

Comparing the Robustness of Trading Systems to Higher Order Uncertainty*

Hyun Song Shin,

Department of Economics, University of Southampton,
Highfield, Southampton, SO17 1BJ, U.K.

Revised, July 1995

Abstract: This paper compares the performance of a decentralized market with that of a dealership market when traders have differential information. Trade occurs as a result of equilibrium actions in a Bayesian game, where uncertainty is captured by a finite state space and information is represented by partitions on this space. In the benchmark case of trade with common knowledge of endowments, the two mechanisms deliver virtually identical outcomes. However, with differential information, the dealership market has strictly higher trading volume, and yields an efficient post-trade allocation in most states. In contrast, the decentralized market suffers from suboptimal trading volume. The reason for this poor performance is the vulnerability of the decentralized market to higher order uncertainty concerning the fundamentals of the market. Traders may know that mutually beneficial trade is feasible, and perhaps know that they know, and yet a failure of common knowledge that this is so precludes efficient trade. The dealership market is robust to this type of uncertainty.

* I would like to record my debt to Stephen Morris for shaping my views on the issues raised here. Ian Jewitt, as managing editor, and three referees guided this paper through several revisions, and I am grateful to them for their perseverance. I am especially grateful to one of the referees for pointing to the importance of limit orders in modifying the results reported here. I have gained from the comments of Helmut Bester, Katie Graddy, Ruben Lee, Ailsa Röell, Daniel Seidmann, Peter Sinclair, John Vickers and seminar participants at Tilburg, Newcastle, CORE, Cambridge, IFS, Southampton and the 1994 European Summer Symposium in Financial Markets at the Studienzentrum Gerzensee. An earlier version of this paper was circulated under the title, "The Volume of Trade in Auction and Dealership Markets with Differential Information".

1. Introduction

This paper is concerned with how trading systems cope with differential information among traders. Traditionally, analysis of trade with differential information has been underpinned by the assumption of price-taking behaviour, beginning with Radner's (1968) analysis of competitive equilibrium with differential information, developing into the now familiar and standard notion of rational expectations equilibrium (Radner (1979), Allen (1981, 1982), Jordan (1982)). This concept has set the standard in the study of differential information, and has been the mainstay of applied research, especially after the introduction of the notion of 'noisy' rational expectations (Grossman and Stiglitz (1980), Hellwig (1980)).

And yet, in spite of its pre-eminent place in applied research, comparatively little is known in terms of the trading institutions which lie behind the analysis. For instance, fully revealing rational expectations equilibrium may involve demand functions which are independent of price, and as such is difficult to view as the outcome of a game in which traders submit demand functions. 'Noisy' rational expectations equilibrium is not prone to this particular weakness, but invites other questions. Dubey, Geanakoplos and Shubik (1987) note that games where traders submit demand functions do not specify a unique outcome when there are several market clearing prices, leading to possibly perverse situations where there is no continuous mechanism which can implement the REE correspondence. The growing literature on market microstructure has been motivated partly by the perception of this gap between allocations and institutions, although the development of this literature has followed rather different lines, which is not always well-suited for the analysis of differential information when the traders' information cannot be ordered.

The point of departure of this paper is the trading institutions themselves. Rather than starting with post-trade allocations and asking which trading institutions implement them, we will begin by describing the institutions themselves, and comparing the allocations which they bring about in equilibrium. The institutions in this paper are extremely simple

and stylized, and are designed to accentuate the differences between a dealership market and a decentralized market. In the former, some of the traders are designated as price-setters, who act in anticipation of the buyers' actions. In the latter, traders play a Shapley-Shubik (1977) market game in which they submit, simultaneously, quantity orders to an auctioneer, who then sets price to clear the market¹.

In comparing these institutions, we will focus on one consequence of differential information: namely, how trading institutions fare in the face of higher order uncertainty. We envisage scenarios where all traders know that mutually beneficial trade is feasible, and perhaps know that they all know it, and yet there is a failure of common knowledge that this is the case. Failures of common knowledge will be the rule rather than the exception in economies with differential information, and yet answers to such questions may elude those working with very general and complex models motivated by other questions.

To see why a failure of common knowledge may be important, it is instructive to contrast the role of uncertainty in strategic situations from that in single person decision problems. In a single person decision problem, payoffs are determined by one's action and the state of the world. When a decision maker receives a message which rules out some states, this information can be utilized directly by disregarding these states in one's deliberations. However, the same is not true in strategic situations in which the payoff of an agent depends on the actions of other agents, as well as the state of the world. Since my payoff depends on your actions and your actions are motivated by your beliefs, I care about the range of possible beliefs you may hold. So, when I receive a message which rules out some states of the world, it may not be possible to disregard these states in my deliberations, since some of these states may carry information concerning *your* beliefs. Furthermore, your beliefs at these neighbouring states may depend on your beliefs concerning *my* beliefs at a further set of states. The reasoning does not stop here. Unless there is some feature of the situation which curtails this sequence of iterated beliefs, higher

¹ The Shapley-Shubik market game may not be a good representation of markets in which traders have access to general limit orders. I am grateful to a referee for a detailed discussion of this issue, and of possible avenues for future research.

order beliefs of all orders will be relevant for my decision now. Thus, in contrast to the single person decision problem, I may be forced to take into consideration states which are *known* not to have occurred.

This last observation is not, of course, original to this paper. Rubinstein (1989) and Monderer and Samet (1989) are early expositions of the theme, and Carlsson and van Damme (1993), Sorin (1993), and Morris, Rob and Shin (1995), have proposed criteria for deciding when such higher order uncertainty may impinge on equilibrium for 2×2 games, or for two player games with finite action sets. Rather, the main contribution of this paper is to offer a set of tools which may be employed in analysing the effects of differential information in more conventional market settings where there is price-mediated trade among many traders.

When unrealized states exercise an influence on the equilibrium allocation, there are unavoidable welfare consequences. Since the optimality or otherwise of the final allocation hinges on what the *fundamentals* are (in terms of preferences and endowments), a well-functioning trading system is one which ensures that post-trade allocations are determined in the appropriate way in relation to the fundamentals. However, when unrealized states exert an influence on the final allocation, this link from fundamentals to the final outcome is subject to interference. The mark of a well-functioning trading system is one which minimizes such interference, and which ensures that the final allocation is as close as possible to that justified by the fundamentals. In contrast, a poorly performing trading system is unable to insulate the equilibrium outcome from the influence of unrealized states. This, in a nutshell, is the theme of this paper.

In order to isolate the key effects, we conduct the analysis in three distinct steps. We first set the stage by describing the fundamentals of an economy with two goods. This allows us to say which allocations are efficient given the fundamentals, and gives us a welfare benchmark against which we may judge the performance of any trading institution.

We then introduce two distinct Bayesian games, whose contrasting rules reflect the contrast between a dealership market and a decentralized, order-driven market. In a preliminary discussion, it is shown that the two institutions deliver virtually identical outcomes when there is no differential information. However, with the introduction of differential information, the performance of the two institutions diverge sharply. The dealership market has strictly higher trading volume, and delivers the Pareto efficient allocation in most states. In contrast, the decentralized market (employing the Shapley-Shubik rules) suffers from low trading volume, and the post-trade allocation is bounded away from the efficient allocation everywhere on the state space. Moreover, the extent to which the trading volume in the decentralized market falls below the efficient level can be large, depending on a parameter which captures the scope for mutually beneficial trade. The apparent fragility of the Shapley-Shubik market to departures from common knowledge can be traced to the large effect of unrealized states on the equilibrium allocation.

The juxtaposition of a decentralized market with that of a dealership market has a parallel in the debate in the literature on the market microstructure of financial markets (for example, Madhavan (1992), Pagano and Röell (1992a, 1992b, 1993) and Biais (1993)). However, trading rules in financial markets are considerably more sophisticated than those examined in this paper, and the analysis offered here ought to be viewed as preliminary observations in a larger research programme whose ultimate goal is to set out more systematically the welfare consequences of alternative trading arrangements when traders face differential information. For our part, let us begin by describing the fundamentals underlying our markets.

2. The Model

Our model is an account of trade between two tribes. One tribe lives by the sea, the other in the mountains. Both tribes grow and consume rice. However, the coastal tribe has been open to the influence of other cultures, and has acquired a taste for rice pudding. Indeed, members of the coastal tribe regard rice and rice pudding as perfect substitutes. Rice

pudding is produced by cooking rice in yaks' milk. Unfortunately for the coastal tribe, yaks find the coastal climate inhospitable so that rice pudding cannot be produced on the coast. However, yaks are plentiful in the mountains where the mountain tribe lives. Members of the mountain tribe place zero value on the consumption of rice pudding, but they welcome the opportunity to trade rice pudding for rice.

There are n rice growers in the coastal tribe and n rice growers in the mountain tribe, where $n \geq 2$. Each rice grower in the mountain tribe has access to a technology which converts rice into rice pudding. If we denote the quantity of rice by x and of rice pudding by y , then the cost (in units of rice) of producing y units of rice pudding is given by :

$$x = y^\beta / \beta, \quad (2.1)$$

where $\beta > 1$. The utility function of a mountain rice grower is $u^M(x, y) = x$, while the utility function of a coastal rice grower is $u^C(x, y) = x + y$. The scope for trade arises from the fact that the coastal growers place a value of 1 on rice pudding (in terms of rice) while the marginal cost for a mountain grower of producing rice pudding is less than one, for $y < 1$.

The rice harvests in the two regions do not suffer from much variability. The crucial factor in the rice harvest in both regions is whether there is any rain in the early growing season. The early growing season lasts for exactly N days. If there is no rain during this period, the rice harvest fails, and yields a harvest of zero. If, however, there are one or more days of rain during this period, the rice harvest yields one unit of rice for every grower in that region. Moreover, the geographical proximity of the two regions means that the number of days of rain in the two regions are highly correlated. The probability of there being q days of rain on the coast and r days of rain in the mountains is zero unless $q = r$ or $q = r - 1$. Furthermore, the probability distribution is uniform so that if we denote by $\text{prob}(q, r)$ the probability of there being q days of rain on the coast and r days of rain in the mountains, we have :

$$\text{prob}(q, r) = \begin{cases} 1/(2N+1) & \text{if } q = r \text{ or } q = r - 1 \\ 0 & \text{otherwise} \end{cases} \quad (2.2)$$

where q and r range over $\{0, 1, 2, \dots, N\}$. Each tribe only observes the number of days of rain in its own region, and must make inferences concerning the rainfall (and hence the harvest) of the other tribe from this distribution.

Having described the preferences, endowments and information of the traders, we now describe two alternative trading mechanisms. The first is an order-driven market in the manner of Shapley and Shubik (1977) in which traders submit quantities to an auctioneer, who then sets the price to equate demand and supply. The second is a dealership market in which the sellers post prices in anticipation of buyers' demands.

Decentralized Market. Every year after the rice harvest, the two tribes have an opportunity to trade rice pudding for rice on a specified day. On the day before the market day, all rice growers from both regions take an action. The mountain rice grower decides how much rice pudding should be produced and taken to the market. In the meanwhile, every coastal grower submits a bid to an auctioneer for rice pudding. The coastal grower submits a bid by putting up a quantity of rice to be exchanged for rice pudding. The task of the auctioneer is to set the price of rice pudding so as to equate demand and supply. Denote by s_i the amount of rice pudding supplied by the i th mountain grower, and denote by d_i the amount of rice put up for exchange by the i th coastal grower. Given the vector of actions $(d_1, d_2, \dots, d_n, s_1, s_2, \dots, s_n)$, the allocation rule followed by the auctioneer on the day of the market is as follows. If $\sum_{j=1}^n s_j > 0$ and $\sum_{j=1}^n d_j > 0$, the auctioneer sets the price of rice pudding (in units of rice) to be :

$$p = \sum_{j=1}^n d_j / \sum_{j=1}^n s_j. \quad (2.3)$$

Then, the i th seller (mountain grower) gets ps_i units of rice in exchange for s_i units of rice pudding, and the i th buyer (coastal grower) gets d_i/p units of rice pudding in exchange for d_i units of rice. However, if either $\sum_{j=1}^n s_j = 0$ or $\sum_{j=1}^n d_j = 0$, then no trade takes place, in which case all buyers consume their initial endowment of rice while the i th seller consumes his initial endowment of rice minus the amount sunk in the production of rice pudding (s_i^β / β).

Given this allocation rule, we can describe the order-driven market as an extensive form game. The game has three stages. In the first stage, Nature chooses the rainfall in the two regions according to the probability distribution given by (2.2). In the second stage, all traders observe their own rainfall and form beliefs on the rainfall (and hence the harvest) of their trading partners. Each trader then takes an action (s_i by the seller, d_i by the buyer). In the third stage, the price of rice pudding is set to clear the market, and all players receive their allocations. The game ends when all traders consume their allocations. The payoff to consumption for the coastal farmer is given by the sum of the quantities of rice and rice pudding, while the payoff of the mountain farmer is just the quantity of rice.

A strategy for the i th buyer (coastal farmer) in this game is a rule which sets d_i as a function of the number of days of rain on the coast. A strategy for the i th seller is a rule which sets s_i as a function of the number of days of rain in the mountains. An equilibrium of the order-driven market is a vector of strategies (one for each trader) such that the action prescribed by a player's strategy given his signal maximizes that player's conditional expected payoff from consumption given that all other traders follow their respective rules. This is the standard notion of Bayesian equilibrium, due to Harsanyi (1967).

Dealership Market. In the dealership market, the producers of rice pudding post prices. On the day before market day, each seller decides how much rice pudding should be produced out of his endowment of rice, and decides on the price to be posted on the following day. On market day, each buyer is allocated a place in a queue. The buyers then trade in sequence. Each buyer observes all the prices posted, and chooses the group of sellers with whom to trade and ranks the sellers in order of preference. If one seller runs out of rice pudding, the buyer is permitted to visit the next seller in the sequence, until demand is satisfied.

As with the order-driven market, the dealership market can be formalized as an extensive form game with three stages. In the first stage, Nature chooses the rainfall in the two regions. In the second stage, the mountain growers observe the rainfall in their own region and form beliefs on the rainfall on the coast. Based on this information, they decide how

much rice pudding to produce and the price to be posted. In the third stage, the buyers observe the actions of all sellers, and trade in sequence. A strategy for a seller is a rule which sets the price-quantity pair for rice pudding as a function of number of days of rain in the mountains. A strategy for a buyer is a rule which selects those sellers with whom to trade, and which ranks them in order of preference, having observed the actions of all sellers. The game ends when all trades are completed and traders consume their final allocations.

Given their preferences, the buyers follow a simple rule of first visiting the seller with the lowest price, provided that the terms of trade are better than one to one. If this seller cannot satisfy the demand for rice pudding, the buyer moves on to the seller with the next highest price, and so on. We will assume that if a buyer is indifferent between trading and not trading, he will choose to trade. Hence, under our assumption, a buyer will trade provided the price is one or less.

This completes the description of the two trading mechanisms. Before we analyse the outcomes under these alternatives, we will digress briefly to compare the outcomes in the two markets when there is common knowledge of the endowments of the traders. This will serve as a benchmark for the main results of our paper which pertain to the case with differential information.

3. Outcomes under Common Knowledge

Suppose that the endowment of rice of both types of rice growers is 1, and that this is common knowledge. The equilibria in the two trading games can be analysed thus.

Order-Driven Market. Consider the problem for the i th seller. This trader has a unit endowment of rice and chooses s_i to maximize consumption of rice. Production of s_i units of rice pudding costs s_i^β/β units of rice but the trader obtains ps_i units of rice in exchange. Thus, the i th seller's consumption of rice is given by :

$$1 + \left(\sum_{j=1}^n d_j / \sum_{j=1}^n s_j \right) s_i - s_i^\beta / \beta. \quad (3.1)$$

The i th seller's problem is to maximize this expression by choosing s_i subject to the endowment of rice. That is, $s_i^\beta/\beta \leq 1$.

Let us now turn to the problem faced by the buyers. A buyer seeks to maximize the sum of the quantities of rice and rice pudding. If the i th buyer puts up d_i units of rice for exchange, then the quantity of rice consumed is $1 - d_i$, unless supply is zero. The quantity of rice pudding consumed is $(\sum_{j=1}^n s_j / \sum_{j=1}^n d_j) d_i$. Thus, the i th buyer seeks to maximize

$$1 - d_i + \left(\sum_{j=1}^n s_j / \sum_{j=1}^n d_j \right) d_i \quad (3.2)$$

by choosing d_i subject to the constraint that $d_i \leq 1$.

This game has a trivial equilibrium in which every trader chooses zero. There is, however, a Pareto superior equilibrium in which traders choose positive quantities. There are potential gains from trade since the marginal cost for a seller of producing rice pudding is $s_i^{\beta-1}$ (which is less than 1 for $s_i < 1$), while the buyers value rice pudding at 1. The efficient outcome would be for all sellers to produce one unit of rice pudding and all the buyers to put up one unit of rice in exchange. In this case, the price and marginal cost are equated at one. With a finite population, each trader exercises some market power so that production in equilibrium falls somewhat short of the first best. However, this discrepancy becomes small as n becomes large. The first order condition for the i th seller is :

$$\left(\sum_j d_j / \sum_j s_j \right) \left(1 - (s_i / \sum_j s_j) \right) - s_i^{\beta-1} = 0 \quad (3.3)$$

while that for the i th buyer is :

$$\left(\sum_j s_j / \sum_j d_j \right) \left(1 - (d_i / \sum_j d_j) \right) - 1 = 0. \quad (3.4)$$

In the symmetric equilibrium where $s_j = s$ and $d_j = d$ for all j , these conditions yield $d = \left(\frac{n-1}{n} \right) s$ and $s^\beta = \left(\frac{n-1}{n} \right) d$. Hence in the symmetric equilibrium,

$$s^* = \left(\frac{n-1}{n} \right)^{\frac{2}{\beta-1}} \quad \text{and} \quad d^* = \left(\frac{n-1}{n} \right)^{\frac{\beta+1}{\beta-1}}. \quad (3.5)$$

Both these quantities fall short of 1, but as n becomes large, both s^* and d^* tend to the first best level of 1. This is a consequence of the diminishing market power of each trader.

Dealership Market. The outcome in the dealership market is simpler to analyse. The buyers follow a simple rule in which they seek the sellers with the lowest price and trade up to their endowment provided that price is one or below. The sellers anticipate this behaviour and choose production and price. In fact in every equilibrium, each seller produces one unit of rice pudding and posts a price of 1. The argument is as follows.

Firstly, note that if a seller produces a positive amount, he will not post a price strictly greater than 1, since the demand facing this seller will be zero. Since the marginal cost of rice pudding is $s_i^{\beta-1}$ while the marginal benefit is the price posted, $s_i^{\beta-1}$ does not exceed price which, in turn, is at most 1. Hence no seller will produce more than one unit of rice pudding. Thus, in any equilibrium,

$$\sum_{j \neq i} s_j \leq n-1. \quad (3.6)$$

Since the total endowment of all buyers is n units of rice, the i th seller can sell at least one unit of rice pudding provided the price is one or below. Thus, the optimal action for the i th seller is to produce one unit of rice pudding and post the price of one. The seller receives one unit of rice in exchange. Therefore, in any equilibrium of the dealership market, every seller produces one unit of rice pudding, which is exchanged for one unit of rice. This is the efficient allocation. Notice also that the argument here relies on two rounds of deletion of strictly dominated strategies, rather than equilibrium reasoning.

When comparing the order-driven market with the dealership market, we see that the differences in post-trade allocations are minimal, provided that the population is large. Thus, in the benchmark case when the endowments are common knowledge, the two trading mechanisms yield similar outcomes. We will now argue that this similarity of outcomes is not preserved when the traders have differential information.

4. Differential Information

The uncertainty in our model stems from the variability of rainfall in the two regions.

Each state ω in our state space Ω will represent a possible combination of rainfall in the two regions. Thus, a pair of natural numbers (q, r) is a state if q and r belong to the set $\{0, 1, 2, \dots, N\}$, and either $q = r$ or $q = r - 1$. Hence, the state space Ω can be represented as in figure 1. The points on the leading diagonal represent those states in which both regions have the same number of days of rain, while those on the subordinate diagonal represent those states in which the mountains have one more day of rain.

[Figure 1 here]

Since each trader only observes the rainfall in his own region, the coastal rice growers cannot distinguish two states which have the same first component, while a mountain grower cannot distinguish two states with the same second component. Hence, the information partition of a coastal grower is generated by the equivalence relation $\stackrel{C}{\sim}$ on Ω defined as :

$$(q, r) \stackrel{C}{\sim} (q', r') \Leftrightarrow q = q' \quad (4.1)$$

The information partition of a mountain grower is generated by the equivalence relation $\stackrel{M}{\sim}$ defined as :

$$(q, r) \stackrel{M}{\sim} (q', r') \Leftrightarrow r = r' \quad (4.2)$$

We define knowledge and iterated knowledge of the traders in the usual way. A trader knows an event E at the state ω if the element of his partition which contains ω is a subset of E . The knowledge operator K_C is defined so that the $K_C E$ is the event that a coastal trader knows E . In other words,

$$K_C E \equiv \{\omega \mid \omega \stackrel{C}{\sim} \omega' \Rightarrow \omega' \in E\} \quad (4.3)$$

The knowledge operator K_M for a mountain grower is defined analogously.

The event which will play an important role in our analysis is that in which the harvests of both regions are good (and hence equal to 1). In our model, the harvest in a particular region is good if the number of days of rain there is one or more. Hence, the event in which both harvests are good is given by :

$$G \equiv \{\omega \mid \omega \geq (1, 1)\}. \quad (4.4)$$

This event is of interest to us since G consists of precisely those states at which the efficient allocation (in which every seller produces one unit of rice pudding and receives one unit of rice in return) is feasible. The “ G ” stands for “good”. If G were common knowledge the equilibria described in the previous section come into play. However, the differential information in our model is such that G is never common knowledge. To explain this point, and to introduce our notion of transparency, let us use the notation $(K_M K_C)^k G$ to denote the event obtained by applying the operator $K_M K_C$ to G k times. This is the event in which a seller knows that a buyer knows that a seller knows that \dots a buyer knows that G is true, where the order of iterated knowledge is k . Figure 2 illustrates some of these events.

[Figure 2 here]

The event $K_M K_C G$ consists of states ω such that the element of the mountain grower’s partition which contains ω is a subset of $K_C G$. Inspection of figure 2 shows that this is the case for all but one state in $K_C G$, namely $(1, 1)$. At state $(1, 1)$, a mountain grower must allow for the possibility that $K_C G$ is false, since he cannot distinguish state $(1, 1)$ from $(0, 1)$. Hence, $K_M K_C G = \{(q, r) \mid q \geq 1, r \geq 2\}$. If we now apply the operator K_C to this event, we lose the state $(1, 2)$, since at $(1, 2)$ a coastal grower cannot distinguish between $(1, 2)$ and $(1, 1)$, and the latter state does not belong to $K_M K_C G$. Hence, $K_C K_M K_C G = \{(q, r) \mid q \geq 2, r \geq 2\}$. Finally, if we apply the operator K_M to the event $K_C K_M K_C G$, we lose the state $(2, 2)$, since at $(2, 2)$ a mountain grower cannot distinguish $(2, 2)$ from $(1, 2)$, and the latter does not belong to $K_C K_M K_C G$. Hence, $(K_M K_C)^2 G = \{(q, r) \mid q \geq 2, r \geq 3\}$. In general, we have :

LEMMA 1. $(K_M K_C)^k G = \{(q, r) \mid q \geq k, r \geq k+1\}$.

PROOF. The proof is by induction on k . We have already noted that the lemma holds for $k = 1$. Suppose that the lemma holds for $k-1$. Then, $(k-1, k) \notin K_C (K_M K_C)^{k-1} G$,

since $(k-1, k) \stackrel{C}{=} (k-1, k-1)$, but $(k-1, k-1) \notin (K_M K_C)^{k-1} G$. But for every state (q, r) such that $(q, r) \geq (k, k)$, if $(q, r) \stackrel{C}{=} (q', r')$ then $(q', r') \in (K_M K_C)^{k-1} G$. Hence,

$$K_C(K_M K_C)^{k-1} G = \{(q, r) \mid q \geq k, r \geq k\}. \quad (4.5)$$

Finally, note that $(k, k) \stackrel{C}{=} (k, k-1)$, but $(k, k-1) \notin K_C(K_M K_C)^{k-1} G$, which implies that $(k, k) \notin K_M K_C(K_M K_C)^{k-1} G$. However, for every $(q, r) \geq (k, k+1)$, if $(q, r) \stackrel{C}{=} (q', r')$ then $(q', r') \in K_C(K_M K_C)^{k-1} G$. Hence,

$$K_M K_C(K_M K_C)^{k-1} G = \{(q, r) \mid q \geq k, r \geq k+1\}. \quad \square$$

A corollary of lemma 1 is that there is no state at which G is common knowledge. If G were common knowledge at a state ω , then $\omega \in (K_M K_C)^k G$ for all k , which is inconsistent with a state being a pair of natural numbers, in view of lemma 1. In the terminology of Morris, Postlewaite and Shin (1993), the depth of knowledge embodied in our information structure is finite.

Lemma 1 sets the agenda for the rest of the paper. We pose the following question. How robust are the two trading mechanisms to higher order uncertainty? If G is known up to degree k but not to degree $k+1$, how far do the outcomes in the two mechanisms diverge from each other, and how far do they diverge from the benchmark case of common knowledge? This question motivates our definition of transparency. For every state ω , we associate a natural number $t(\omega)$, called the transparency of the state ω , where,

$$t(\omega) = \begin{cases} 0 & \text{if } \omega \notin K_M K_C G \\ k & \text{if } \omega \in (K_M K_C)^k G \text{ but } \omega \notin (K_M K_C)^{k+1} G \end{cases} \quad (4.6)$$

We now investigate the effect of transparency on the volume of trade. The performance of each trading mechanism will be assessed in terms of its ability to generate liquidity, both in the market and in the pudding bowl. We begin with the order-driven market.

5. Trade in Order-Driven Market

Since the traders choose their quantities after observing the rainfall in their own region, the strategies of the traders are functions from the set $\{0, 1, 2, \dots, N\}$ to their respective

action sets. A strategy for the i th seller is a function s_i which maps the number of days of rain in the mountains to the quantity of rice pudding supplied. Similarly, a strategy for the i th buyer is a function d_i which maps the number of days of rain on the coast to the amount of rice put up for exchange.

Since a state ω specifies the rainfall in both regions, we may also see the strategies as functions which map each state ω to the action at that state, with the proviso that if a trader cannot distinguish state ω from ω' , then the action is the same at the two states (i.e.) a strategy is measurable on the information partition of the trader. In what follows, we will use both formalizations of a strategy. Thus, $s_i(\omega)$ is the supply of the i th trader at state ω , while for integer k , $s_i(k)$ is the supply of the i th trader given k days of rain in the mountains. $d_i(\omega)$ and $d_i(k)$ are defined analogously for the buyer.

An equilibrium of the order-driven market is a vector of strategies $(s_1, \dots, s_n, d_1, \dots, d_n)$ such that for every state ω and every trader, the action prescribed by that trader's strategy maximizes his expected payoff conditional on the information at ω , given that all other traders follow their respective strategies.

The first of our pair of theorems highlights the role of transparency on the volume of trade. For any symmetric equilibrium and at every state ω , supply always falls short of s^* , the equilibrium supply with common knowledge, as given by (3.5). Trade is increasing in transparency, but no amount of transparency (short of being infinite) can restore the outcome under common knowledge.

THEOREM 1. For any symmetric equilibrium supply strategy s , $s(\omega) < s^*$ for every ω . Moreover, if supply is positive at some state, then for all $\omega, \omega' \in \Omega$, if $t(\omega) < t(\omega')$, then $s(\omega) < s(\omega')$.

The second result shows that the damage to trade at potentially “good” states can be quite serious. The extent of the damage depends on the cost parameter β . If β is close to 1 (so that the cost of rice pudding production is high), then irrespective of the number of traders,

the supply of rice pudding is close to zero at every state. More precisely, our second theorem reads :

THEOREM 2. For any $\epsilon > 0$, there is a benchmark level $\beta^* > 1$ of the cost parameter such that, if $\beta < \beta^*$, then for any population n and any symmetric equilibrium supply strategy s , $s(\omega) < \epsilon$ everywhere.

The important point to drive home is that the low trading volume described in theorem 2 is not the kind which will disappear when the population becomes large. The benchmark β^* applies to any population n . Hence, the mechanism at work cannot be that of market power. Rather, as we shall see below, it is the absence of common knowledge of the event G which feeds into higher order uncertainty, and which in turn, starves the market of liquidity. The closer β is to 1, the stronger is this effect.

Before presenting the formal arguments, it is instructive to pause to think about the forces at work generating these results. To fix ideas, let us consider the reasoning of a seller who has observed 17 days of rain. Should he produce a lot of rice pudding? The answer is yes if he believes that demand will be high. He knows that buyers have had 16 or 17 days of rain, and so have the full endowment of rice, but this by itself does not guarantee that demand will be high, since buyers' actions are chosen rationally given their beliefs. Thus, in order to decide what the seller should do, it is necessary to think about the beliefs of buyers who have observed 16 or 17 days of rain.

A buyer who has observed 16 days of rain knows that the sellers have had 16 or 17 days of rain, and so have the full endowment of rice. A buyer will submit a large order if he believes that supply will be high, but the mere fact that sellers are capable of producing a lot is not enough for high supply. The supply depends on the beliefs of the sellers who have observed 16 or 17 days of rain.

Thus, a seller who has observed 17 days of rain must worry about what he would have done given 16 days of rain, since the buyers care about this, and the seller cares about the beliefs of the buyers. But then, the reasoning does not stop there, since the seller's optimal

action given 16 days of rain depends on the optimal action given 15 days, and by iteration, all lower numbers. Ultimately, the belief hierarchy must include beliefs about actions given zero days of rain, at which supply is constrained to be zero, since there is no rice around. This bad event contaminates the belief hierarchy, since at no state is there common knowledge that this is not the case, and it is this which reduces trading volume. The parameter β gives an indication as to how much of a “cushion” there is in the system to absorb this higher order uncertainty. The closer β is to 1, the less there is to shield the actions at states with high rainfall, and hence the lower is the trading volume.

The following proofs of theorems 1 and 2 are *ad hoc* to our model, and involve inductions on the number of days of rain. However, in section 7, an alternative proof of theorem 2 is given which relies on general features of the information structure.

Let us first consider the problem faced by the i th buyer who has observed that the number of days of rain on the coast is k , where $k < N$. Two states are consistent with this observation, (k, k) and $(k, k+1)$. Since the probability distribution over Ω is uniform, the two states are given equal weight. If (k, k) is the true state, and the traders follow strategies $(s_1, \dots, s_n, d_1, \dots, d_n)$, then the utility of the i th buyer is $1 - d_i(k) + \frac{\sum_{j=1}^n s_j(k)}{\sum_{j=1}^n d_j(k)} d_i(k)$. If $(k, k+1)$ is the true state, utility has a similar expression, except that $s_j(k)$ is replaced by $s_j(k+1)$. Thus, expected utility given the message k is :

$$1 - d_i(k) + \left(\frac{\sum_{j=1}^n s_j(k) + \sum_{j=1}^n s_j(k+1)}{2} \right) \frac{d_i(k)}{\sum_{j=1}^n d_j(k)} \quad (5.1)$$

The i th buyer seeks to maximize this expression by choosing $d_i(k)$ subject to the resource constraint. The constraint is $d_i(k) \leq 1$ for $k \geq 1$, while for $k = 0$, d_i is constrained to be zero due to the zero harvest. The buyer’s problem given message N is simpler, since only one state is consistent with this message – namely, (N, N) . The i th buyer chooses $d_i(N)$ to maximize :

$$1 - d_i(N) + \left(\sum_{j=1}^n s_j(N) \right) \frac{d_i(N)}{\sum_{j=1}^n d_j(N)} \quad (5.2)$$

subject to the constraint $d_i(N) \leq 1$.

Let us now turn to the problem facing the i th seller, given the message k , where $k \geq 1$. This message is consistent with two states - namely, $(k-1, k)$ and (k, k) . Both states receive equal weight. Thus, the i th seller chooses $s_i(k)$ to maximize :

$$1 + \left(\frac{\sum_{j=1}^n d_j(k-1) + \sum_{j=1}^n d_j(k)}{2} \right) \left(\frac{s_i(k)}{\sum_{j=1}^n s_j(k)} \right) - s_i(k)^\beta / \beta \quad (5.3)$$

subject to the constraint that the amount of rice used in the production of rice pudding cannot exceed 1. That is, $s_i(k)^\beta / \beta \leq 1$. Given message $k = 0$, only the state $(0, 0)$ is consistent with this message. However, the zero harvest constrains the seller to set $s_i(0) = 0$.

Ignoring resource constraints for the moment, the symmetric solution to the first-order conditions yield :

$$d(k) = \left(\frac{n-1}{n} \right) \left(\frac{s(k) + s(k+1)}{2} \right) \quad \text{if } k < N \quad (5.4)$$

$$= \left(\frac{n-1}{n} \right) s(N) \quad \text{if } k = N$$

$$s(k)^\beta = \left(\frac{n-1}{n} \right) \left(\frac{d(k-1) + d(k)}{2} \right) \quad \text{if } k \geq 1 \quad (5.5)$$

We should now check that the resource constraints are satisfied. It turns out that for $k \geq 1$, the constraints do not bind. We can argue by contradiction. Suppose that the constraint $d(k) \leq 1$ binds for some non-empty set of integers $J \subseteq \{1, 2, \dots, N\}$. Then $d(k) = 1$ for $k \in J$, and $d(k) < 1$ for $k \notin J$. In any case, $d(k) \leq 1$ for all k . If $k \in J$, then (5.4) implies that $s(k) > 1$ or $s(k+1) > 1$ or both. In either case, (5.5) implies that one or more of $d(k-1)$, $d(k)$, and $d(k+1)$ is strictly larger than 1, which is a contradiction. But then, if the buyers' constraints do not bind, (5.5) implies that $s(k) < 1$, so that the sellers' constraints are non-binding also. However, for $k = 0$, the zero harvest implies that $s(0) = d(0) = 0$.

Substituting (5.4) into (5.5) we obtain a second-order, non-linear difference equation for $s(k)$ for $k \in \{1, 2, \dots, N-1\}$ and a terminal condition for $s(\cdot)$. They are, respectively,

$$s(k)^\beta = \left(\frac{n-1}{n}\right)^2 \left(\frac{s(k-1) + 2s(k) + s(k+1)}{4}\right) \quad (5.6)$$

and

$$s(N)^\beta = \left(\frac{n-1}{n}\right)^2 \left(\frac{s(N-1) + 3s(N)}{4}\right) \quad (5.7)$$

We can now demonstrate theorem 1 by means of the following lemma. Recall that s^* is the equilibrium supply when G is common knowledge.

LEMMA 2. If $\{s(k)\}$ satisfies (5.6) and (5.7) for $k \in \{1, 2, \dots, N\}$, and $s(N) \neq 0$, then

$$s(N) \geq s^* \quad \Rightarrow \quad s(0) \geq s(1) \geq s(2) \geq \dots \geq s(N-1) \geq s(N)$$

$$s(N) < s^* \quad \Rightarrow \quad s(0) < s(1) < s(2) < \dots < s(N-1) < s(N).$$

PROOF. Suppose $s(N) \geq s^*$. We shall argue by induction, backwards from N . Begin by recalling that $s^* = \left(\frac{n-1}{n}\right)^{\frac{2}{\beta-1}}$, so that $s(N) \geq s^*$ implies:

$$\left(\frac{n}{n-1}\right)^2 s(N)^\beta \geq s(N). \quad (5.8)$$

(5.7) then implies that $s(N-1) \geq s(N) \geq s^*$. Thus, suppose that $s(k) \geq s(k+1) \geq s^*$.

From (5.6),

$$\begin{aligned} s(k) &= \frac{1}{2} \left(s(k-1) + s(k+1) \right) + 2 \left(s(k) - \left(\frac{n}{n-1}\right)^2 s(k)^\beta \right) & (5.9) \\ &\leq \frac{1}{2} \left(s(k-1) + s(k+1) \right) & \text{(since } s(k) \geq s^*) \\ &\leq s(k-1) & \text{(since } s(k) \geq s(k+1)) \end{aligned}$$

so that $s(k) \leq s(k-1)$, as desired. An analogous argument holds for $s(N) < s^*$. \square

We can now complete the proof of theorem 1. First of all, note that if $s(k) = 0$ for some k , then (5.4) and (5.5) imply that $s(k) = d(k) = 0$ for every k . This is the trivial equilibrium in which every trader submits zero. Thus, if $s(k) > 0$ for some k , then $s(k) > 0$ for every k . Under the hypothesis of theorem 1, it must be the case that $s(N) < s^*$, since otherwise lemma 2 implies that $s(0) \geq s^*$ which violates the resource constraint $s(0) = 0$. Given that $s(N) < s^*$, lemma 2 tells us that $s(k)$ is strictly increasing

in k . It remains for us to note that if ω is the state (q, r) and ω' is the state (q', r') , then $t(\omega) < t(\omega')$ only if $r < r'$. Thus, for any symmetric equilibrium in which supply is positive at some state, $t(\omega) < t(\omega')$ only if $s(\omega) < s(\omega')$. This completes the proof of theorem 1.

We now turn to the proof of theorem 2. Let us define $\Delta(\beta)$ as :

$$\Delta(\beta) \equiv \max \left\{ s - s^\beta \mid s \in [0, 1] \right\}. \quad (5.10)$$

$\Delta(\beta)$ is positive for $\beta > 1$ and $\Delta(\beta)$ tends to zero as β tends to 1. We then have :

LEMMA 3. For any symmetric equilibrium supply strategy s ,

$$s(0) \geq s(N) - 2N(N+1)\Delta(\beta). \quad (5.11)$$

PROOF. If $s(0) \geq s(N)$ there is nothing to prove. Thus, suppose $s(0) < s(N)$. Then, from theorem 1, $s(k) < s^*$ for all k , so that $s(k) \geq \left(\frac{n}{n-1}\right)^2 s(k)^\beta$. Then, (5.6) gives

$$\begin{aligned} s(k-1) &= 2s(k) - s(k+1) - 4 \left(s(k) - \left(\frac{n}{n-1}\right)^2 s(k)^\beta \right) \\ &\geq 2s(k) - s(k+1) - 4\Delta(\beta). \end{aligned} \quad (5.12)$$

We now claim that $s(0) \geq (k+1)s(k) - ks(k+1) - 4(1 + 2 + 3 + \dots + k)\Delta(\beta)$ for any $k \in \{1, 2, \dots, N\}$. The proof is by induction. For $k = 1$, substitution into (5.12) yields the desired inequality. For the inductive step, suppose that $s(0) \geq ks(k-1) - (k-1)s(k) - 4(1 + 2 + 3 + \dots + k-1)\Delta(\beta)$. Then, substituting out $s(k-1)$ by using (5.12) yields the desired inequality, which proves the claim. For the case when $k = N-1$,

$$\begin{aligned} s(0) &\geq Ns(N-1) - (N-1)s(N) - 4(1 + 2 + 3 + \dots + N-1)\Delta(\beta) \\ &= s(N) + N(s(N-1) - s(N)) - 4(1 + 2 + 3 + \dots + N-1)\Delta(\beta) \\ &\geq s(N) - 4N\Delta(\beta) - 4(1 + 2 + 3 + \dots + N-1)\Delta(\beta) \quad (\text{from (5.7)}) \\ &= s(N) - 2N(N+1)\Delta(\beta). \quad \square \end{aligned}$$

To complete the proof of theorem 2, we choose β^* so that $\Delta(\beta^*) \leq \epsilon/2N(N+1)$. Then for $\beta < \beta^*$, we have $2N(N+1)\Delta(\beta) < \epsilon$ so that from lemma 3, and the fact that $s(0) = 0$, we have $s(N) < \epsilon$. But since $s(k) < s(N)$ for all $k < N$, supply is below ϵ everywhere.

Many questions suggest themselves concerning the generality of these results. For now, we note that although our model has a story concerning production, it is clear that a similar story could be told in a pure exchange economy. The production of rice pudding plays a role only to the extent that the seller's marginal rate of substitution between rice and rice pudding in the exchange differs from the buyer's. For instance, the utility function $x + y - y^\beta/\beta$ for the seller in a pure exchange economy will reproduce all the necessary steps in the argument. One small addition to the rules would be to specify what the final allocations are if one side of the market has no traders. In Shapley and Shubik (1977), infinite prices are allowed, so that if there is no one on the other side of the market, one loses the submitted quantity. This ensures that allocations are continuous at extreme prices, and the argument above goes through.

6. Trade in Dealership Market

In contrast to the order-driven market, trading volume in the dealership market does not suffer from the fragility to higher-order uncertainty exhibited by its order-driven cousin. In a dealership market, a strategy for the i th seller is given by an ordered pair of functions $(s_i(\cdot), p_i(\cdot))$, where $s_i(\omega)$ is the amount of rice pudding produced at state ω and $p_i(\omega)$ is the price posted at ω , with the restriction that both s_i and p_i are measurable on a seller's information partition. For a buyer, the information available before taking an action includes not only that yielded by his information partition, but also includes the vector of actions taken by all the sellers. A buyer's strategy is a rule which, based on all this information, selects a subset of sellers with whom to trade, and which ranks this set of sellers in order of preference. Our result on the dealership market can be stated as follows.

THEOREM 3. In any equilibrium of the dealership market and for any state ω , if $t(\omega) \geq 1$, then $(s_i(\omega), p_i(\omega)) = (1, 1)$ for every i .

In other words, if $\omega \in K_M K_C G$ then the outcome under common knowledge will transpire at ω in equilibrium. Iterated knowledge of G of order 1 is sufficient to implement the efficient trade. Since the event $K_M K_C G$ includes every state other than $(0, 0)$, $(0, 1)$ and $(1, 1)$, the dealership market performs well in preserving liquidity (both in the market and in the pudding bowl). The contrast with the order-driven market could hardly be more stark. Before examining some of the forces at work which generate these differences, we shall demonstrate theorem 3.

We solve the game backwards, starting with the actions of the buyers at the last stage of the game. A buyer selects all sellers who have posted a price of one or below, and ranks them in reverse order in terms of price. For a buyer's action to be optimal, it must never be the case that the buyer trades with a seller who has posted a price strictly higher than another seller who has unsold stock of rice pudding.

Now, suppose that $\omega \in K_M K_C G$, and consider the j th seller's supply $s_j(\omega)$ and price $p_j(\omega)$ at ω . Since $K_M K_C G \subset K_C G = G$ and $\omega \in K_M K_C G$, two features hold at ω . Firstly, all buyers have the maximum endowment of rice of 1. Secondly, all sellers know that every buyer has this endowment. Any action $(s_j(\omega), p_j(\omega))$ of the j th seller for which $s_j(\omega) > 0$ and $p_j(\omega) > 1$ is strictly dominated by any action for which $s_j(\omega)$ is zero, since no buyer will trade with a seller with a posted price greater than 1. Hence, if $(s_j(\omega), p_j(\omega))$ is an equilibrium action of the j th seller, then

$$s_j(\omega) > 0 \Rightarrow p_j(\omega) \leq 1. \quad (6.1)$$

Since the marginal cost of producing rice pudding is $s_j^{\beta-1}(\omega)$ while the marginal benefit is at most the price posted, we have $s_j^{\beta-1}(\omega) \leq p_j(\omega)$, which together with (6.1) implies $s_j(\omega) \leq 1$. Hence, in any equilibrium, $\sum_{j \neq i} s_j(\omega) \leq n-1$, so that the residual demand facing the i th seller is given by :

$$D = \begin{cases} n - \sum_{j \neq i} s_j(\omega) \geq 1 & \text{if } p_i \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (6.2)$$

Thus, the i th seller's consumption of rice (and hence utility) is given by :

$$\begin{aligned}
u_M &= 1 + p_i(\omega)s_i(\omega) - s_i(\omega)^\beta/\beta && \text{if } s_i(\omega) \leq D \text{ and } p_i(\omega) \leq 1 \\
&= 1 + Ds_i(\omega) - s_i(\omega)^\beta/\beta && \text{if } s_i(\omega) > D \text{ and } p_i(\omega) \leq 1 \\
&= 1 - s_i(\omega)^\beta/\beta && \text{if } p_i(\omega) > 1
\end{aligned}$$

which is maximized when $(s_i(\omega), p_i(\omega)) = (1, 1)$. Since the argument is symmetric for all sellers, every seller produces 1 unit of rice pudding and posts a price of 1. In any equilibrium, n units of rice pudding are exchanged for n units of rice.

In the efficient allocation, the gains from trade are appropriated by the sellers. The buyers are no better off than before trade. This has been assumed in order to simplify the analysis, and it would not be difficult to consider preferences for the buyers which will make the efficient allocation strictly Pareto-superior. A fuller analysis ought to consider distributive issues as well as efficiency, and should be the subject of further research. Also, let us note that the efficiency of the dealership market is not sensitive to who sets prices. If the buyers were to set prices instead, price competition will drive up prices to one, and efficient production takes place. The only difference will be that the single layer of knowledge needed will involve the operator $K_C K_M$ rather than $K_M K_C$.

7. Assessing the Argument

The superiority of the dealership market in terms of greater allocative efficiency is in need of explanation. The divergence in the performance of the two markets is especially noteworthy in view of the similarity of the outcomes in these markets in the absence of uncertainty. At a superficial level, it is certainly correct to say that uncertain transaction price is the culprit in causing inefficiency in the order-driven market. However price uncertainty in our model is endogenous. The interesting question, therefore, is why the traders cannot act in such a way as to remove this type of risk. Even more narrowly, what is to stop the traders in the order-driven market mimicking the workings of the dealership market?

The answers to these questions lie in the details of the trading game, especially the move order of the traders and the information available to traders when they take their respective

actions. The order-driven market requires that all traders take their actions simultaneously, whereas in the dealership market, the price setters move first and buyers take their actions having observed the actions of the sellers. Although this distinction may seem unimportant at first sight, differential information has rather different effects in the two markets.

Our modelling of a simultaneous move game with differential information has followed the standard technique, due to Harsanyi (1967), of casting it as a Bayesian game in which distinct “types” of each player play their equilibrium strategies. In the terminology of “types”, each trader could be one of $N+1$ possible types, corresponding to each possible piece of information which arrives before taking an action. Each cell in a trader’s information partition corresponds to a type of that trader. In such a setting, although two types of a single trader are mutually exclusive, it is quite possible that the actions of these types are constrained through the best-reply structure of the game. The clearest illustration of this point is equation (5.6), which states that :

$$s(k)^\beta = \left(\frac{n-1}{n}\right)^2 \left(\frac{s(k-1) + 2s(k) + s(k+1)}{4}\right).$$

Here, the actions of three distinct types (a seller who has observed $k-1$, k and $k+1$ days of rain respectively) are constrained by this relationship, even though a seller who observes one of these messages can exclude the other two possibilities. The point is that a seller of type k is concerned with the actions of buyers of types k and $k-1$, who in turn care about the actions of the sellers of type $k-1$, k , and $k+1$.

Such restrictions across the cells of a decision maker’s information partition is in marked contrast to single person decision theory. In a single person decision problem, it is possible to divide the world neatly into mutually exclusive parcels labelled by the message received, and then take the optimal action given each message, quite independently of the actions taken at other states. However, as soon as we have a game, this facility for dividing the world into neat parcels no longer exists, and actions at one part of the state space will influence (and be influenced by) the actions at other parts of the state space. In the context of two player Bayesian games with finite action sets these effects can be studied in a

reasonably complete manner, as shown by Morris, Rob and Shin (1995), which in turn builds on the earlier work of Rubinstein (1989) and Monderer and Samet (1989), who noted the importance of higher order uncertainty, and the work of Carlsson and van Damme (1993) on the notion of global uncertainty (see also Sorin (1993) on a definition of the impact of an event). Geanakoplos (1992) is a survey of the issues raised by some of the earlier papers.

The poor performance of the order-driven market should be understood in these terms. The influence of states with the poor harvest is transmitted through the best-reply structure of the game to states which, potentially, could have high trading volume. The actual size of this effect depends on how sensitive the actions of one type are to the actions of neighbouring types. Theorem 2 showed that the cost parameter β is an important determinant of this sensitivity.

The reasons for the relative superiority of the dealership market then become more transparent. By allowing the buyers to observe the actions of the sellers before taking their actions, the dealership market cuts the link between the equilibrium actions of the different types of the seller. When a seller is setting the price-quantity pair, there may be uncertainty about the type of the buyer, but the seller knows that a buyer will play a best reply to the seller's action having observed that action. Crucially, the seller need not worry about what he would have done had he been a different type. Such considerations of counterfactual propositions are redundant given that the buyers choose their action only after observing the seller's action. In this sense, the sequential move structure of the dealership market restores the feature, present in the single person decision problem, of allowing the decision maker to disregard the optimal actions at those states which are excluded by the current message.

How general are the results reported here for the Shapley-Shubik mechanism? The *ad hoc* argument given in section 5 relied on the 'stationary' nature of the restrictions, in which the same qualitative restriction applies to type k as it does to type $k+1$. Although this stationarity gave the model sufficient tractability to use an induction argument, it is not

essential, as we will see with the following alternative argument for theorem 2. This argument will be important in suggesting possible extensions of our analysis.

An Alternative Argument. Consider a finite Markov chain with $N+1$ states, labelled by the set $\{0, 1, 2, \dots, N\}$. The one-step transition probability from i to j is denoted by $\phi(i, j)$. State 0 is an absorbing state, and is the only such state. Hence, $\phi(0, 0) = 1$, and $\phi(0, j) = 0$ for any $j \neq 0$. From state $i \in \{1, 2, \dots, N-1\}$, three transitions are possible. The system returns to i with probability $1/2$, while it progresses to $i-1$ and $i+1$ with probability $1/4$ each. In other words, for $i \in \{1, 2, \dots, N-1\}$,

$$\phi(i, j) = \begin{cases} \frac{1}{2} & \text{if } i = j \\ \frac{1}{4} & \text{if } j = i+1 \text{ or } j = i-1 \\ 0 & \text{otherwise} \end{cases} \quad (7.1)$$

Finally, from state N , the system returns to N with probability $3/4$, and makes a transition to $N-1$ with probability $1/4$.

This is a Markov chain, all of whose states are transient, except for the single absorbing state 0. Then, equilibrium $s(\cdot)$ satisfies the following key properties.

LEMMA 4. Let s be a symmetric equilibrium supply strategy. Then, for any $i \neq 0$,

(i) $s(i) \geq \sum_j \phi(i, j)s(j)$,

(ii) for any $\epsilon > 0$, there is $\beta^* > 0$ such that, for any $\beta < \beta^*$ and any n ,

$$s(i) - \sum_j \phi(i, j)s(j) < \epsilon. \quad (7.2)$$

PROOF. To see (i), recall that $s(i) < s^*$ from lemma 1, so that $s(i) > \left(\frac{n}{n-1}\right)^2 s(i)^\beta$. Then (i) is just a restatement of equations (5.6) and (5.7). For (ii), let $\epsilon > 0$ be given. Choose β^* sufficiently close to 1 so that for the function Δ defined in (5.10), $\Delta(\beta^*) \leq \epsilon$. Then, for $\beta < \beta^*$ we have $s(i) - \left(\frac{n}{n-1}\right)^2 s(i)^\beta < s(i) - s(i)^\beta < \epsilon$, so that (7.2) follows from (5.6) and (5.7). \square

If we regard $s(\cdot)$ as a function of the underlying Markov chain $\{X_t\}$ taking values in $\{0, 1, 2, \dots, N\}$, clause (i) above states that $s(\cdot)$ is a *supermartingale* with respect to this Markov chain. This is because $E(s(X_{t+1}) \mid X_t = i) = \sum_j \phi(i, j)s(j) \leq s(i)$, so that

$$s(X_t) \geq E(s(X_{t+1}) | X_t). \quad (7.3)$$

Functions which satisfy (i) are also referred to as *superregular* (see, for example, Karlin and Taylor (1981, p.45)). A precursor to the following argument is in Shin and Williamson (1992), which derives the more restrictive condition that equilibrium actions in a coordination game is a *martingale* with respect to an analogous Markov chain.

The matrix of transition probabilities Φ whose (i, j) th entry is $\phi(i, j)$, can be partitioned as:

$$\Phi = \begin{pmatrix} I & 0 \\ R & Q \end{pmatrix}, \quad (7.4)$$

where I is the trivial 1×1 identity matrix, 0 is a row vector of zeros, R is a column vector with N entries all of which are zero except the top entry which is $1/4$, and Q is an $N \times N$ matrix. Equilibrium supply can be expressed in a succinct way in terms of the matrix Q . This is because $s(0) = 0$, so that if we denote by s the column vector :

$$s \equiv \begin{pmatrix} s(1) \\ s(2) \\ \vdots \\ s(N) \end{pmatrix},$$

of supplies for positive number of days of rain, lemma 4(i) gives us the vector inequality $s \geq Qs$. Also, for any positive integer k , we have the identity :

$$s = Q^{k+1}s + (I + Q + \dots + Q^k)(s - Qs). \quad (7.5)$$

Since the matrix Q describes the transitions among the transient states, the series $(I + Q + \dots + Q^k)$ converges to some finite $N \times N$ matrix $V \equiv \sum_{k=0}^{\infty} Q^k$, whose (i, j) th entry is the expected number of visits to state j given that the process starts at state i . (Since the number of states is finite, $I - Q$ is non-singular, and $V = (I - Q)^{-1}$). Clearly, Q^k tends to the zero matrix as k becomes large. Since $s \geq Qs$, the second term on the right hand side of (7.5) is increasing in k . Thus, we may appeal to monotone convergence in passing to the limit in (7.5), to yield :

$$s = Vh, \quad (7.6)$$

where $h \equiv s - Qs \geq 0$.

Equation (7.6) gives us an intriguing interpretation of equilibrium supply. Imagine the Markov process starting at state $i > 0$. Each time the process visits some transient state j , there is a “prize” of $h_j \geq 0$, where h_j is the j th component of h . The equilibrium supply at state i is then the expected aggregate prize earned when the process starts from state i . Since all states other than 0 are transient, the expected aggregate prize is finite. Equilibrium supply at i is just this expected aggregate prize.

When (7.6) is combined with lemma 4(ii), we have an immediate proof of theorem 2, which uses the idea that the “prize” for visiting a state goes to zero as β approaches 1. Given that equilibrium supply is the expected aggregate prize, equilibrium supply can be made as small as we like by choosing β sufficiently close to 1. Formally, lemma 4(ii) tells us that h tends to the zero vector as β tends to 1, irrespective of the number of traders. Since $s = Vh$, s tends to zero also. This is theorem 2.

The above proof of theorem 2 suggests that results analogous to theorem 2 may hold in contexts which are considerably more general than the model in this paper. For instance, the only property of the Markov chain which is needed in the above argument is that state 0 can be reached from all other states, but that all states other than 0 are transient. It would appear that the specific assumptions on the joint distribution over rainfall in the two regions do not play any role, other than to ensure that there is a failure of common knowledge sufficient to ensure that state 0 can be reached from all other states.

8. Some Tentative Conclusions

Although the argument using the decomposition of superregular functions holds out hope for more general results concerning Shapley-Shubik games, it does not answer the logically prior question of whether such mechanisms shed any light on the workings of realistic markets.

The Shapley-Shubik trading game has parallels with both ‘market orders’ and ‘limit orders’, but corresponds to neither exactly. The similarity with market orders lies in the

fact that buyers are uncertain as to the transaction price. However, unlike market orders, the buyer in the Shapley-Shubik game is uncertain as to how many units of the good he will obtain. Indeed, the buyer's action can be seen as the submission of a downward sloping demand curve in the price-quantity space for rice pudding, with the restriction that the demand curve be a rectangular hyperbola. This is because the buyer submits a sum of money (rice) in the game, and this action is equivalent to the submission of a set of price-quantity pairs for the good (rice pudding) which, when multiplied, yields the sum of money in question. To the extent that buyers submit downward sloping demand curves, there is a parallel with limit orders in financial markets. However, the set of admissible demand curves is severely curtailed, the restriction being that it must be a rectangular hyperbola. For the sellers, the rules of the Shapley-Shubik game correspond exactly to standard market orders. In effect, the seller is constrained to submit vertical supply curves, only.

For the specific arguments used in this paper, restrictions of the strategies to vertical supply curves and rectangular hyperbola demand curves have played an essential role. When traders are free to submit general limit orders the inefficiencies will not be large as that for the Shapley-Shubik game (at least, this would be so for the fundamentals in our model). Since the buyers have a constant valuation for rice pudding, and since all traders know the value of their own endowment, traders can prevent any trade which would make them worse off than consuming their own endowment by submitting the appropriate limit orders which would permit trade only if a net gain in utility is guaranteed². However, in more general settings in which these features do not hold, it is far from clear as to whether limit orders would ensure efficiency. For instance, if traders have common or similar valuations for an object but have imperfect information of this value (as when buying a project yielding an uncertain payoff), it is not clear that limit orders will be immune from inefficient trading volume. The performance of trading institutions will be sensitive to the fundamentals and the nature of the differential information.

² This observation is due to a referee.

There are some indications of what the ingredients of a more general approach might be. We have already commented in the previous section on the role played by the failure of common knowledge, and how this constrains actions across the cells of a trader's information partition. Equally important is the nature of these restrictions. It is important for our argument that the actions of the traders are strategic complements (Bulow, Geanakoplos and Klemperer (1985), Milgrom and Roberts (1990)). For the relevant range, the marginal benefit of a trader's action is increasing in the actions of traders on the other side of the market. More directly, the best-reply function of a trader is upward sloping in the aggregate actions of the other side of the market. Thus, buyers submit large orders if they expect sellers to submit large orders, and similarly, the sellers submit large orders if they expect the buyers to submit large orders. This is why, if some of the traders at some of the states are constrained to submit small orders, this has repercussions across all the states, as all traders react by lowering their orders, even though better outcomes are feasible.

The strategic complementarity of actions reflects the view of the market as a coordination device in which the incentive of each individual trader to engage in trade is increasing in the degree to which other traders participate. However, a difficulty in tying down this idea is that it becomes necessary to define an ordering on the action sets of traders. Defining this ordering on action sets as complex as the set of all demand functions may be problematic. Clues may lie in simpler games which are known to have strategic complementarity of actions. For instance, Morris (1992) shows that in 'acceptance games' where traders reply 'yes' or 'no' to trades proposed by a referee, the cut-off point on the space of signals are strategic complements. One of the challenges thrown up by our analysis is to demonstrate how such ideas may be applicable to realistic market institutions.

References

- Allen, B. (1981) “Generic Existence of Completely Revealing Equilibria for Economies with Uncertainty when Prices Convey Information”, *Econometrica*, 1173 – 1199.
- Allen, B. (1982) “Approximate Equilibria in Microeconomic Rational Expectations Models” *Journal of Economic Theory*, 26, 261 – 278.
- Biais, B. (1993) “Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets”, *Journal of Finance*, 48, 157 – 185.
- Bulow, J., Geanakoplos, J. and Klemperer, P. (1985) “Multi-Market Oligopoly: Strategic Substitutes and Complements”, *Journal of Political Economy*, 93, 488 – 511.
- Carlsson, H. and E. van Damme (1993) “Global Games and Equilibrium Selection” *Econometrica*, 61, 989 – 1018.
- Dubey, P., Geanakoplos, J. and Shubik, M. (1987) “The Revelation of Information in Strategic Market Games: a Critique of Rational Expectations Equilibrium”, *Journal of Mathematical Economics*, 16, 105 – 137.
- Geanakoplos, J. (1992) “Common Knowledge” *Journal of Economic Perspectives*, 6, 53 – 82.
- Grossman, S. J. and Stiglitz, J. (1980) “On the Impossibility of Informationally Efficient Markets”, *American Economic Review*, 70, 393 – 408.
- Harsanyi, J. (1967) “Games with Incomplete Information Played by Bayesian Players, parts I, II, III” *Management Science*, 14, 159 – 182, 320 – 334, 486 – 502.
- Hellwig, M. (1980) “On the Aggregation of Information in Competitive Markets”, *Journal of Economic Theory*, 22, 477 – 498.
- Jordan, J. (1982) “The Generic Existence of Rational Expectations Equilibrium in the Higher Dimensional Case”, *Journal Economic Theory*, 26, 224 – 243.
- Karlin, S. and Taylor, H. (1981) *A Second Course in Stochastic Processes*, Academic Press, New York.
- Madhavan, A. (1992) “Trading Mechanisms in Securities Markets”, *Journal of Finance*, 47, 607 – 641.

- Milgrom, P. and Roberts, J. (1990) “Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities”, *Econometrica*, 58, 1255 – 1278.
- Monderer, D. and D. Samet (1989) “Approximating Common Knowledge with Common Beliefs”, *Games and Economic Behaviour*, 1, 170 – 190.
- Morris, S. (1992) “Acceptance Games and Protocols”, CARESS discussion paper #92 – 22, University of Pennsylvania.
- Morris, S., Postlewaite, A., and Shin, H. S. (1993) “Depth of Knowledge and the Effect of Higher Order Uncertainty”, CARESS discussion paper #93-15, University of Pennsylvania, forthcoming in *Economic Theory*.
- Morris, S., Rob, R., and Shin, H. S. (1995) “ p -Dominance and Belief Potential”, *Econometrica*, 63, 145 – 158.
- Pagano, M. and Röell, A. (1992a) “Auction and Dealership Markets: What is the Difference?”, *European Economic Review*, 613 – 623.
- Pagano, M. and Röell, A. (1992b) “Transparency and Liquidity: a Comparison of Auction and Dealer Markets with Informed Trading”, mimeo, London School of Economics.
- Pagano, M. and Röell, A. (1993) “Auction Markets, Dealership Markets and Execution Risk” in Vittorio Conti and Rony Hamoui (eds.) *Financial Markets’ Liberalization and the Role of Banks*, Cambridge University Press, 1993.
- Radner, R. (1968) “Competitive Equilibrium under Uncertainty”, *Econometrica*, 36, 31 – 58.
- Radner, R. (1979) “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices”, *Econometrica*, 47, 655 – 678.
- Shapley, L. and Shubik, M. (1977) “Trade Using One Commodity as a Means of Payment”, *Journal of Political Economy*, 85, 937 – 968.
- Rubinstein, A. (1989) “The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge” *American Economic Review*, 79, 385 – 391.
- Shin, H. S. and T. Williamson (1992) “How Much Common Belief is Necessary for a Convention?” CARESS discussion paper #92-32, University of Pennsylvania, forthcoming in *Games and Economic Behavior*.
- Sorin, S. (1993) “On the Impact of an Event”, mimeo, Ecole Normale Supérieure.

