geq 2$,

$$\delta q'(\delta) (f(k, \delta, q'(\delta), \epsilon) - f(k + 1, \delta, q'(\delta), \epsilon)) < (1 - \delta)g.$$

As in the proof of Proposition 1, the right hand side of this expression is the short term loss when a player plays C instead of D in phase II and is matched with someone who plays C. The expectation over k of the left side is the expected future gain. Clearly, the future gain is too small to make a deviation profitable.

The discussion of phase I play is more complicated than before because a player in phase I must assign probability $r_k > 0$ to the event that unbeknownst to him, k

other players are already playing according to phase II or will tremble and play D in the current period. Keep in mind that r_k is a function both of ϵ and of the history of the game. I show, however, that for ϵ sufficiently small this uncertainty is small regardless of the history of the game.

To show that player 1's best response whenever he is in phase I in period t is to play C, I will *not* show directly that his expected payoff from playing C in period t and then following his equilibrium strategy is better than his expected payoff from playing D in period t then following his equilibrium strategy. Instead, I compare the payoff from playing C in period t then switching to phase II play in period $t+1$ to the payoff from playing D in period t and continuing according to phase II. Player 1's period t action has no effect on play after any period $t+s$ in which $q_{t+s} > q'(\delta)$. We have already seen that playing D in phase II is a best response so that the latter strategy gives the greatest possible expected payoff to a player who plays D in period t . If the former is greater, the best response must involve playing C in period t .

To compare the payoffs of the two strategies, look first at the period t outcome. If player 1 plays D in period t he gains g whenever $o_1(t)$ plays C and avoids a loss of ℓ whenever $o_1(t)$ plays D. Hence the short term gain is

$$\sum_{k=0}^{M-1} r_k \left(\frac{M-k-1}{M-1} g + \frac{k}{M-1} \ell \right).$$

In the future, a player who plays D in period t can never be better off because both strategies prescribe the same play from period $t+1$ on and there are always either the same number or more players in phase II in period $t+1$. When $k=0$ and there are also no ϵ -probability trembles in period $t+1$, the player who plays D in period t is worse off, obtaining a continuation payoff of $f(2, \delta, q'(\delta), \epsilon)$ instead of $f(1, \delta, q'(\delta), \epsilon)$. The discounted expected loss is then at least

$$r_0(1-\epsilon)^{M-1} \frac{\delta q'(\delta)}{1-\delta} (f(1, \delta, q'(\delta), \epsilon) - f(2, \delta, q'(\delta), \epsilon)).$$

To show that playing C is better in period t it thus suffices to show that this loss

outweighs the short term gain. Using (8), it will suffice to show

$$r_0(1 - \epsilon)^{M-1}(g + \eta/2) \geq r_0g + (1 - r_0) \max(g, \ell).$$

We can choose $\bar{\epsilon}_2$ such that

$$(1 - \epsilon)^{M-1}(g + \eta/2) \geq g + \eta/4$$

for all $\epsilon < \bar{\epsilon}_2$. It then only remains to establish

$$(A5) \quad r_0 \eta/4 \geq (1 - r_0) \max(g, \ell)$$

for ϵ sufficiently small.

At first look one might think that if the game has been going on long enough, then player 1 will be fairly sure that someone must have trembled. This reasoning suggests that the r_0 term might not dominate in (A5). However, it is important to keep in mind that r_0 is not an unconditional probability, but rather the conditional probability that no one has trembled since the last time s that $q_s > q'(\delta)$ occurred *given* that no opponent of player 1 has played D since that time. To show in fact that

$$(A6) \quad \lim_{\epsilon \rightarrow 0} \inf_{h_t} r_0 = 1$$

take any $\zeta > 0$. We can choose T_1 so that

$$\text{Prob}\{\text{Player 1 is still in phase I} \mid \text{Some player was in phase II } T_1 \text{ periods ago}\} < \zeta/2$$

Next choose $\bar{\epsilon}_3$ so that

$$(1 - \bar{\epsilon}_3)^{T_1 M} > 1 - \zeta/2$$

Then for any $\epsilon < \bar{\epsilon}_3$, the probability that there has been a tremble in the last T_1 periods given that none has been observed is less than $\zeta/2$. Hence, $r_0 > 1 - \zeta$ regardless of the number of periods that have elapsed since the last time $q_s > q'(\delta)$

occurred. This implies (A6), and hence choosing $\bar{\epsilon}$ smaller than $\bar{\epsilon}_1$, $\bar{\epsilon}_2$, and $\bar{\epsilon}_3$ we get the sufficient condition (A5) for no deviations in phase I. This concludes the proof that $s^*(\delta, \epsilon)$ is a perfect equilibrium of the ϵ -constrained game.

Finally, the proof of 3., that we get efficiency in the limit, is relatively easy. Consider the largest possible effect that a single tremble by player j in period t can have on player i 's total payoff anywhere on the path of the equilibrium with noise. This tremble can only affect player i 's payoff in period t and in any future period until the first time $q_{t+s} > q'(\delta)$. Thus, the expected loss caused by this single tremble is at most

$$\delta^t \sum_{s=0}^{\infty} (1-\delta) \delta^s q'(\delta)^s (1+g+\ell) = (1+g+\ell) \frac{\delta^t (1-\delta)}{1-\underline{\delta}'}$$

Player i 's expected per period payoff is equal to 1 minus the expected loss from each possible tremble times the probability of that tremble occurring. This gives

$$\begin{aligned} f(0, \delta, q'(\delta), \epsilon) &\geq 1 - (1-\delta) \sum_{t=0}^{\infty} (1+g+\ell) \frac{\delta^t}{1-\underline{\delta}'} M \epsilon \\ &= 1 - \frac{(1+g+\ell) M \epsilon}{1-\underline{\delta}'} \end{aligned}$$

Clearly

$$\lim_{\epsilon \rightarrow 0} \lim_{\delta \rightarrow 1} f(0, \delta, q'(\delta), \epsilon) = 1.$$

QED

Chapter 3

Markov Models of Price Wars in the Joint Executive Committee

3.1 Introduction

The “Joint Executive Committee” (JEC) was a railroad cartel organized in 1879 to set prices for transport between Chicago and the East Coast. The detailed records it kept have already been the basis for several insightful studies of cartel structure. Several theoretical papers discuss the problem of cartel stability in a repeated game framework (Green and Porter, (1984); Rotemberg and Saloner, (1986); Abreu et. al., (1986)). While these models are all based on the idea that the threat of a price war can enforce cartel discipline, each leads naturally to different conclusions about the apparent causes of price wars. In this paper, I focus on these dynamic predictions in order to assess the applicability of the theory to the experience of the JEC.

The first section of this paper briefly reviews the historical background of the JEC. Next, I review two main theories of cartel stability and discuss the modifications necessary to apply each to the situation of the JEC. The first theory, due to Green and Porter (1984), views price wars as necessary to maintain incentives with imperfect information. The second, due to Rotemberg and Saloner (1986), suggests that price wars may result when the gains to deviation become large and is associated with the view that price wars are likely to break out during booms. Proponents of both

theories have cited the JEC as providing empirical support for their views (Lee and Porter, 1984; Porter, 1983; Porter, 1985; Rotemberg and Saloner, 1986).

The econometric approach of this study is to apply a switching regressions model in which the shifts in regimes follow a Markov process. (Coslett and Lee (1985) introduced such a model to test whether the regimes were independent, and Hajivassiliou (1989) has recently discussed extensions similar to those of this paper.) I further allow the transition probabilities for the regime shifts to be influenced by predetermined variables in a way which allows us to test whether the pattern of price wars is consistent with the predictions of each theory. In contrast with the two step approach of Porter (1985), where each week's behavior is first classified as collusive or non-collusive based on structural variables alone, I simultaneously estimate both the structural parameters of supply and demand *and* the dynamics of regime shifts. In the final section of this paper, I discuss the estimates of several related Markov models which yield reasonable estimates of the structural parameters and exhibit the expected pattern of long collusive periods separated by occasional price wars. Generally, the empirical evidence supports the existence of a Green-Porter type mechanism with unusual demand patterns triggering price wars. Some results are also consistent with the Rotemberg-Saloner model, although the evidence here is not strong.

3.2 History

MacAvoy (1965) and Ulen (1978) have given thorough descriptions of the "Joint Executive Committee" and its attempts to control prices for grain transport between Chicago and the Eastern seaboard, so I shall only give a brief overview of the time period including a few facts relevant to arguments made later in this paper.

In the 1870's, railroads began to carry grain and other goods from Chicago to the Eastern seaboard. Larger quantities of grain were shipped by steamer over the Great Lakes during the summer months. In 1874, the Baltimore and Ohio railroad extended its lines into Chicago and began competing with the New York Central and Pennsylvania railroads. After several unsuccessful attempts to control prices, the

railroads agreed in April of 1879 on a mechanism for cartel control. The newly formed “Joint Executive Committee” was given the authority to establish a uniform set of rates for all three railroads, and an administrative structure was set up to handle any disputes. Throughout the period analyzed in this paper, 1880-1886, the JEC continued to establish both official rates and market share allotments for traffic out of Chicago. At times, the cartel successfully maintained rates as high as 40 c/100 lbs. In parts of 1881, 1884, and 1885, the cartel officially endorsed price wars, and prices dropped as low as 12.5 c/100 lbs.¹

The JEC collected and disseminated the data which we analyze here. The nature of those data is important to arguments made later in this paper. The JEC polled the member firms to produce an index of prices, and collected data on weekly grain shipments by each firm out of Chicago. While the reported prices probably do not reflect secret price cuts, we do have reason to believe that the quantities are reliable.² We know that in 1882, the cartel agreement was modified to require cash payments by any firms exceeding their assigned market shares. As the cartel evidently lacked the power to collect these payments, this plan proved unsuccessful.³ However, that the JEC would agree to rely on its market share data for cash payments indicates that they had sufficient auditing capability to ensure that firms could not underreport shipments.

3.3 Theory

The switches between collusive and noncooperative behavior in the JEC were associated with dramatically different pricing strategies and will allow us to study the regime shifts econometrically. The theoretical models which have been developed to describe pricing in optimal cartels have implications for the nature of regime shifts. I

¹For example, MacAvoy cites an 1885 report that “All official rate-setting procedures were suspended and all roads authorized to meet any cut rates.” MacAvoy, 1965, p. 102

²The contemporary trade press contains many reports of secret price cutting. See the *Daily Commercial Bulletin* of June 16, 1881, September 15, 1884, and February 2, 1886 for examples.

³MacAvoy reports that the Grand Trunk railway failed to make any payments in 1884, and left the cartel in 1885 owing over \$100,000. MacAvoy, 1965, p. 190

review two main lines of thought and discuss some specific considerations necessary to apply these models to the JEC.

Porter (1983b) formulated his original paper on the JEC as a test of the Green-Porter (1984) model. The original Green-Porter model considers the case of repeated Cournot competition between identical firms. Production levels are unobservable, and there is noise in demand; therefore firms can only observe an imperfect signal of their competitors' behavior. Green and Porter consider whether collusion can be sustained by trigger strategies which involve switches between collusive states and price wars. In a price war the firms produce at the Cournot level, the unique equilibrium of a single stage game. In a collusive state the firms produce at a point between the Cournot and monopoly levels. Because of the noise in demand, low prices in a collusive state can result either from low demand or from cheating on the production limits. The optimal equilibrium in trigger strategies has firms produce at the collusive level until the market price falls below a specified level at which point a T period price war begins. In equilibrium, no firm deviates from the collusive production levels, but there are occasional price wars due to demand shocks.

Abreu, Pearce, and Stachetti (1986) extend the Green-Porter model to more general games and characterize the optimal equilibrium among all strongly symmetric strategies, not just trigger strategies.⁴ They find a similar pattern, but with the switches between the collusive and price war states following a first order Markov process involving only two states rather than using T-period punishments (which require $T + 1$ states).

Several modifications are necessary to apply this theory to the JEC. It is certainly more reasonable to think of the railroads not as setting quantities but rather as setting prices with reputations for quality of service and differing route structures creating differentiated demands. As a rough approximation, we can think of price wars as involving marginal cost pricing, although with increasing costs and differentiated goods the optimal equilibrium could very well involve either higher or lower prices in a price war. In collusive states prices increase to a point closer to the monopoly level. What

⁴Strategies are strongly symmetric if all players take the same action after any history.

is and is not observable also changes. With secret price cuts and no single market price, price is not observable. As noted earlier, firm-specific quantities presumably are observable. Also unlike the standard theory, we do not have a perfectly symmetric model. Observed market shares would allow the JEC to infer which firms may have cheated so the optimal equilibrium would involve punishing deviating firms more severely, not using symmetric price wars.⁵ However, I have only limited data with which to test dynamic predictions and strongly symmetric strategies do appear to have been used, so I shall stick with the simplest cartel structures.

Because a variety of trigger mechanisms can be used to sustain a nearly optimal cartel, it may be uninformative to focus too sharply on testing whether the optimal equilibrium is used. Instead, I attempt to test whether an equilibrium similar to that of Green and Porter (1984) is played. In such an equilibrium, we would expect to see price wars triggered by one of several possible signals of cheating. The best signal with observable market shares would probably be some pattern of a high market share for one firm and lower market shares for all the rest. A second possible signal is *high* aggregate demand, as when one or more firms lower prices, there is not only a transfer of market shares, but also an expansion of demand. Finally, a firm could use its own sales or market share as a signal because unusually low sales may result from cheating by another firm.

A different perspective on cartel stability is found in Rotemberg and Saloner (1986). They look at firms competing in prices, and consider the possible equilibria that can be supported by reversion to the Bertrand outcome. While demand is subject to i.i.d. random shocks, the random component of demand is observable in advance of each period so there is no problem in detecting collusion. They note that monopoly pricing may not always be sustainable if potential punishments are short or firms have small discount factors. In this case, the discounted value of future losses may not be sufficient to outweigh the short term gains from deviation. In particular, when the random component of demand is large, short term gains are larger, but future losses

⁵The results of Fudenberg, Levine, and Maskin (1991) imply that for sufficiently patient players, we can achieve the fully collusive outcome.

are unchanged. In order to lessen the incentive to deviate, the optimal cartel must lower prices when the random component of demand is high.

In order to apply rigorously the ideas of Rotemberg and Saloner to the situation of the JEC, it would be necessary to add time varying demand to a Green-Porter style model with unobservable prices. Certainly, it must be true that firms will find it harder to collude when demand is expected to decline. However, in adapting a cartel design to a situation of declining demand a variety of changes are possible. Prices in each phase may be raised or lowered, price wars may be lengthened or shortened, and the sensitivity of the triggers may be adjusted. Price wars will be more common during booms than at other times only if the optimal adjustment makes triggers more sensitive. In the absence of a formal model, it is possible to reason by analogy regarding the problem of booms as similar to a lowering of firms' discount factors. The results of Porter(1983a) then suggest that price wars might be more likely to occur during booms, but that any results to this effect will be sensitive to assumptions on the distribution of the error terms.

We can hope to test empirically whether price wars in the JEC tended to occur when the incentives to deviate were unusually high. The incentive to deviate will be high whenever high current demand is expected to be followed by lower future demand. In studying the nature of the railroad industry, we can identify two factors which create such a situation. First, we have the seasonal pattern of demand influenced by the closing and opening of the Great Lakes. When the lakes were closed, the JEC received both higher demand and higher prices. Near the end of the winter, the short term gain to deviation is high, but future losses decline as part of these losses would take place after the lakes have melted and prices and output would have been lower anyway. Second, under the specific assumptions made about the nature of the random demand shocks, booms die out slowly over time and a large random component of demand today increases the present gains to deviation with a less than proportional increase in future losses. This prediction is similar to the commonly discussed interpretation of the Rotemberg-Saloner model as predicting price wars during booms. However, with increasing costs, slow changes in demand, and short

price wars, this effect may not be large.

3.4 Econometric Model

The models of this paper are based on that of Porter (1983b). The primary difference is that where Porter assumes that each period's regime is independent of both past and future regimes, I impose a Markov structure on the regime switches and maximize the joint likelihood of both the observed prices and quantities and of the inferred regime transitions. The estimation of models with different variables influencing transition probabilities allows us to test predictions of the theories above.

To facilitate comparisons, I follow Porter's specification of the model closely. The demand for grain shipments in week t is assumed to be given by

$$(1) \quad \log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 Lakes_t + \alpha_{3-14} Seasxx_t + U_{1t}$$

where $Lakes_t$ and $Seasxx_t$ are intended to capture seasonal variation in demand (see section 6). In a departure from Porter's model, I usually assume that U_{1t} follows an AR(1) process,

$$(2) \quad U_{1t} = \rho U_{1t-1} + V_{1t} \quad |\rho| < 1.$$

The price elasticity of demand is α_1 . We expect that α_2 will be negative to reflect the diversion of shipments when the lakes are open.

Porter and others have noted that if N identical oligopolists with costs

$$C(Q) = aQ^\delta$$

face isoelastic demand as above, their pricing behavior is summarized by

$$P_t(1 + \theta/\alpha_1) = a\delta Q_t^{\delta-1}$$

with θ equal to 1 with monopoly pricing, $1/N$ with Cournot competition, and 0 with marginal cost pricing. Our estimates suggest that δ is greater than one so that marginal costs are increasing and an exogenous increase in demand leads to price increases.

In our standard model of the cartel, we thus assume the supply equation to have the form

$$(3) \quad \log P_t = \beta_0 + \beta_1 \log Q_t + \beta_2 I_t + \beta_{3-6} DMx_t + U_{2t}$$

where I_t is an indicator of collusion, and the DMx are four dummy variables for changes in the cartel composition. The errors V_{1t} and U_{2t} are assumed to have a multivariate normal distribution. While the possibility of secret price cuts makes P_t unobservable, we do have available the official prices firms were supposed to have charged. While using this series may introduce an errors-in-variables problem, the optimal cartel theory predicts that firms should in fact never deviate from the official prices, so we can hope that these errors are small.

I assume that there are two possible behavioral states in each time period indexed by I_t . $I_t = 1$ indicates a more collusive state, and $I_t = 0$ indicates a state of lesser collusion which I shall describe interchangeably as a noncooperative or price-war state. If the values of θ in the two states, θ_1 and θ_0 are both unknown, estimation of the supply equation yields

$$\exp \beta_2 = \frac{(\alpha_1 + \theta_0)}{(\alpha_1 + \theta_1)}.$$

Thus, we can only estimate a nonlinear function of the two θ s. Note that

$$\theta_1 - \theta_0 = -\alpha_1(1 - e^{-\beta_2}) - \theta_0(1 - e^{-\beta_2})$$

When price wars involve marginal cost pricing θ_0 is zero. Hence, to approximate the difference between θ_1 and θ_0 I follow Porter in using

$$(4) \quad \theta \equiv -\alpha_1(1 - e^{-\beta_2})$$

as a rough measure of the differing degrees of collusion between the two states.

The standard model of this paper assumes that the cartel adopts a behavior which involves unobserved transitions between two states. Given the state I_t , we assume that price and quantity are determined by (1) and (3). Writing the model as

$$YB = X\Gamma + S\Delta + V$$

$$\begin{aligned} Y_t &= (\log Q_t \log P_t) \\ X_t &= (1 \text{ Lakes}_t \text{ DM}x_t \text{ Seas}xx_t U_{1t-1}) \\ S_t &= (0 \ I_t) \\ V_t &= (V_{1t} \ U_{2t}) \end{aligned}$$

with B , Γ , and Δ defined appropriately, the likelihood function for $\log Q_t$ and $\log P_t$ (I shall write Y_t for $(\log Q_t \log P_t)$ to save space) given that state I_t arises and given all predetermined variables Z_t is

$$f(Y_t|I_t, Z_t) = \frac{1}{2\pi} |\det(\Sigma)|^{-\frac{1}{2}} |\det(B)| e^{-\frac{1}{2} V_t \Sigma^{-1} V_t'}$$

To determine the log likelihood of the entire sample, we assume the regime switches follow a Markov process as shown in Figure 3-1.

Within each time period, there are two possible states, the upper state labelled collusive ($I_t = 1$) and the lower state labelled price war ($I_t = 0$). Within each state, P_t and Q_t will be generated randomly according to (1), (2), and (3). At the end of each time period, a probabilistic transition determines the state I_{t+1} of the following period. While the optimal cartel would likely have discontinuous transition probabilities based on cutoff values for observed variables, I assume that the likelihood of the transition is given by a logit model

$$(7) \quad \text{Prob}\{I_{t+1} = 1|I_t, Z_t\} = \frac{e^{\gamma W_t}}{(1 + e^{\gamma W_t})}$$

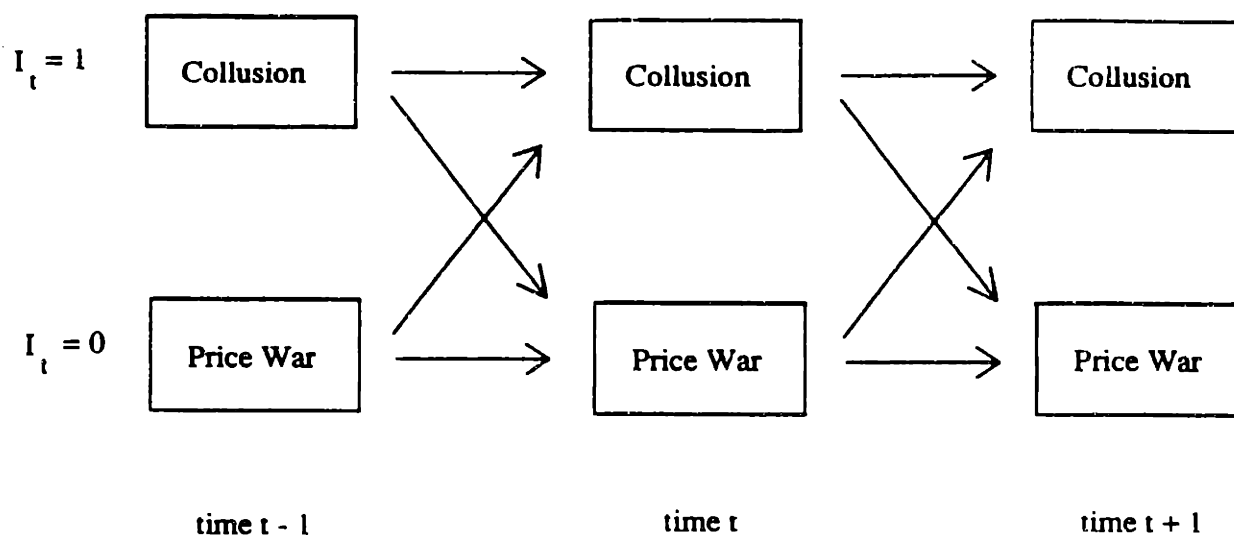


Figure 3-1: Regime Transitions

Taking W_t to be a constant independent of I_t gives a standard model with independent regime switches as in Porter (1983b). Allowing W_t to contain I_t adds a Markov structure to the price wars as in Coslett and Lee (1985) so that a noncooperative state today can be likely to lead to another noncooperative state next week. I also include several other variables in W_t to test the predictions of the Green-Porter and Rotemberg-Saloner models that certain factors should trigger price wars.

The likelihood of the sequence $(Y_t, Y_{t+1}, \dots, Y_T)$ is given by the the sum of the joint likelihoods of $(Y_t, Y_{t+1}, \dots, Y_T)$ and $(I_t, I_{t+1}, \dots, I_T)$ over all possible paths $(I_t, I_{t+1}, \dots, I_T)$. As described in Coslett and Lee (1985), the first order nature of the Markov process allows a simple calculation using a recurrence relation. At $t = 1$, the firms are assumed to be in a collusive state so that

$$f(Y_1, I_1 = 1 | R_1) = f(Y_1 | I_1 = 1, R_1)$$

$$f(Y_1, I_1 = 0 | R_1) = 0$$

(with R_t the set of exogenous variables known at time $\leq t$). At $t > 1$, the joint likelihood is obtained by summing the likelihood given the two possible transitions

from the state I_{t-1} :

$$f(Y_1, \dots, Y_t, I_t | R_t) = \sum_{i=0}^1 \{f(Y_t | I_t, Z_t) \times Prob(I_t | I_{t-1} = i, Z_{t-1}, Y_{t-1}) \times f(Y_1, \dots, Y_{t-1}, I_{t-1} = i | R_t)\}.$$

Maximum likelihood estimates of the parameters were computed by an iterative algorithm from this equation. Numerical derivatives were used to compute standard errors for all estimates.

In interpreting the results of the models, it is useful to be able to examine the classification of the sample into collusive and price-war states. Unlike the independent switching regression model, the likelihood of I_t given that I_{t-1} is not observed depends not only on past values but also on future values of P_t and Q_t . The maximum likelihood classification is computed by Bayes rule as

$$f(I_t | Y_T, Z_T) = f(I_t | Y_t, Z_t) f(Y_{t+1}, \dots, Y_T | I_t, Z_{t+1}, R_T) / K_t.$$

The first term on the right side is computed by Bayes rule from the values

$$f(Y_1, \dots, Y_t, I_t | R_t).$$

The second term is computed iteratively working backwards from the end of the data by

$$f(Y_t, \dots, Y_T | I_{t-1}, Z_t, R_T) = \sum_{i=0}^1 Prob(I_t = i | I_{t-1}, Z_{t-1}, Y_{t-1}) \times f(Y_t | I_t = i, Z_t) \times f(Y_{t+1}, \dots, Y_T | I_t, Z_{t+1}, R_T).$$

The constant K_t makes the probabilities sum to 1.

3.5 Previous Results

The data of this chapter have been analyzed several times in the past. Before proceeding with the promised results of the models above, I discuss the nature and

conclusions of some prior works.

Porter's (1983b) original paper on the subject developed the structural supply and demand framework given above, and investigated the appropriateness of regime switching models. He estimates the model by maximum likelihood assuming *independent* regime determination. (This is extended by Lee and Porter (1984) to consider the probabilistic information provided by a series of reported regime shifts as well.) The main thrust of Porter is then simply to establish that switches between "collusive" and "non-cooperative" regimes did occur, and to examine the effect on prices. Porter's model yields reasonable parameter estimates. Most importantly, the value of the collusion dummy in the supply equation is 0.545 indicating that there were large price changes associated with the regime shifts. The measure θ of collusion is only 0.336 indicating that the changes in behavior were far smaller in magnitude than a switch from monopolistic to marginal cost pricing. One puzzling finding is that the price elasticity of demand is only -0.800. Lee and Porter (1984, p. 412) propose that this could result from a cartel unable to raise prices to fully collusive levels because of incentive problems.

The papers above discuss the causes of cartel breakdown only briefly. In a later paper, Porter (1985) focuses on the causes and durations of price wars. The estimation method is a two stage process. First, he computes the maximum likelihood regime classification, I_t , as in the Lee-Porter model. Then, to examine the causes of price wars, he runs a probit regression of I_t on several explanatory variables over a sample of points where $I_{t-1} = 1$. The coefficient on the previous period's demand, Q_{t-1} , is negative and significant. He finds this curious because he expects price wars to be triggered by low market shares, not high aggregate demand. A variable intended to measure deviations in market share has the right sign, but is only marginally significant.

Several papers have added some structure to the price wars. Berry and Briggs (1988) compare simple frequency counts for the same I_t series and show that the regime of the previous week is important in determining collusion, but are unable to prove the importance of any further history. Coslett and Lee (1985) had previously

discussed the estimation of a model like ours with first order Markov transitions, and found that the regimes are not independently determined. While 77% of the periods exhibit cooperation, a cooperative state occurs with probability 0.96 after a cooperative state, but with probability only 0.12 after a price-war state. However, in allowing for serial correlation in the pricing equation without correcting for the discreteness of prices, they produced a model with many unreasonable parameter values. Like Berry and Briggs, they are also concerned only with independence and do not discuss the causes of price wars.

Finally, Hajivassiliou (1989) has independently detailed the estimation of a sophisticated model which is quite similar to the one used here in that it allows for a Markov structure on the transition probabilities influenced by predetermined variables, and which further allows for several forms of measurement error I have not included. Hajivassiliou focuses on the econometric issues and is not as successful in applying his methodology to testing theories of cartel stability. To test the Abreu et. al. model, he again verifies that his model indicates a Markov structure on the transition probabilities, but does not look any further to identify what strategies are being used to trigger the price wars. When he does include explanatory variables on the transition probabilities in an attempt to test the Rotemberg and Saloner theory two difficulties arise. First, because the variables are modeled as equally affecting the probabilities of both entering and remaining within a price war, it is hard to interpret the estimates as identifying causes of price wars. Second, recall that the Rotemberg and Saloner theory predicts that collusion is more difficult if high current demand is likely to be followed by reduced demand throughout the duration of the ensuing price war (here about 15 weeks). It seems unlikely that the measure of Midwest grain production he includes will indicate high demand for shipments one week followed by greatly reduced demand in the next few weeks.

Rotemberg and Saloner (1986) themselves cite the JEC as providing evidence that price wars tend to occur in years of high demand. However, their analysis is based only on annual averages and hence does not necessarily reflect any relationship between the onset of price wars and the demand conditions prevailing in those weeks.

3.6 Data

Thanks to the data collection efforts of the JEC, weekly data on prices and firm specific shipments are available on a week by week basis for a 328 week period from the first week of 1880 to 1886. The data in this paper were originally compiled by Ulen (1978) and have been used in all the papers cited above. Table 3.1 gives summary statistics for all of the variables.

Table 3.1: Summary Statistics

Variable	Mean	Std. Dev.	Minimum	Maximum
<i>Quantity</i>	25384.	11632.	4810.	76407.
<i>Price</i>	0.2465	0.0665	0.125	0.400
<i>Lakes</i>	0.5732	0.4954	0	1
<i>DM1</i>	0.4238	0.4949	0	1
<i>DM2</i>	0.0457	0.2092	0	1
<i>DM3</i>	0.4329	0.4962	0	1
<i>DM4</i>	0.0152	0.1227	0	1
<i>RS_Seasonal</i>	1.0153	0.2416	0.6707	1.5223
<i>BigShare1</i>	1.0908	0.5690	0.0397	2.9711
<i>Bigshare2</i>	1.2409	0.7102	0.1849	4.2345
<i>BigshareQ</i>	1.1744	0.5162	0.1578	2.9728
<i>SmallShare</i>	1.2303	0.7373	0.1161	5.7565

Quantity gives the total eastbound grain shipments of the railroads in the JEC in tons per week. *Price* is given in dollars per 100 lbs. of grain. Note also that except for a brief price of 12.5 cents/100 lbs., the price was always set at a multiple of 5 cents/100 lbs. The *Lakes* variable is a dummy variable set to one when the Great Lakes were navigable and steamers could compete with the railroads. We expect lower prices and quantities when the lakes are open. The four structural dummies *DM1-DM4* are taken directly from Porter (1983b) and mark significant changes in the membership of the cartel. These changes could affect both the cost structure of the cartel and the degree of sustainable collusion, but are included in the model simply as shifts in the supply equation. Seasonal dummy variables labelled *SEAS1-SEAS12* are used for the first twelve four-week periods of each year. All of these variables are exactly as

in Porter (1983b).

In addition, five more complicated variables were created to test the Green-Porter and Rotemberg-Saloner models. In the case of the Green-Porter model, the variables are not guided by a definitive theory, but rather are a rough attempt to capture a few of the possible workings of a class of cartels. The variable *BigShare1* is intended to be a plausible trigger in a cartel which switches to a price war when one firm obtains a suspiciously high market share. The measure is based on deviations in $\log q_{it}$ rather than q_{it} to try to provide incentives which are roughly independent of firm size. Recognizing that the measure is arbitrary, it seemed prudent to construct two variants on the theme, *BigShare2* and *BigShareQ*. Each of the variables is the largest amount by which the strength s_{it} of any firm's demand exceeds its predicted value \bar{s}_{it} (after normalizing by the sample standard deviation), *i.e.* each is of the form $\max_i (s_{it} - \bar{s}_{it}) / \sigma_i$. In the case of *BigShare1* and *BigShareQ*, $s_{it} = \log q_{it} - \frac{1}{n} \sum_j \log q_{jt}$ is used to represent the strength of firm *i*'s demand. In the case of *BigShare2*, the ordinary market share is used. The predicted values \bar{s}_{it} used in constructing *BigShare1* and *BigShare2* are taken to be the average of the same measure for the previous 12 weeks. The predicted value for *BigShareQ* is computed from the assigned quotas. While the use of quotas as expectations initially seems appealing, in practice the quotas were often set unrealistically and did not reflect the market shares that would result from adherence to a uniform price.⁶ Adjustments to the predictions had to be made at the start of the series and as the cartel structure shifted. A fourth variable, *SmallShare1*, is intended to be a plausible trigger for a cartel in which price wars are triggered by an unusually small market share for one firm. Using the same measures as in *BigShare1*, it is the largest amount by which the strength of any firm's demand falls short of its predicted value.

The final variable, *RS_Seasonal*, is created to test the Rotemberg-Saloner theory. It is designed to capture the seasonal variation in the ratio of current demand to the average level of demand which will prevail throughout the duration of an ensuing

⁶For example, MacAvoy (1965, p. 88) reports that the Chicago and Grand Trunk Railway was admitted with a 10% quota while receiving only a 2-7% share at the cartel price.

price war. The measure is

$$RS_Seasonal_t = K \times \frac{\exp(\hat{\alpha}_2 Lakes_t + \hat{\alpha}_{3-14} Seas\pi\pi_t)}{\sum_{s=1}^{52} (\hat{p}_{w,t+s} - \hat{p}_{c,t+s}) \exp(\hat{\alpha}_2 Lakes_{t+s} + \hat{\alpha}_{3-14} Seas\pi\pi_{t+s})}$$

The term used for discounting, $(\hat{p}_{w,t+s} - \hat{p}_{c,t+s})$, is the extra probability that the system will be in a price war at time $t+s$ if a price war is begun at time t . Coefficient estimates from the standard model in Table 3.2 are used for both these probabilities and the $\hat{\alpha}_i$.

3.7 Results

We are now finally in a position to examine the results of the models described above. First, I present the results of a “standard” Markov model, and discuss the departures from previous models. The incorporation of an autoregressive error term yields a larger estimate of the degree of collusion. Next, I comment on the support the model provides for a Markov process of state transitions and examine the extent to which the Markov structure leads to changes in classification of the time period into collusive and noncooperative regimes. Looking at a number of different models, I then discuss the causes of price wars and whether the data support the predictions of the Green-Porter and Rotemberg-Saloner theories of cartel stability. Finally, it will certainly have occurred to the reader that the models employed here could be quite sensitive to functional form specifications, so I present some results on the robustness of the estimates.

3.7.1 The “standard” model

The results of what I shall call the “standard” Markov model are given in the first column of Table 3.2. The model differs from those of Porter (1983b) and Lee and Porter (1984) in two ways. First, I allow for serial correlation in the error term of the demand equation. Table 3.2 compares the results of the standard model to one without the ρU_{1t-1} term in the demand equation. Note that the coefficient on U_{1t-1} is

highly significant, and several other coefficients change as well. The most substantive is that the estimated price elasticity increases from 0.84 to 1.80 indicating that the JEC did indeed face an elastic demand curve. As a consequence, the estimated measure of collusion increases from 0.40 to 0.85, an indication that the JEC was more aggressive in its pursuit of the potential gains of collusion than is reported in Porter (1983b). While the *Lakes* effect seems to disappear, its place has been taken over by the dummy variables for the warm weather months.⁷

The other change incorporated into the standard model is the Markov nature of the regime shifts. As in Coslett and Lee (1985), we find that the regimes are not independently determined. The estimated probability of collusion in the standard model of Table 3.2 is $0.975 (e^{3.67}/(1 + e^{3.67}))$ after a collusive state and 0.067 after a price-war state. The expected duration of a price war is a reasonable 15 weeks. The primary effect of the change from independent to Markov regime shifts is on the classification of the states into collusive and non-cooperative regimes. When one is trying to determine the immediate causes of price wars and has only six to ten price wars to work with, the results can be very sensitive to changes in regime classifications. Table 3.3 gives three different series for the timing of the price wars in the JEC. The first is the maximum likelihood classification based on the standard Markov model. The second is from a model with identical structural parameters and independent regime shifts. (For this model, only the regime shift parameter was estimated because the maximum likelihood estimation algorithm on the full set of coefficients did not converge.) The third is that compiled by Ulen (1965) from contemporary reports in the *Railway Review*. The three series show six, nine, and eleven price wars. The timing of a given price war often differs. Note also that the timing may change again as we add parameters to model the causes of the price wars. This flexibility may allow the detection of causal connections which might be missed

⁷Unlike Coslett and Lee (1985), I do not allow for serial correlation in the residuals of the pricing equation. Such a specification seems unreasonable for a model of cartel behavior. The serial correlation they find reflects in part the discreteness of the JEC's price choices, not an AR(1) error process. When Coslett and Lee allow U_{2t} to follow an AR(1) process, they obtain a coefficient near unity on U_{2t-1} and most of the explanatory variables become insignificant. Such a model is unlikely to shed light on the issues I address.

Table 3.2: The "Standard" Model

Demand:				
$\log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 \text{Lakes}_t + \alpha_{3-14} \text{Seas}x_t + U_{1t}$				
Price:				
$\log P_t = \beta_0 + \beta_1 \log Q_t + \beta_2 I_t + \beta_{3-6} \text{DM}x_t + U_{2t}$				
Regimes:				
$\text{Prob}\{I_{t+1} = 1 I_t, Z_t\} = \frac{e^{\gamma W_t}}{(1+e^{\gamma W_t})}$				
"Standard" Model			No Serial Correlation	
Variable	Estimate	Std. Error	Estimate	Std. Error
Demand:				
<i>Constant</i>	7.677	1.882	9.019	0.361
<i>log P</i>	-1.802	1.287	-0.843	0.193
<i>Lakes</i>	-0.009	0.112	-0.460	0.348
<i>Seas1</i>	-0.103	0.086	-0.117	0.157
<i>Seas2</i>	0.146	0.145	0.167	0.180
<i>Seas3</i>	0.147	0.138	0.149	0.166
<i>Seas4</i>	-0.011	0.157	0.145	0.242
<i>Seas5</i>	-0.315	0.165	0.062	0.164
<i>Seas6</i>	-0.550	0.179	0.077	0.170
<i>Seas7</i>	-0.446	0.198	0.081	0.176
<i>Seas8</i>	-0.504	0.194	-1.116	0.374
<i>Seas9</i>	-0.395	0.165	0.048	0.185
<i>Seas10</i>	-0.545	0.164	0.102	0.191
<i>Seas11</i>	-0.521	0.180	0.085	0.304
<i>Seas12</i>	-0.397	0.173	0.183	0.241
Supply:				
<i>Constant</i>	-4.764	1.863	-5.649	9.461
<i>log Q</i>	0.306	0.178	0.398	0.928
<i>DM1</i>	-0.154	0.075	-0.211	0.124
<i>DM2</i>	-0.246	0.064	-0.283	0.160
<i>DM3</i>	-0.317	0.076	-0.373	0.242
<i>DM4</i>	-0.198	0.119	-0.419	0.422
<i>I_t</i>	0.637	0.104	0.660	0.406
Regimes:				
<i>Const.(I_t = 1)</i>	3.675	0.474	3.661	0.513
<i>Const.(I_t = 0)</i>	-2.641	0.404	-2.620	0.476
Other:				
σ_1	0.290	0.061	0.396	0.029
σ_{12}	-0.007	0.004	-0.045	0.142
σ_2	0.160	0.045	0.191	0.313
<i>U_{1t-1}</i>	0.832	0.085		

in a two stage estimation process if the first stage classification is incorrect.

Table 3.3: Timing of Price Wars

Predicted from Markov Model	Predicted from Switching Model	As reported in <i>Railway Review</i>
	1. 1881:10	1. 1881:10–1881:17
1. 1881:27–1882:10	2. 1881:26–1881:45	2. 1881:24–1882:4
	3. 1881:47–1882:9	3. 1883:8
		4. 1883:17–1883:18
		5. 1883:36–1883:37
2. 1884:13–1884:27	4. 1884:13–1884:27	6. 1884:7–1884:25
3. 1884:36–1884:41	5. 1884:36–1884:41	7. 1884:31–1884:40
		8. 1884:46
4. 1885:5–1885:6	6. 1885:4–1885:6	9. 1885:3–1885:5
5. 1885:12–1885:30	7. 1885:12–1885:18	
	8. 1885:20–1885:30	10. 1885:11–1885:47
6. 1885:34–1885:42	9. 1885:34–1885:41	
		11. 1886:3–1886:11

3.7.2 The Green-Porter and Rotemberg-Saloner models

At this point, we are prepared to discuss whether econometric models within the framework of Section 4 can provide empirical support for the theoretical models discussed in Section 3. This is a much more difficult problem, for while there are 328 weeks of data in the sample, there are only between six and eleven price wars. In a sense, there is really far less information on the causes of price wars, and we cannot expect to see the same significance levels as above.

Table 3.4 summarizes the results of seven models designed to test the predictions of the Green-Porter and Rotemberg-Saloner models. Each is a Markov model containing all of the variables of the “standard” Markov model. In addition, it is assumed that the probability of a transition from a collusive to a price-war state is not a constant, but rather is given by $\frac{1}{(1+e^{\gamma W_t})}$ where we use seven different sets of explanatory variables for W_t , and γ is a vector of unknown coefficients. The explanatory variables have been defined so that in each case a larger value is predicted by the theory to be more

likely to lead to a breakdown of collusion. The coefficient estimates then are expected to be negative.

Table 3.4: Causes of Price Wars*

Regimes:	$Prob\{I_{t+1} = 1 I_t = 1, Z_t\} = \frac{e^{\gamma W_t}}{(1+e^{\gamma W_t})}$						
	Model						
Parameter	1	2	3	4	5	6	7
<i>Constant</i>	6.33 (1.95)	4.63 (0.79)	5.56 (2.02)	5.61 (1.85)	6.43 (2.03)	4.37 (0.95)	4.93 (2.29)
<i>BigShare1</i>	-0.75 (0.48)	-0.78 (0.49)					
<i>BigShare2</i>			-0.34 (0.39)				
<i>BigShareQ</i>				-0.14 (1.06)			
V_{1t}					-4.27 (2.70)	-5.09 (3.70)	
<i>SmallShare</i>							0.86 (0.82)
<i>RS_Seasonal</i>	-1.60 (1.80)		-1.29 (1.83)	-1.65 (1.79)	-1.79 (1.76)		-2.05 (1.59)
U_{1t}		0.12 (2.09)				1.22 (1.13)	

* Estimated standard errors in parentheses.

As described earlier, we examine the Green-Porter model by testing whether the price wars are preceded by one of several possible triggers for a cartel with imperfect information. The first four models in the table focus on what should be the the strongest signal, that of one firm obtaining an unusually high demand relative to the others. The first two include the *BigShare1* variable created to test this mechanism. The two coefficient estimates are each negative as predicted, although each falls just short of being significant at the 5% level in a one tailed test. The next two columns of Table 3.4 give estimates of models with the two other variants on this variable, *BigShare2* and *BigShareQ*. The estimates are again negative, but are smaller and less significant.

An instructive question to ask is whether the observed effect is of a magni-

tude capable of supporting the collusive outcome. Absent cost data and a model of firm level demand, I give a back-of-the-envelope calculation. Suppose that there is very little differentiation so that firm i is able to lower its price by a negligible amount in period t and thereby increase its demand by 1%. The measure of demand, $\max_k(\log q_{kt} - \frac{1}{n} \sum_j \log q_{jt})/\sigma_k$, might then be expected to be increased by $\log(1.01)/\sigma_i$ whenever firm i would otherwise have had the highest market share. A simple computer simulation with reasonable parameter values gives the probability of a price war increasing by about 0.0001. If the losses from a price war are a little less than 15 times the per period profits, the future loss is about 0.1% or 0.2% of current profits. The estimated effect then is about an order of magnitude too weak to support collusion. Of course, if the firms had information not contained in our data, and/or the true trigger for price wars is only imperfectly correlated with *BigShare1*, the actual incentives would likely be much stronger than the effects we have identified. On the whole, the results resemble the predictions of the theory, but do not clearly establish the existence of a mechanism capable of sustaining collusion.

With differentiated demands and price as the strategic variable, secret price cuts cause not only a redistribution of market shares but also an increase in total sales. Hence, unexpectedly *high* demand could also be used as a trigger in a Green-Porter type cartel. Such a trigger might be particularly appealing if firms were concerned with preventing simultaneous deviations by more than one firm. The fifth and sixth models in Table 3.4 test whether large residuals V_{1t} in the aggregate demand equation trigger price wars. The coefficient estimates are large and negative as the theory predicts and again fall just short of significance at the 5% level.

I should note that there is also a potential difficulty in interpreting the results. We are faced with the same informational limitation as the JEC in that we do not know whether secret price cuts were given. Thus, the result may indicate not a Green-Porter mechanism with a price war triggered by a large random shock, but instead widespread price cutting by many firms (an undeclared price war) before the new lower price is officially acknowledged. I have already noted that there is ample anecdotal evidence of cheating in the JEC. The fourth chapter of this thesis applies

a more complicated hidden regime model and argues that price cutting may have been quite common. However, that conclusion is far from definitive, and is only likely to account for the large values of V_{1t} which precede price wars if in addition firms increased the size of their price cuts immediately before the price wars. In any case, we should keep in mind that errors in the timing of price wars are possible and may make it harder to interpret estimates of the immediate causes of price wars.

The final model in the Table 3.4 includes the *SmallShare* variable to test whether price wars are triggered by an unusually small demand for one firm. The coefficient estimate has the wrong sign and is not significant.

As described in section 3, the Rotemberg-Saloner theory predicts that the likelihood of a price war increases when current demand is high relative to expected future demand. I have already mentioned that we can identify two periods of increased incentives to deviate. First, the *RS_Seasonal* variable is used to test whether price wars reflect the seasonal incentive to deviate. The coefficient estimate is negative in each model although the results are not significant. While this provides little support for the Rotemberg-Saloner theory, an examination of Table 3.3 shows that in each of the last two years of the cartel price wars (those numbered 2. and 5. in the first column) did break out when the seasonal incentive was at its peak prior to the opening of the Great Lakes.

Second, given the observed serial correlation of the random component of demand, when U_{1t} is large, firms will expect demand to slowly decline over time. This again leads to an increased incentive to deviate. The coefficient estimates reported in models 2 and 6 in the table provide no evidence that these incentives are related to the onset of price wars in the JEC. I would like to emphasize that noise in the demand equation may lead to price wars via two distinct mechanisms. The Green-Porter theory holds that unanticipated increases, i.e. large values of V_{1t} , may be mistaken for secret price cuts. The Rotemberg-Saloner theory holds that price wars are more likely whenever demand is likely to decline in the future, and hence likely when U_{1t} is large. The results of model 6, which includes both V_{1t} and U_{1t} as explanatory variables, point to the presence only of the first effect. The absence of an effect here runs contrary to

the observations Rotemberg and Saloner make on the JEC.

3.7.3 Sensitivity of the results to functional form assumptions

As always, one wonders how the character of the results depends on the functional form assumptions. Compelling arguments can be made for many minor changes in the structural equations. Rather than discuss the results of several such changes, I examine one extreme change in the model and regard results which carry over as being fairly robust. The obvious choice is to estimate a model with linear demand. Porter (1983b) mentions that he performed this test and obtained reasonable results.

I assume that demand has the linear form

$$Q_t = \alpha_0 + \alpha_1 P_t + \alpha_2 Lakes_t + \alpha_{3-14} Seasx_t + \alpha_{15} Lakes_t P_t + U_{1t}$$

A monopolist with constant marginal cost c would respond by setting

$$P_t = c - \frac{1}{\alpha_1} Q_t - \left(\frac{1}{\alpha_1 + \alpha_{15}} - \frac{1}{\alpha_1} \right) Q_t Lakes_t$$

Hence, I estimate a model with demand as above and with supply

$$P_t = \beta_0 + \beta_1 Q_t I_t + \beta_2 Q_t Lakes_t I_t + \beta_{3-6} DMx_t + U_{2t}$$

where I_t is an indicator of collusion and the error structure is as before. The results of the model (with Q measured in 10000 tons) are presented in Table 3.5. Most of the coefficients have the expected sign. The *Price* term in the demand equation is negative and significant as are many of the dummy variables for the warmer months. The collusion term is positive and significant, and collusion has less of an effect on the price when the lakes are open.

Table 3.6 compares the standard Markov model to the linear model of Table 3.5. It is intended to give a rough idea of the effect of the change in specification on estimates

Table 3.5: The Linear Specification

Demand:		
$Q_t = \alpha_0 + \alpha_1 P_t + \alpha_2 Lakes_t + \alpha_{3-14} Seas_{xt} + \alpha_{15} Lakes_t P_t + U_{1t}$		
Price:		
$P_t = \beta_0 + \beta_1 Q_t I_t + \beta_2 Q_t Lakes_t I_t + \beta_{3-6} DM_{xt} + U_{2t}$		
Regimes:		
$Prob\{I_{t+1} = 1 I_t, Z_t\} = \frac{e^{\gamma W_t}}{(1 + e^{\gamma W_t})}$		
Variable	Estimate	Std. Error
Demand:		
<i>Constant</i>	5.932	1.789
<i>Price</i>	-11.133	7.012
<i>Lakes</i>	-0.363	1.048
<i>Seas1</i>	-0.182	0.218
<i>Seas2</i>	0.086	0.305
<i>Seas3</i>	0.088	0.330
<i>Seas4</i>	0.058	0.401
<i>Seas5</i>	-0.649	0.372
<i>Seas6</i>	-1.178	0.415
<i>Seas7</i>	-0.864	0.491
<i>Seas8</i>	-1.150	0.434
<i>Seas9</i>	-1.005	0.407
<i>Seas10</i>	-1.201	0.450
<i>Seas11</i>	-1.060	0.427
<i>Seas12</i>	-0.821	0.420
<i>Lakes P</i>	0.807	3.859
Supply:		
<i>Constant</i>	0.227	0.021
<i>Q I_t</i>	0.022	0.009
<i>Q I_t Lakes</i>	-0.008	0.006
<i>DM1</i>	-0.044	0.020
<i>DM2</i>	-0.063	0.018
<i>DM3</i>	-0.066	0.019
<i>DM4</i>	0.062	0.030
Regimes:		
<i>Const.(I_t = 1)</i>	3.542	0.427
<i>Const.(I_t = 0)</i>	-3.041	0.447
Other:		
σ_1	0.651	0.053
σ_{12}	-0.002	0.002
σ_2	0.033	0.004
<i>U_{1t-1}</i>	0.789	0.049

of the price elasticity and the degree of collusion. The price elasticity for the linear model is computed with Q and P set to their mean values. It is somewhat lower than the elasticity from the standard model. The “increase in price from collusion” is the increase in the predicted price of each model when all of the *Lakes*, *Seasax*, and *DMx* dummy variables are set to 0 and I_t switches from 0 to 1. It shows smaller price increases for the linear model. The linear model also gives a lower estimate of the measure θ of collusion (here defined as $-\alpha_1\beta_1$).

Table 3.6: Functional Form Comparison

	“Standard” Model	Linear Model
Price Elasticity of Demand	-1.802	-1.081
Increase in Price from Collusion	50.7%	26.2%
Measure θ of Collusion	0.849	0.242

The most relevant issue for our purposes is whether changes in the functional form of the model affect our conclusions on the causes of price wars. Conceivably, small changes in the structural model could change the classification of the price wars and thereby greatly alter our conclusions.⁸ Fortunately, this is not the case. Table 3.7 gives estimates of the logit parameters for regime transitions from models similar to those in Table 3.4, but with linear supply and demand equations. The results are very close to those in Table 3.4. We have roughly the same overall probability of collusion, and 19 of the 21 coefficients reported have the same sign under both specifications. The *BigShare1* variable loses some significance although the estimates do not change greatly. The V_{1t} variable becomes even more significant, while its overall

⁸As a separate test of the robustness of the results to errors in the timing of price wars, the logit models for the transition probabilities were estimated assuming Ulen’s (1978) contemporarily reported classification J_t to be correct (as in Porter (1985)). The results are given in Table 3.8 and are surprisingly consistent with the results presented earlier in this paper.

effect declines slightly. The other variables used to test the Green-Porter theory remain insignificant. Meanwhile, the primary variable used to test the Rotemberg-Saloner theory, *RS_Seasonal*, increases in both magnitude and significance under the linear specification.

Table 3.7: Price War Causes in the Linear Model*

Regimes:	$Prob\{I_{t+1} = 1 I_t = 1, Z_t\} = \frac{e^{\gamma W_t}}{(1 + e^{\gamma W_t})}$						
Parameter	Model						
	1	2	3	4	5	6	7
<i>Constant</i>	7.68 (2.64)	5.04 (1.17)	6.97 (2.34)	6.18 (2.28)	7.02 (2.19)	3.95 (0.56)	6.84 (2.33)
<i>BigShare1</i>	-0.69 (0.61)	-1.03 (0.64)					
<i>BigShare2</i>			-0.14 (0.43)				
<i>BigShareQ</i>				0.75 (0.89)			
<i>V_{1t}</i>					-1.19 (0.55)	-1.72 (0.61)	
<i>SmallShare</i>							0.07 (0.37)
<i>RS_Seasonal</i>	-2.94 (1.90)		-2.90 (1.75)	-3.11 (1.78)	-2.77 (1.83)		-3.04 (1.76)
<i>U_{1t}</i>		-0.66 (0.64)				0.47 (0.41)	

* Estimated standard errors in parentheses.

Overall, the change in specification seems to lessen the estimated extent of collusion, while leaving the structure of the price wars largely unchanged. It provides reassuring evidence of the robustness of our results, slightly bolstering the Rotemberg-Saloner theory.

3.8 Conclusion

In this paper, I have examined the structure of price wars in the Joint Executive Committee by estimating simultaneously the structural model of competition and a

Table 3.8: Price War Causes using Reported Regimes*

Regimes:	$Prob\{J_{t+1} = 1 J_t = 1, Z_t\} = \frac{e^{\gamma W_t}}{(1 + e^{\gamma W_t})}$			
	Model			
Parameter	1	2	3	4
<i>Constant</i>	5.20 (1.57)	4.62 (1.38)	4.79 (1.37)	4.46 (1.43)
<i>BigShare1</i>	-0.68 (0.48)			
<i>BigShare2</i>		-0.40 (0.44)		
<i>BigShareQ</i>			-0.29 (0.46)	
<i>SmallShare</i>				0.06 (0.39)
<i>RS_Seasonal</i>	-1.47 (1.31)	-1.17 (1.42)	-1.52 (1.28)	-1.60 (1.24)

* Estimated standard errors in parentheses.

dynamic model of behavioral regimes. The structural estimates are similar to those in previous papers, although I find a greater degree of collusion than is usually reported. The estimation of a dynamic model of behavior is made difficult by the small number of price wars we have to work with. Nonetheless, it is possible to formulate tests of the Green-Porter and Rotemberg-Saloner theories of price wars.

The Markov model yields reasonable transition probabilities and should have some flexibility in classification to help identify the causes of price wars. Models designed to test the Green-Porter theory are compatible with a Green-Porter type mechanism being used to sustain collusion. We find that two possible triggers for such a cartel, unusual market share patterns and unusually high aggregate demand, tend to precede price wars. Two caveats are, however, necessary. First, we have not identified a regime transition rule which is strong enough to enforce cartel discipline. Second, it is hard to tell whether we are observing a Green-Porter type mechanism or whether we are simply observing cheating by the firms. More reliable price data would clearly help answer this question.

The Rotemberg-Saloner view that price wars may result from unusually high short

term gains from cheating should be easier to investigate. Both seasonal patterns and the serial correlation of demand residuals create identifiable periods of increased incentives to deviate. The observed price wars are consistent with the prediction based on the seasonality of demand, although the results are not significant. We find no evidence for the more widely discussed implication that price wars are likely to occur during booms.

While the structural model involves a strong functional form assumption, the results on the causes of price wars are robust to a major change in the structural model.

References

Abreu, D., Pearce, D., and Stachetti, E., "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory* 39(1986), 251-269.

Berry, S., and Briggs, H., "A Non-parametric Test of a First-order Markov Process for Regimes in a Non-cooperatively Collusive Industry," *Economics Letters* 27(1988) 73-77.

Cosslett, S., and Lee, L., "Serial Correlation in Latent Discrete Variable Models," *Journal of Econometrics* 27(1985) 79-97.

Daily Commercial Bulletin. Chicago: Howard White & Co., 1880-1886.

Ellison, G., "What Does Not Seeing Something Look Like? Secret Price Cuts in the Joint Executive Committee," Mimeo, 1991.

Fudenberg, D., Levine, D., and Maskin, E., "The Folk Theorem with Imperfect Public Information," Mimeo, 1991.

Green, E., and Porter, R., "Noncooperative Collusion under Imperfect Price Competition," *Econometrica* 52(1984), 87-100.

Hajivassiliou, V., "Testing Game-theoretic Models of Price-fixing Behaviour," Cowles Foundation Discussion Paper No. 935, 1989.

Lee, L., and Porter, R., "Switching Regression Models with Imperfect Sample Separation," *Econometrica* 52(1984), 391-418.

MacAvoy, P., *The Economic Effects of Regulation*. Cambridge: MIT Press, 1965.

Porter, R., "Optimal Cartel Trigger Price Strategies," *Journal of Economic Theory* 29(1983a), 313–338.

Porter, R., "A Study of Cartel Stability: The Joint Executive Committee, 1880–1886," *Bell Journal of Economics* 14(1983b), 301–314.

Porter, R., "On the Incidence and Duration of Price Wars," *The Journal of Industrial Economics* 33(1985), 415–426.

Rotemberg, J., and Saloner, G., "A Supergame-Theoretic Model of Price Wars During Booms," *The American Economic Review* 76(1986), 390–407.

Ulen, T., *Cartels and Regulation*. Unpublished Ph.D. dissertation, Stanford University, 1978.

Chapter 4

What Does Not Seeing Something Look Like? Secret Price Cuts in the Joint Executive Committee

4.1 Introduction

The previous chapter argued that in some ways the Joint Executive Committee (JEC) resembles an optimally designed cartel. However, the conclusions were mixed in that the firms' responses did not appear to generate incentives strong enough to support a cooperative equilibrium. It is this mixed result which motivates the investigations of this essay. Certainly, the evidence of the previous chapter is not incompatible with the operation of an optimal cartel. For example, if the Joint Executive Committee relied on both price data and industrial espionage to determine whether to begin a price war, then the model of the previous chapter is misspecified. In analyzing triggers only weakly correlated with the true triggers, we would then expect to find a weak effect. On the other hand, it might also be that the Joint Executive Committee was not an example of an optimal cartel. In that case, firms may in fact have had an incentive to deviate from the agreements.

In an optimal Green-Porter cartel, no firm will deviate and price wars result only from random demand shocks. It is the fact that no firms should deviate which justifies

the use of list prices in previous analyses. On the other hand, if the cartel did not provide firms with sufficient incentives to cooperate, it is reasonable to suppose that some firms would have taken advantage of the mistake and offered secret price cuts. We would therefore like to know if such price cuts were offered. What would a secret price cut look like? If such a secret price cut were offered only once for one week, an econometrician, like the JEC members, would have no chance of finding it. On the other hand, if both periods of price cutting and periods of strict adherence to official prices were common, there is hope. When secret price cuts are given, demand will appear to be unusually high relative to the official prices. Models with hidden regimes might then be able to identify secret price cuts through periods of unusually high demand. The use of hidden regimes to identify periods of secret price cuts relies on heroic functional form and distributional assumptions. For this reason, I feel it necessary to devote a considerable part of this paper to discussing the properties of the regimes I estimate in order to say whether they can truly be described as periods of secret price cutting.

The first section of this paper discusses generally the use of models with hidden regimes. I argue that the interpretation of the results deserves more attention than it often receives and outline the types of descriptive evidence which are often available. The second section begins with the estimation of models of hidden regimes with Markov state transitions which show to a high degree of significance that demand is sometimes high and sometimes low. It continues with a discussion of the available descriptive evidence, and I conclude that one of the models may provide an estimate of the degree of price cutting, but that the results are not nearly as conclusive as they first appear.

4.2 Models of Unobserved Regimes

Rather than beginning right away with a discussion of secret price cuts, I first discuss the general problem of interpreting the results of models with unobserved regimes. Suppose we have hypothesized that a particular unobserved variable is an important

determinant of an economic system. An econometric study of the problem must address two problems. First, do models of hidden regimes identify an unobserved factor as being important? Second, if our estimates do show an unobserved factor to be significant, can we justifiably say that it is the particular factor we were looking for? Both questions must be answered affirmatively before we can claim to have found empirical support for our theory. I shall now discuss each question in a little more detail.

Certainly, an affirmative response to the first question is necessary. If we find no evidence of any hidden regimes, we cannot conclude that a particular unobserved factor is important. Suppose we believe the true model to be

$$Y = \beta X + \gamma Z + \epsilon$$

where X is a vector of observed exogenous variables, and Z is a vector of unobserved exogenous variables. If Z and ϵ have a multivariate normal distribution, there is no hope of identifying the effect of the unobserved variables. If we make other distributional assumptions, however, we may be able to estimate γ . For example, we may believe that ϵ is normally distributed but that Z is a discrete variable. This is the case in Porter's (1983) study of the JEC, where Z is an indicator for collusive or competitive pricing on the part of the firms. Of course, it is certainly possible that a cartel would choose a pricing strategy with more than two states. When we have no clear prediction for the distribution of Z , (as is the case for secret price cuts) we may nonetheless wish to try models which assume that Z may take on several discrete values in hopes that we will find values around which Z is concentrated.

An inherent danger in these models is that the models are identified by both functional form and distributional assumptions. If we assume normal errors, but the true model features non-normal errors we can find significant evidence for unobserved regimes when none actually exist. While this is obvious for a bimodal error distribution, it can also easily result if the distribution is uniform or has thick tails. In the data of this paper, the official prices set by the JEC are in fact discrete, so the

price equation clearly will not have normal errors. What appears to be an unobserved regime could easily be a reflection either of this non-normality or of other misspecifications. False results can also arise if there is simply a significant unobserved variable which we had not anticipated. In studying demand shifts in the JEC, many factors other than price cuts could be responsible. Among these are heavy snowstorms, mechanical problems, strikes, unusually fierce or light competition from steamships, and irregular departure schedules. There are simply too many alternative explanations for hidden regimes for us to talk of having identified a particular unobserved factor without giving corroborating evidence.

This leads us to the second question: how can we be sure that the unobserved factor we have found is the particular one we were looking for? While there has been little explicit discussion of this problem, most of the studies I cite have in fact dealt with this question fairly well. Given the unknown character of alternative explanations, answering this question is not usually just a process of testing competing theories. In building a description of the effect we have found, no single piece of evidence will be available to establish that we have found the effect we were looking for. On the other hand, many tests can prove that the effect we have found looks very different from and hence is probably not the one we were looking for. I try below to group the commonly used forms of evidence into three broad classes.

The first and most commonly used class of descriptive evidence is a simple comparison of the sign and magnitude of the parameter estimates to theoretical or other predictions. This, for example, is the main justification on which Dickens and Lang (1985) rely in their test of dual labor market theory. They first show that two wage equations fit better than one, then conclude that one equation represents the secondary sector because it applies to a small segment of the market in which wages are low and do not increase with education and experience. I would like to suggest though that with unobserved regimes such comparisons often are not very powerful. One common problem is the lack sharp predictions for the effects of variables we do not observe. In Porter's (1983) estimate of the degree of collusion in the JEC, any pricing regimes between which prices differ by less than the difference between

monopolistic and competitive prices could be attributed to collusion. In a model of secret price cuts, the sign of the coefficient for demand shifts can even be useless if we have no a priori idea of whether price cuts or strict adherence is more common. Further, such a comparison only increases our faith that we have found the effect we were looking for if the coefficient estimate we find is more probable for that effect than for alternative unobserved effects. Any coefficient attributable to secret price cuts might just as easily be attributed to effects of the weather. Of course, a completely unreasonable parameter estimate could tell us that we are not seeing the effects of secret price cuts.

A second main class of descriptive evidence consists of relationships between the estimated regime classification and other observed variables. Often, we have available fragmentary evidence related to the unobservable effect we have proposed.¹ Eberly (1990) gives a justification of this type in her study of automobile purchases. In arguing that credit constraints are responsible for the differences in behavior between the two groups she identifies she notes that those who are classified as credit constrained are far more likely to report having been denied credit and are less likely to obtain a car loan. If this were not true, we would hesitate to call the people credit constrained. Porter (1983) similarly notes a high correlation between his estimated regime classification and a series of contemporary reports of price wars. While contemporary commentary assists our attempt to name the unobserved regimes, we must recognize that the lay use of terms may not match the economic use. If a reporter, for example, simply writes that a price war is occurring whenever prices are low, he is providing us with no new information on the behavior of the firms. Ideally, we should hope for fragmentary evidence which is correlated with the proposed unobserved factor but uncorrelated with the observed variables which have been used to estimate the model.

A third class of descriptive evidence consists of the dynamic properties of the estimated regime classification series. Often, theory predicts that the regime changes due to a particular unobserved variable will follow a known dynamic pattern. For ex-

¹Lee and Porter (1984) discuss how such evidence can be directly incorporated into the model estimation.

ample, the Green-Porter theory predicts that the price wars in Porter's model should follow a pattern which looks like reasonably long price wars triggered by demand conditions which signal possible cheating by the firms. Note again that a single negative result here would make us doubt that the regimes represented price wars. If the model of the previous chapter had indicated that the low price regime prevailed in isolated randomly scattered weeks, we could not regard it to be the price wars predicted by Green and Porter (1984). When one is most interested in evidence for a particular unobserved factor, it seems prudent to use a two-step approach where the model is first estimated without dynamic restrictions and the resulting classification is then examined. If the dynamic properties of the hidden regimes are incorporated as in the previous chapter, the same parameters explain both the non-normalities in the residuals and the dynamic structure of the regime shifts. Interpretation is therefore more difficult.

The reasoning above serves as an outline for the following section in which I first try to identify a possible regime of secret price cuts and then examine descriptive evidence of each of the three types discussed above.

4.3 Secret Price Cuts in the JEC

Between 1880 and 1886, the JEC operated as a cartel to control rail rates between Chicago and the East Coast. Thorough descriptions can be found in MacAvoy (1965) and Ulen (1978). In this section, I ask whether firms secretly offered cut rates below the official rates for grain transport from Chicago to New York. That secret price cutting was possible and that some such price cuts were actually given is not really in doubt. The contemporary trade press contains many reports of such price cuts.² What I hope to add are estimates of both the frequency and the magnitude of the secret price cuts.

²The *Daily Commercial Bulletin* of June 16, 1881 contains a typical report, "Some agents claim to insist on 20c Grain to New York, while others state that 15c has no doubt been accepted." An article from February 2, 1886 expresses more confidence, "There is little doubt but rates are being cut on shipments of Grain to the East by one or two lines, ..."

4.3.1 Looking for price cuts via hidden regimes

As I mentioned in the introduction, when secret price cuts are given demand will appear unusually high relative to the official prices. As the theoretical model of Green and Porter (1984) predicts that no such price cuts should occur, we do not have a precise idea of what we are looking for. While price cuts could potentially be continuously distributed over a wide range, I look for the easiest thing to find, occasional price cuts all of the same magnitude. Such price cuts would give two demand regimes, with the demand curve shifted upward when secret price cuts are given. As the whole exercise is doomed if there is no evidence at all of demand shifts, I look for demand shifts first and worry later about whether they are caused by secret price cuts.

The previous chapter discusses a model of the JEC based on that of Porter (1983). It assumes that weekly demand is of the form

$$(1) \quad \log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 \text{Lakes}_t + \alpha_{3-14} \text{Seasxx}_t + U_{1t}$$

with *Lakes* an indicator for whether the Great Lakes are navigable by the competing steamships and the *Seasxx* a set of monthly dummy variables. The cartel's pricing decision is assumed to generate the supply curve

$$(2) \quad \log P_t = \beta_0 + \beta_1 \log Q_t + \beta_2 I_t + \beta_{3-6} \text{DMx}_t + U_{2t}$$

with the *DMx* dummy variables for changes in the cartel composition and *I* an indicator for collusive pricing. It is assumed that U_{1t} follows an AR(1) process,

$$(3) \quad U_{1t} = \rho U_{1t-1} + V_{1t} \quad |\rho| < 1.$$

and that V_{1t} and U_{2t} have a multivariate normal distribution.

I begin by motivating the search for secret price cuts with a simple illustration drawn from the model above. Taking the parameter estimates reported in the previous chapter, I construct the sequence of demand residuals \hat{V}_{1t} . Under correct specification,

the \hat{V}_{1t} should be normally distributed. The density estimate shown in Figure 4-1 was obtained from the residuals by kernel density estimation. I think the reader will agree that this picture strongly suggests that the residuals are not normally distributed and that there may be two or three distinct demand regimes each of which yields residuals centered near one of the peaks seen in the figure. The small peak on the right side of the figure seems like the best candidate to represent occasional periods of increased demand from unobserved price cuts.




Figure 4-1: Estimated Density of Demand Residuals

To formally explore the possibility of multiple demand regimes, I apply a model of hidden regimes with Markov transitions which is more complicated and somewhat more farfetched than that of the previous chapter. Essentially, I add indicator variables for one or two additional demand regimes to the demand relation (1). The model I refer to as Model 2 allows for two demand regimes so that the demand equation is

$$(4) \quad \log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 \text{Lakes}_t + \alpha_{3-14} \text{Seaspr}_t + \alpha_{15} \text{Regime1}_t + U_{1t}.$$

I intend *Regime1* to be an indicator for periods of price cuts, so I assume that this high demand state occurs with probability p_1 whenever firms are trying to collude, but never occurs during price wars when firms are already presumably pricing at marginal cost. I do not add any additional dynamic structure so that the demand regime which prevails at time t does not affect the likelihood of a price war at time $t+1$. I ignore this possibility for now in order to be sure that regime estimates reflect only unexplained shifts in demand, and not dynamic features like causes of price wars.

Model 3 allows for three demand regimes so that the demand equation becomes (5)

$$\log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 Lakes_t + \alpha_{3-14} Seasxx_t + \alpha_{15} Regime1_t + \alpha_{16} Regime2_t + U_{1t}.$$

Again I am envisioning that *Regime1* will be an indicator for a high demand regime due to secret price cuts. The lowest demand regime might represent a period of unusually low demand due to severe weather, labor problems, unusually low steamship rates, etc. As it is not crucial to the problem at hand, I do not try to provide an interpretation for this regime. In a collusive state, I assume that the high demand *Regime1* arises with probability p_1 , the medium demand *Regime2* with probability p_2 , and the third regime with probability $1 - (p_1 + p_2)$. I again assume that *Regime1* never arises during a price war and that the other two regimes occur with the same relative probabilities under collusive and price-war pricing.

Figures 4-2 and 4-3 illustrate the structure of regime transitions in Model 2 and Model 3. The large boxes represent the different regimes possible at time t . In Model 2, there are two regimes with collusive prices (the two possible demand regimes) and one with competitive prices. The arrows indicate probabilistic transitions and are labeled with the probabilities with which they are taken. For example, at the end of period $t-1$, we have a probabilistic transition which determines whether period t behavior will be described by one of the upper collusive states or by the lower price-war state. The transition to collusive pricing occurs with probability q_{12} after either collusive state and with probability q_0 after a price-war state. The probability q_{12} will be large and q_0 will be small. Within period t , another probabilistic transition

determines which of the collusive demand states arises. The high demand *Regime1* occurs with probability p_1 , and the low demand *Regime2* occurs with probability $1 - p_1$. Within each of the regimes, prices and quantities are determined by the supply and demand relations (2) and (4), with the appropriate values of the regime dummies, e.g. with $Regime1_t = 1$ and $I_t = 1$ within the uppermost period-t state in the figure. The diagram for Model 3 is similar, but with the five states representing the three possible demand regimes under collusive pricing and the two possible demand regimes in a price war.

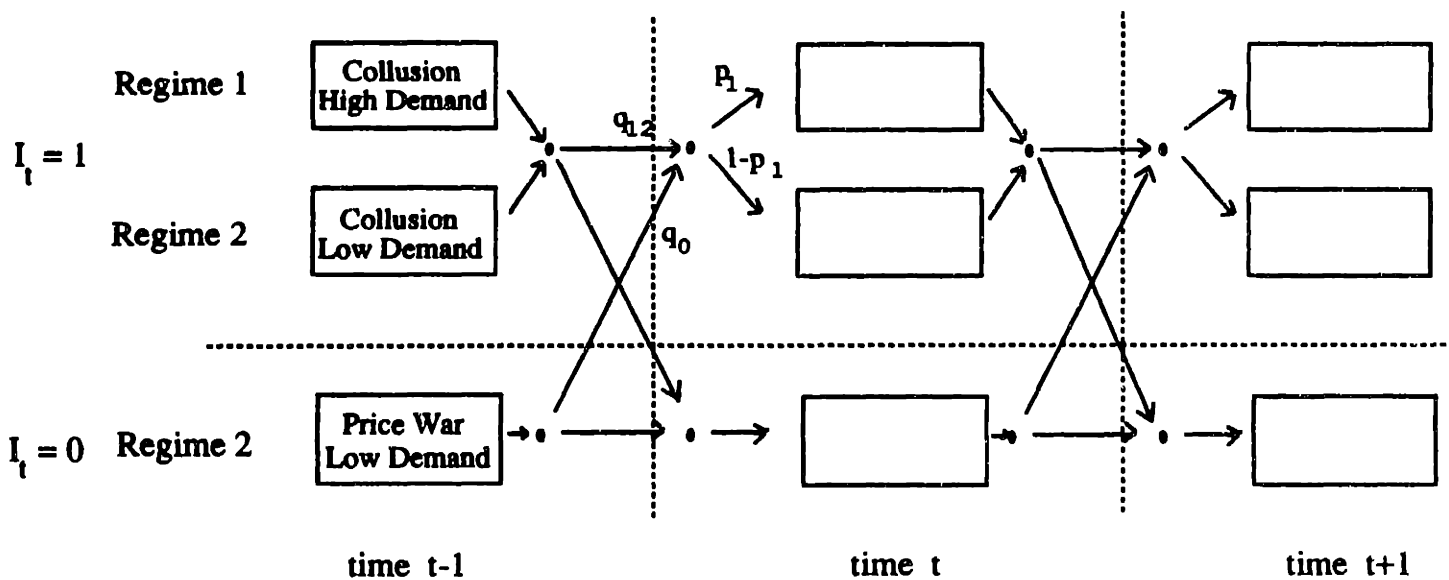


Figure 4-2: Regimes in Model 2

Table 4.1 gives the maximum likelihood coefficient estimates for Models 2 and 3. Most coefficient estimates are reasonable and quite similar to those reported in the previous chapter. Many of the key parameters, including the price elasticity of demand and the effect of collusion on prices are more significant, particularly in Model 3. Perhaps the only troubling change is that the estimated serial correlation of demand becomes quite close to one in Model 3. The most striking results is the extremely high degree of significance of the dummy variables *Regime1* for the hidden demand regimes we were looking for. The t-statistics on the *Regime1* parameter

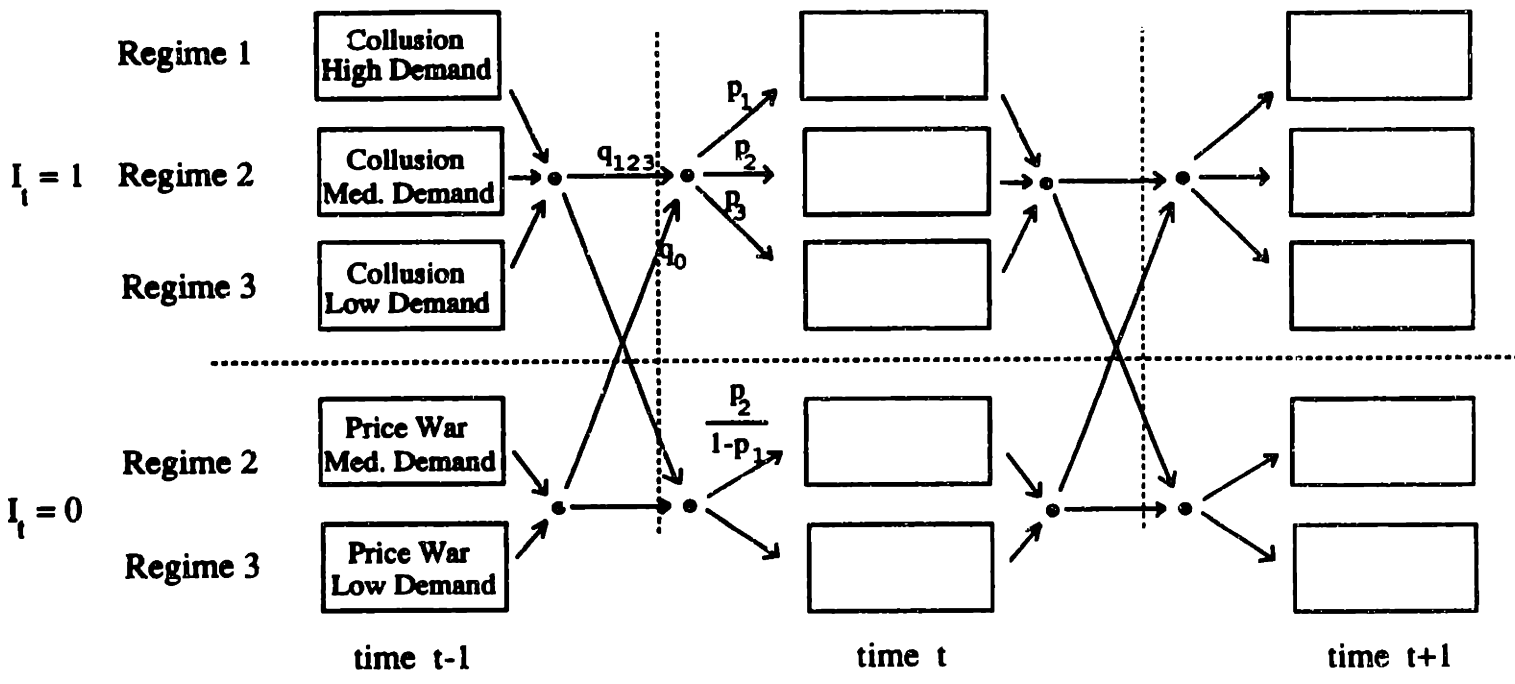


Figure 4-3: Regimes in Model 3

estimates, 12.1 and 26.7 for Models 2 and 3 respectively, are almost unreasonably high given a sample size of 328 weeks. In Model 2, the high demand *Regime1* has demand increased by about 55%. In Model 3, the high demand *Regime1* has demand 41% above the level of the medium demand *Regime2*, while the low demand regime has demand 32% below that of *Regime2*. From the standard errors of the parameter estimates, one is tempted to claim that each model has provided conclusive evidence for the existence of secret price cuts. Note, however, that the frequency p_1 with which the high demand regime occurs differs greatly between the two models. In Model 2, the main central peak of Figure 4-1 is seen as the result of the high demand regime. In Model 3, the smaller rightmost peak of Figure 4-1 becomes the high demand regime. If the reader was not already convinced, I hope that the contrast between the two models highlights the danger of jumping to the conclusion that a significant estimate for a hidden regime means that we have found the effect we were looking for. We must look far more closely at the available descriptive evidence before we can say whether the unobserved high demand regimes of either model can be said to have resulted from secret price cuts.

Table 4.1: Estimates of Hidden Regimes

Demand:				
$\log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 \text{Lakes}_t + \alpha_{3-14} \text{Seas}x_t + \alpha_{15} \text{Regime1}_t + \alpha_{16} \text{Regime2}_t + U_{1t}$				
Price:				
$\log P_t = \beta_0 + \beta_1 \log Q_t + \beta_2 I_t + \beta_{3-6} \text{DM}x_t + U_{2t}$				
Model 2			Model 3	
Variable	Estimate	Std. Error	Estimate	Std. Error
Demand:				
<i>Constant</i>	8.147	0.870	7.651	0.176
<i>log P</i>	-1.354	0.711	-1.587	0.092
<i>Lakes</i>	-0.222	0.073	-0.229	0.092
<i>Seas1</i>	-0.046	0.122	-0.114	0.087
<i>Seas2</i>	0.166	0.147	0.163	0.095
<i>Seas3</i>	0.016	0.115	0.106	0.108
<i>Seas4</i>	0.056	0.137	-0.041	0.122
<i>Seas5</i>	-0.225	0.169	-0.321	0.127
<i>Seas6</i>	-0.271	0.158	-0.374	0.162
<i>Seas7</i>	-0.163	0.194	-0.139	0.159
<i>Seas8</i>	-0.350	0.205	-0.332	0.157
<i>Seas9</i>	-0.179	0.182	-0.184	0.167
<i>Seas10</i>	-0.215	0.122	-0.219	0.148
<i>Seas11</i>	-0.194	0.121	-0.089	0.128
<i>Seas12</i>	-0.216	0.083	-0.091	0.114
<i>Regime1</i>	0.439	0.036	0.728	0.027
<i>Regime2</i>			0.384	0.022
Supply:				
<i>Constant</i>	-3.861	1.577	-4.101	0.412
<i>log Q</i>	0.220	0.153	0.242	0.038
<i>DM1</i>	-0.165	0.069	-0.162	0.072
<i>DM2</i>	-0.244	0.062	-0.239	0.068
<i>DM3</i>	-0.331	0.069	-0.323	0.063
<i>DM4</i>	-0.228	0.082	-0.198	0.119
<i>I_t</i>	0.595	0.079	0.608	0.047
Transition Probabilities:				
<i>q₁₂₃</i>	0.979	0.012	0.976	0.037
<i>q₀</i>	0.061	0.028	0.070	0.076
<i>p₁</i>	0.650	0.071	0.256	0.083
<i>p₂</i>			0.519	0.086
Other:				
<i>σ₁</i>	0.121	0.015	0.089	0.006
<i>σ₁₂</i>	-0.001	0.001	-0.0002	0.0007
<i>σ₂</i>	0.097	0.021	0.099	0.011
<i>U_{t-1}</i>	0.888	0.031	0.947	0.018

4.3.2 Do the estimated regimes look like secret price cuts?

In the following discussion, I apply each of the types of descriptive evidence mentioned in Section 2 to the problem of determining whether secret price cuts are responsible for the regimes estimated above. I begin simply by looking at the parameter estimates themselves, which suggest that we should focus our attention on the high demand regime of Model 3. A difficulty here is that we do not have sharp predictions for the magnitude and frequency of price cuts that may have been given. The Green-Porter theory which has guided the analysis so far predicts that there should be no price cuts. Perhaps the best estimates for the magnitude of potential price cuts are contemporary reports which indicate that a reasonable size might be between 10% and 25% of the official prices.³ The frequency of price cuts is harder to predict. While anecdotal evidence indicates that some price cuts were given, the fact that the cartel went as long as two years without a price war suggests that price cuts were not too common.

In Model 2, the high demand *Regime1* occurs with estimated probability $\hat{p}_1 = 0.65$ when firms are colluding. Given the estimated price elasticity of 1.35, a regime shift of the indicated magnitude would result from price cuts of between 25% and 30%. It seems unlikely that price cuts of this size could have occurred so often without clear historical evidence and more frequent price wars. The parameters estimated seem more likely to reflect instead occasional periods of low demand caused by some other unobserved variable. I think that the estimates themselves are sufficient evidence to conclude that the high demand regime of Model 2 does not reflect secret price cutting. Nonetheless, I continue to present evidence on Model 2 in the tables so that the reader may compare the descriptions of *Regime1* in the two models. Such a comparison suggests that the evidence I present is of some use in determining whether secret price cuts are indicated.

In Model 3, the high demand *Regime1* occurs with estimated probability $\hat{p}_1 = 0.26$. Given the estimated price elasticity of 1.59, the increase in demand from

³Some such reports appear in the *Daily Commercial Bulletin* of June 16, 1881, September 25, 1884, and January 22, 1885.

Regime2 to *Regime1* would require a price cut of about 20%. This certainly seems like a reasonable size and frequency for price cuts, although given our lack of precise predictions about price cutting and alternative explanations, we cannot say that the magnitude of the estimates makes us much more confident that we have found the effect of price cuts.

The second main class of evidence I mentioned in Section 2 is the relationship between the estimated regime classification and available evidence about the proposed explanation, secret price cuts. The Chicago Board of Trade's *Daily Commercial Bulletin* published market reports which contained both price quotations for rail freight transportation and occasional reports of rumored price cuts. A new price series P_{dcb} was constructed by averaging the daily prices listed there. When a range of prices were quoted or a rumored price cut specified, the lowest figure was used. Prices were assumed to be unchanged from the previous day on any day for which no quote appears. The variable P_{dcb} has a correlation of 0.95 with the official price P , being higher in 31 and lower in 50 of the 328 weeks. To identify periods of significant price cutting, the indicator variable *CutListed* was set to one whenever the quoted price P_{dcb} was at least three cents below the official price. A second indicator variable, *CutReported*, was set to one for any week in which an article stated that secret price cuts (often below the level reflected in P_{dcb}) were rumored to be given. Table 4.2 gives summary statistics for these variables. Keep in mind that we have ample reason to expect that each indicator only imperfectly indicates periods of price cutting. When only some shippers are being secret price cuts, the shippers may want to keep the price cuts secret in order to maintain a competitive advantage, so the press may have no reliable source of information.⁴ At other times, shippers might find it desirable to start rumors of price cutting when none is actually taking place in hopes of triggering a price war. While these indicators are only imperfectly correlated with price cutting, we have no reason to believe that hidden regimes attributable to other factors like the weather would be at all correlated with these indicators.

⁴Though not directly in connection to the JEC, Brown (1925, p. 176) discusses the "underhanded and evasive ways" in which price cuts could be given.

Table 4.2: Summary Statistics

Variable	Mean	Std. Dev.	Minimum	Maximum
P_{dcb}	0.2462	0.0674	0.0983	0.4000
$P - P_{dcb}$	0.0003	0.0212	-0.1000	0.1083
<i>CutListed</i>	0.0549	0.2277	0	1
<i>CutReported</i>	0.0366	0.1877	0	1

Given the parameter estimates for Models 2 and 3, we can construct the maximum likelihood classification of each week into one of the hidden regimes. Table 4.3 compares the conditional mean of each indicator variable given the estimated classification of demand regimes. Only periods with collusive pricing are included in the calculation. In the case of Model 3, expectations of the two indicators for price cutting increase from 0.040 and 0.024 to 0.048 and 0.063 when the firms are estimated to be in *Regime1*. While each indicator does appear to be positively correlated with the estimated high demand regime, these differences are not nearly as large as we might have wished. As would probably be expected given the small sample sizes, the differences also are not statistically significant. The fragmentary evidence supports only very weakly that *Regime1* of Model 3 may be measuring secret price cuts.

Table 4.3: Historical Evidence

Regime	Conditional Expectation $E(\cdot \text{Regime})$		
	<i>CutListed</i>	<i>CutReported</i>	# obs
Model 2 <i>Regime1</i>	0.050	0.031	(n=161)
<i>Regime2</i>	0.024	0.060	(n=83)
Model 3 <i>Regime1</i>	0.048	0.063	(n=63)
<i>Regime2</i>	0.040	0.024	(n=125)
<i>Regime3</i>	0.036	0.036	(n=56)

The final class of evidence I consider consists of the dynamic properties of the estimated regimes. The most obvious dynamic property of price cuts is that they should tend to cause price wars. We know that the JEC specifically empowered its

commissioner to declare price wars when price cutting occurred.⁵ To test whether any of the regimes we have found tend to cause price wars, I estimate two new models which I shall refer to as Model 2A and Model 3A. These models are identical to Models 2 and 3, except that I exploit more fully the possibility of Markov transition probabilities between the states. The probability of a price war at time $t+1$ is now allowed to depend on the demand regime which prevailed at time t . Secret price cuts should make price wars more likely. On the other hand, any shifts in the demand curve due to variables unknown to us but understood and observed by the JEC should not cause price wars. (Irregularities in shipping schedules might be one such factor.) Hence, this test should allow us to distinguish between price cuts and at least some alternate explanations for the high demand regimes.

Figures 4-4 and 4-5 illustrate the structure of the transition probabilities in the new models. The difference compared to Figures 4-2 and 4-3 is that the probability that a state with collusive pricing occurs at time $t+1$ is now q_1 after *Regime1* prevails at time t , q_2 after *Regime2*, and q_3 after *Regime3*. If the high demand *Regime1* of either model reflects a period of price cutting, we would expect q_1 to be smaller than either q_2 or q_3 .

To examine the properties of the previously identified regimes, only the transition probabilities q_1 , q_2 , q_3 , and q_0 are estimated with all other parameters held fixed at the values shown in Table 4.1. Table 4.4 gives the maximum likelihood parameter estimates for Models 2A and 3A. The estimates of Model 3A show that collusion follows the high demand *Regime1* with probability $1 - \hat{q}_1 = 0.064$ compared with 0.015 after *Regime2* and 0.00 after *Regime3*. This is not as strong an effect as we might have liked, but again is statistically significant and is consistent with the idea that the high demand *Regime1* of Model 3 is reflecting at least partially the effect of secret price cuts.

Dynamic evidence can also be usefully applied to look for properties we do not expect to find if the high demand regime reflects secret price cuts. One potential

⁵Daggett (1908, p. 6) states "By agreement of March 11, 1881, the chairman of the JEC, Mr. Fink, was given authority to proclaim a general reduction in published rates when it should be shown that any pool line had been accepting traffic at less than the regular rate."

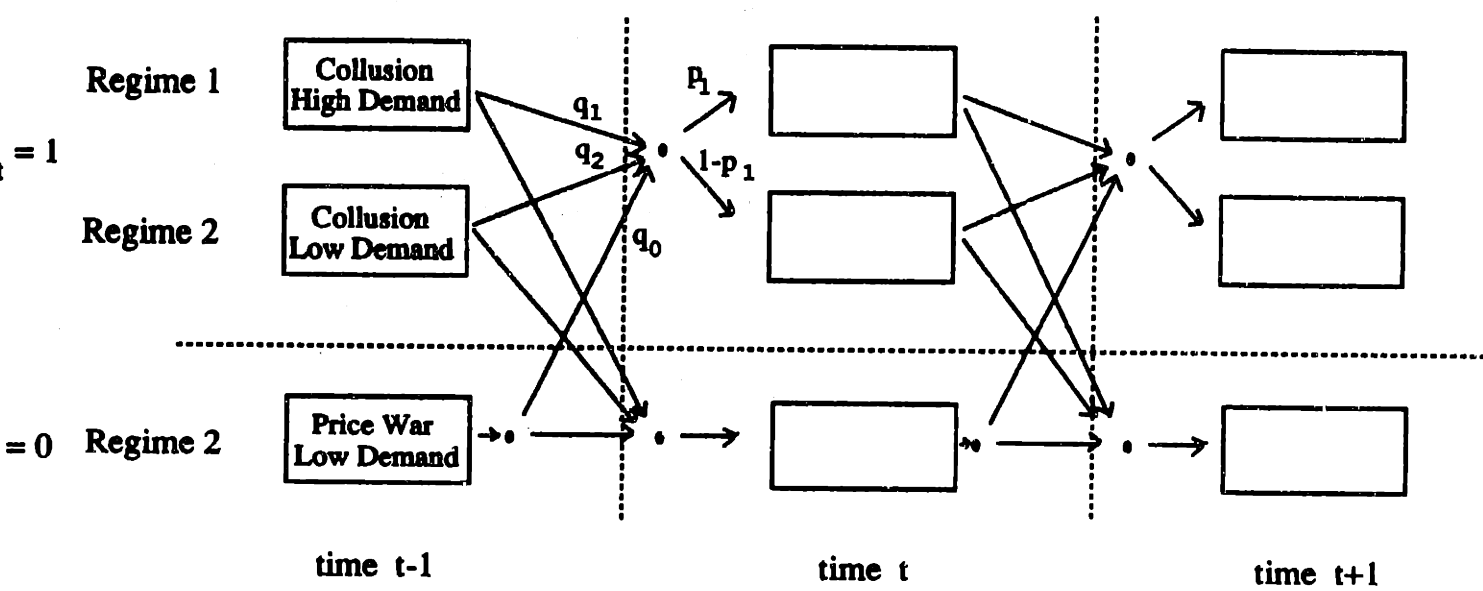


Figure 4-4: Regimes in Model 2A

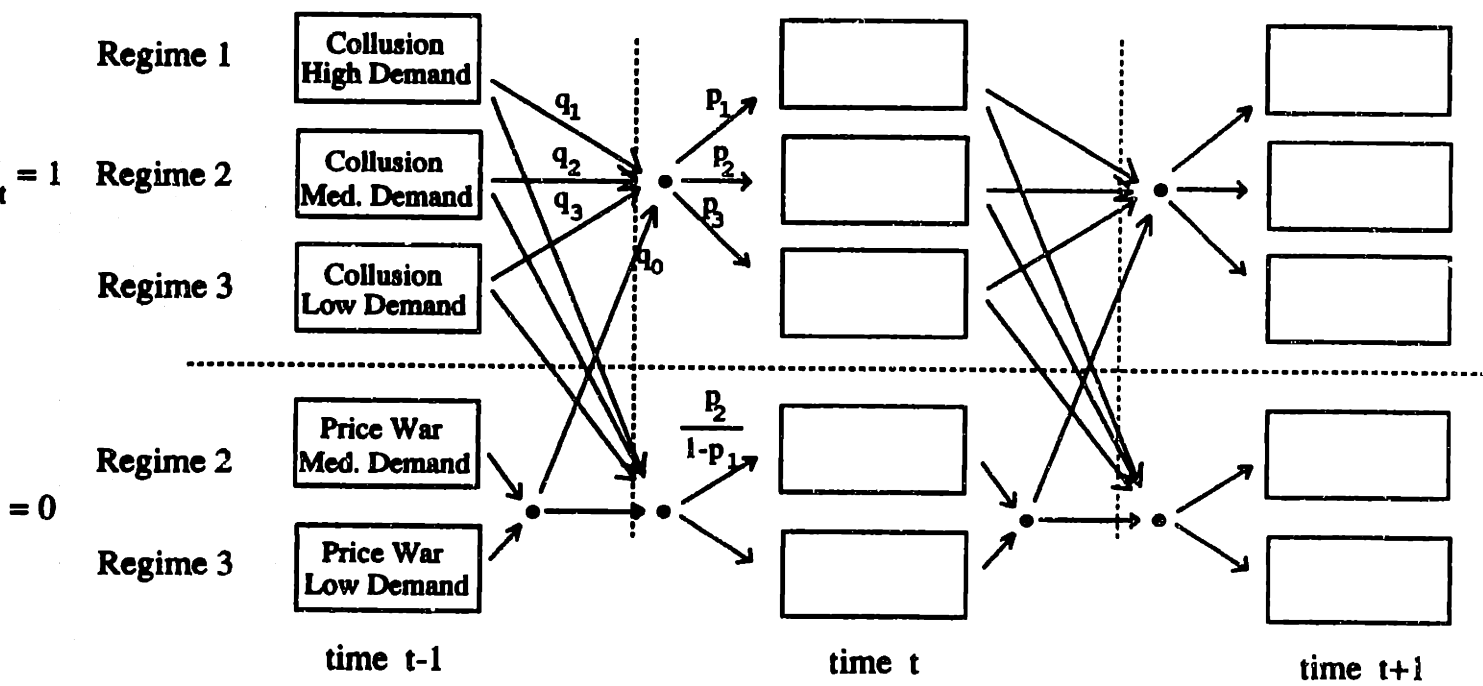


Figure 4-5: Regimes in Model 3A

Table 4.4: Dynamic Properties of Hidden Regimes

Probability of collusion following each regime				
Variable	Model 2A		Model 3A	
	Estimate	Std. Error	Estimate	Std. Error
q_1	0.968	0.015	0.936	0.032
q_2	1.000	—	0.985	0.011
q_3			1.000	—
q_0	0.061	0.007	0.070	0.027

alternate explanation for the demand regimes is that they might reflect seasonal variation in demand not captured by the monthly dummy variables I have used. If no seasonal pattern is apparent in the regime classification this competing explanation is discredited. (Many findings might be consistent with secret price cuts. A seasonal pattern could result from the seasonal incentives to cheat discussed in chapter 3.) Table 4.5 examines a very simple indicator for seasonal patterns, the relationship between the demand regimes in weeks which are exactly one year apart. For each possible demand regime at time t , the table gives the frequency with which *Regime1* occurs at time $t+52$ in the estimated regime classification series. Only periods for which both t and $t+52$ involve collusive pricing are included in the calculation. The estimates for Model 3 show all probabilities to be fairly close, and we can not reject that all probabilities are equal. By ruling out one alternative explanation, this result again increases our confidence that the regime reflects secret price cuts.

Table 4.5: Seasonal Pattern of the Estimated Regimes

Frequency of <i>Regime1</i> at time t				
$Regime_{t-52}$	Model 2		Model 3	
	$E(Regime1_t)$	# obs	$E(Regime1_t)$	# obs
<i>Regime1</i>	0.705	(n=105)	0.326	(n=43)
<i>Regime2</i>	0.533	(n=45)	0.324	(n=74)
<i>Regime3</i>			0.242	(n=25)

4.4 Conclusion

In this chapter, I have applied models of hidden regimes to look for periods of unusually high demand, hoping that these regimes might reflect and hence provide estimates of the degree of secret price cutting in the JEC. I find that hidden demand regimes are very strongly indicated. However, it is much harder to determine whether the hidden regimes which have been identified can properly be attributed to the proposed explanation, secret price cutting. The results are not nearly as clear cut as the first t-statistics suggest.

One of the models identifies a high demand regime which fails to exhibit any properties which might convince us that it represents periods of price cutting. A second high demand regime tends to lead to price wars, and may be correlated with the historical evidence on price cutting. Given these properties, I very tentatively conclude that this regime estimate reflects at least in part the effects of secret price cuts. If so, price cuts of about 20% were occasionally given by the JEC. It is harder to be confident about the frequency of price cuts, because I certainly cannot rule out the possibility that other unobserved factors are also causing similar demand increases and might be contributing to the estimated frequency of the high demand regime. Nonetheless, the estimated frequency of this regime, 26% of the collusive periods, is at least an upper bound on the frequency of secret price cutting if not a tentative estimate.

The previous chapter indicated that the apparent causes of price wars in the JEC were fairly consistent with the predictions of the Green-Porter theory. The theory further predicts that firms do not offer secret price cuts. To the extent that this prediction is not supported by the historical data, the JEC may be an example of a cartel whose design failed to provide firms with the incentives necessary to guarantee cooperation.

References

- Brown, H., *Transportation Rates and Their Regulation*. New York: MacMillan, 1925.
- Daggett, S., *Railroad Reorganization*. New York: Houghton, Mifflin & Co., 1908.
- Daily Commercial Bulletin*. Chicago: Howard, White & Co., 1880–1886.
- Dickens, W., and Lang, K., “A Test of Dual Labor Market Theory,” *The American Economic Review* 75(1985), 792–805.
- Eberly, J., “Adjustment of Consumers Durable Stocks: Evidence from Automobile Purchases,” Mimeo, MIT, 1990.
- Ellison, G., “Markov Models of Price Wars in the Joint Executive Committee,” Mimeo, MIT, 1991.
- Green, E., and Porter, R., “Noncooperative Collusion under Imperfect Price Competition,” *Econometrica* 52(1984), 87–100.
- Lee, L., and Porter, R., “Switching Regression Models with Imperfect Sample Separation,” *Econometrica* 52(1984), 391–418.
- MacAvoy, P., *The Economic Effects of Regulation*. Cambridge: MIT Press, 1965.
- Porter, R., “A Study of Cartel Stability: The Joint Executive Committee, 1880–1886,” *Bell Journal of Economics* 14(1983), 301–314.

Ulen, T., *Cartels and Regulation*. Unpublished Ph.D. dissertation, Stanford University, 1978.