Asset price manipulation with several traders*

Ansgar Walther, University of Cambridge†

Work in progress, October 2011

Abstract

In markets with asymmetric information, traders may have an incentive to forgo profitable deals today in order to preserve their informational advantage for future deals. This is established for monopolistic traders (Kyle 1985, Chakraborty and Yilmaz 2004). The effect is slower social learning and increased price volatility. Using an extension of Glosten and Milgrom's (1985) trading model, we investigate whether this effect persists when the number of traders increases. One might expect it to disappear as competition between traders decreases the manipulative power of an individual. However, we argue that there is strategic complementarity which may lead price manipulation to increase. Two contrasting conclusions emerge: On the one hand, in oligopolistic markets with a few traders, the strategic complementarity effect may dominate, leading to more price manipulation. On the other hand, in large markets, the competition effect dominates and price manipulation in equilibrium converges to zero. In particular, the effect of manipulation on price volatility is strongest for markets with two or three traders.

Keywords: Price manipulation, asset pricing, asymmetric information, Glosten-Milgrom model.

JEL classifications: D80, D82, G10, G14.

---

*Thanks are due to Hamid Sabourian for valuable research guidance, to Selma Telalagic for helpful comments and suggestions, and to the ESRC for financial support.

†ansgar.walther@googlemail.com

1 Introduction

When a trader receives a favourable tip about a stock, she will want to buy it. Buying a lot straight away would give a strong signal of her tip to the market and erode her informational advantage. Thus she may hold, or even sell a little, forgoing profitable deals today to preserve more of her advantage for future deals. This sort of behaviour could be described as price manipulation. If it prevails, it may considerably slow down learning and information revelation in financial markets with asymmetric information.

If the trader in question is the only trader dealing with an uninformed market maker, the incentive to manipulate prices is strong. This is the intuition behind the outcome of Kyle's (1985) repeated share auction. Also, Chakraborty and Yilmaz (2004) show that in the sequential trade model by Glosten and Milgrom (1985), a monopolistic informed trader may engage in such manipulative behaviour if the time horizon is long enough, giving him enough opportunities to profit in the future.

This paper asks whether this type of price manipulation can prevail when the number of informed traders in the market (N) increases. At first sight, it might seem obvious that the incentives to manipulate should decrease. After all, there is a 'competition effect': it only pays to manipulate when my impact on the market provides a strong enough incentive to forgo a profitable deal today. My market impact declines as more traders enter.

However, there might be an effect working the other way, based on *strategic complementarity.* Even in the presence of other traders, my market impact can be large when the actions of the others are not very informative. But the information content of the others' actions is reduced precisely when they use manipulative strategies. In this sense, more manipulation by others encourages me to manipulate more. Our results show that there is scope for this effect to dominate in the equilibrium of oligopolistic markets with intermediate values of N. The volatility of prices can increase significantly as N increases as manipulation by more traders introduces extra uncertainty. However, we also show that as N gets large, price manipulation in equilibrium disappears, casting some doubt on the conclusions of Chakraborty and Yilmaz (2004).

We study an extended version of the Glosten-Milgrom model, where N traders each get two chances to deal with competitive but uninformed market makers.

Either they are informed, and receive good or bad news, or they are noise traders. The market makers do not know what type they are dealing with. Consider a situation without price manipulation incentives. This would arise if we replaced our N traders by two identical generations of short-lived traders. In the case of good news, they would all buy in both periods to exploit their informational advantage over the market makers. Similarly, in the case of bad news, everyone would sell in both periods. Consequently, we define manipulative behaviour as a situation where informed traders choose not to take advantage of their information on both trading days.

Our main results are as follows: Firstly, we prove that as N grows large, we converge exponentially fast to a situtation where non-manipulative behaviour is impossible. This confirms and, perhaps surprisingly, strengthens the intuition in Chakraborty and Yilmaz (2004).

Secondly, we prove that it is possible for the strategic complementarity effect to dominate. For any choice of N, there is always an open subset of the parameter space such that there is more price manipulation in a market with $N+1$ traders than in a market with N traders.

Thirdly, we use simulations to show that equilibrium price manipulation converges to zero as N gets large. This insight mitigates our first result and the analysis of Chakraborty and Yilmaz (2004). Non-manipulative behaviour may be ruled out for large N, but this does not imply that the degree of manpulation actually observed in equilibrium is significant. It is also implied that our second result is only numerically significant for oligopolistic markets with small or intermediate values of N.

Finally, we use simulations to show the impact of manipulative behaviour in equilibrium on average prices and volatility. The impact on average prices and volatility is highest in oligopolistic markets with intermediate amounts of noise. Again, we find that the impact may increase in N as the strategic complementarity effect dominates. The results are particularly strong for the impact on volatility. In fact, volatility may increase in N *even if* equilibrium price manipulation does not.

The strategic complementarity effect is reminiscent of games of delayed investment (see Chamley 2004, Chapter 6). These games feature a group of agents with different private information about the profitability of a project, who decide whether to invest now or wait one period. Chamley points out that there can be strategic complementarity, as more delay by others raises the information content of their actions and encourages me to delay more. Note, however, that this strategic complementarity is driven by an informational externality (everybody wants to learn from each other), whereas in our model it results from a payoff externality (everybody wants to take advantage of the same uninformed players).

Price manipulation we study can slow down learning in financial markets. News gets absorbed into prices slowly when traders follow manipulative strategies. This has interesting implications for real world markets. For example, consider a situation where rational traders discover a stock price bubble. Abreu and Brunnermeier (2003) show that they may fail to coordinate, thus allowing the bubble to last for a while. However, once the rational traders do attack in their model, the bubble bursts instantly. Historical experience (Kindleberger and Aliber 2011) suggests that the deflation of bubbles is a drawn-out process that starts slowly and then accelerates, rather than an instant crash. Slow learning, as generated by our model, could play a role in explaining this discrepancy.

Of course, slow learning has been explained through many other channels. We refer the reader to Chamley (2004) and Brunnermeier (2001) for more comprehensive surveys. Most prominently, rational traders may 'follow the herd', ignoring their own signals and slowing down the revelation of information. The most famous illustration is Bikhchandani, Hirshleifer and Welch (1992). It was long believed that the herding effect is wiped out in a financial market with efficient pricing, such as the Glosten-Milgrom model, due to the negative results of Avery and Zemsky (1998). However, Park and Sabourian (2011) shows that herding may prevail in a Glosten-Milgrom model with a general signal structure under special circumstances, which are fully characterised in their paper. In our model, herding effects are switched off due to a binary signal structure. Allowing for herding effects and price manipulation incentives simultaneously would be desirable, but may not be tractable. Our analysis should be seen as complementary to herding models in explaining slow learning in financial markets.

The definition of price manipulation considered here is very weak compared to other definitions in the literature. Importantly, the manipulative strategies used by players in our model would not be considered unlawful (see Kyle (2008) for a discussion of the legal dimension). Also, they do not involve insider trading or the release of false information (see Allen and Gale (1992) for a review of papers on such strategies).

A related definition to ours is considered by Allen and Gale (1992). They show that uninformed traders may manipulate prices by emulating the behaviour of informed traders. The effect we examine is the opposite, and complementary for understanding information in markets. Whereas in their model, bogus information is 'learned' by the market, real information fails to be learned in ours.

The most closely related study to ours is Chakraborty and Yilmaz (2004), who also allow for repeated traders by informed trades in a Glosten-Milgrom framework. Their definition of manipulative behaviour is slightly stronger than ours. They only consider an agent to manipulate prices when he engages in loss-making deals, whereas here it is enough for him to forgo a profitable deal. Their main result rules out non-manipulative behaviour under certain conditions. Our results suggest that there is a problem with this approach, both in their monopoly model and in larger markets. After all, behaviour may be effectively non-manipulative in equilibrium, even though strictly non-manipulative behaviour can be ruled out. The stronger type of manipulation cannot occur in our model because we restrict the time horizon to two periods. We effectively accept a weakening of the definition for the sake of being able to characterise equilibrium play. We expect that similar results to ours would prevail with the stronger definition if the time horizon were longer, but such a model would quickly become hard to solve.

The paper is structured as follows. Section 2 sets up our extension of the standard Glosten-Milgrom framework, and defines the relevant equilibrium concept. Section 3 provides a general characterisation of symmetric equilibria. Section 4 shows the conditions under which non-manipulative strategies can prevail in equilibrium. Section 5 analyses the degree of price manipulation that is observed in equilibrium, and how it is affected by the number of traders. Section 6 illustrates the impact of manipulative behaviour in equilibrium on average prices and volatility. Section 7 concludes and outlines extensions currently in progress that will be included in future drafts of this paper.

2 The environment

Our market model generalises the classic sequential trade model by Glosten and Milgrom (1985) to a case where N traders each get two chances to deal in a financial market. There are many ways of achieving such a generalisation. In particular, one has to decide on the signal structure among traders, the sequence of moves, and the nature of the competitive price setting process. Studying the most general case makes the model untractable, so we make the following assumptions: Firstly, all traders receive the same signal. Secondly, there are two sequential trading days, but traders move simultaneously on each day. Finally, there is a pool of competitive market makers, each of whom can only deal with one trader per day. They get randomly matched to traders and have to quote uniform bid and ask prices before they know which trader they will deal with.

2.1 Traders and information

There are N traders $i \in \{1, 2, ..., N\}$ who are alive on two trading days, $t \in \{1, 2\}$. There is one risky asset with random liquidation value $V \in \{0, 1\}$, and $P[V = 1] = \pi$. Information is perfectly correlated across traders. With probability μ, the traders are all informed, and we denote this event by $I = 1$. Informed traders observe a signal $S \in \{0, 1\}$ with precision $q \in \left(\frac{1}{2}, 1\right)$, where $q = P[S = j|V = j]$ for $j \in \{0, 1\}$. With probability $1-\mu$, they are noise traders, who have no information. The random variables V, S, and I are mutually independent. All of the above is common knowledge.

We denote the probability of good news as $\gamma(\pi, q) = P[S = 1|I = 1] = \pi q + (1 - \pi)(1 - q)$. Note that $q > \gamma(\pi, q)$. We define the conditional probabilities of $V = 1$ upon observing S as

$$\pi_1 = P[V = 1|S = 1] = \frac{\pi q}{\gamma(\pi, q)}$$

$$\pi_0 = P[V = 1|S = 0] = \frac{\pi(1 - q)}{1 - \gamma(\pi, q)}$$

We have $\pi_1 > \pi > \pi_0$, since $S = 1$ is good news.

2.2 Market makers and price setting

In addition to the traders, there is a pool of $M > N$ competitive market makers. A market maker can only deal with one trader on each trading day. When dealing, traders and market makers know the history of past deals and prices H_t. Traders' orders, denoted a_t^i, are chosen strategically if the traders are informed, whereas the noise traders buy, hold or sell with equal probability, independently of each other.

On trading day t, the timing of interactions between traders and the market is as follows:

- Each market maker sets *competitive* ask and bid prices A_t and B_t.[1] We will explain shortly what is meant by competitive.
- Each trader is randomly matched to one market maker, so that the ex ante probability that market maker m deals with trader i is equal to M^{-1} for all i and m. Note that market makers have to set prices *before* they know who they will be matched to.
- Each trader submits an order to her market maker. She can either buy one unit of the asset at the ask price, sell one unit at the bid price, or hold.
- The prices and executed orders become common knowledge.
- Day $t + 1$ starts.

We will take the competitive price setting process as exogenously given, but it could be endogenised by introducing Bertrand competition between market makers.[2] We will abstract from strategic moves by market makers, and simply impose that for all t and all realisations of H_t, it is common knowledge among traders that the bid and ask prices quoted to all of them satisfy the condition

[1] Bear in mind that even though we simply write A_t and B_t, the prices will be random variables, depending on the history H_t. We do not make this dependence explicit to save on notation.

[2] For instance, consider a game where $M > 2N$, and each trader is randomly matched to *two* market makers and gets to decide which one to submit his order to, having observed their prices. The competitive price setting rules we conjecture below are an equilibrium, although there are also equilibria where some market makers do not observe these rules and get business with probability zero. The important point, however, is that any market maker who executes an order will follow the rules we conjecture.

that market maker M makes zero expected profits on both buy and sell orders, conditional on the history H_t. This translates to the following two conditions, which are simple generalisations of the zero profit conditions in Glosten and Milgrom (1985).

$$
\begin{aligned}
\sum_{i=1}^{N} \{ M^{-1} P \left[a_t^i = \text{buy at } A_t | H_t \right] & \\
\times \left[A_t - E \left[V | H_t, a_t^i = \text{buy at } A_t \right] \right] \} &= 0 \\
\sum_{i=1}^{N} \{ M^{-1} P \left[a_t^i = \text{sell at } B_t | H_t \right] &= 0 \\
\times \left[E \left[V | H_t, a_t^i = \text{sell at } B_t \right] - B_t \right] \} &= 0
\end{aligned}
$$

Implicit in these conditions is the notion that the conditional probabilities and expectations involved are obtained by Bayesian updating. Thus they depend on the action rules of both informed and noise traders, which are assumed to be known by market makers. This is because action rules conditional on H_t and the value of S influences the information content of a_t^i .

Rearranging, we have

$$
\begin{aligned}
A_t &= \sum_{i=1}^{N} \left\{ \frac{P \left[a_t^i = \text{buy at } A_t | H_t \right]}{\sum_{j=1}^{N} P \left[a_t^j = \text{buy at } A_t | H_t \right]} \right. && (1) \\
& \left. \times E \left[V | H_t, a_t^i = \text{buy at } A_t \right] \right\} && (2) \\
B_t &= \sum_{i=1}^{N} \left\{ \frac{P \left[a_t^i = \text{sell at } B_t | H_t \right]}{\sum_{j=1}^{N} P \left[a_t^j = \text{sell at } B_t | H_t \right]} \right. && (3) \\
& \left. \times E \left[V | H_t, a_t^i = \text{sell at } B_t \right] \right\} && (4)
\end{aligned}
$$

Note also that each history occurs with strictly positive probability, regardless of the action rules of informed traders, due to the presence of noise traders. Hence, there will be no need to specify out-of-equilibrium beliefs of market makers. The competitive bid and ask prices can always be calculated using Bayes' rule.

We assume that it is common knowledge that the competitive price setting process satisfies (1) and (3) for all t and all realisations of the history H_t. This way, the history of orders $\left(a_t^i \text{ for all } i,\, s < t \right)$ is a sufficient statistic for the information contained in the history H_t, and we will treat the two as equivalent when constructing conditional probabilities and expectations.

In the analysis below, we will focus on traders who are ex ante identical and hence we look for symmetric equilibria where all informed agents follow the same action rule. In this special case, $P\left[a_t^i = \text{buy at } A_t | H_t\right]$ and $E\left[V | H_t, a_t^i = \text{buy at } A_t\right]$ are the same for all i, and the expressions in (1) and (3) reduce to

$$\begin{aligned} A_t &= E\left[V | H_t, a_t^i = \text{buy at } A_t\right] \\ B_t &= E\left[V | H_t, a_t^i = \text{sell at } B_t\right] \end{aligned}$$

2.3 Perfect Bayesian Equilibrium

Since we take the market making and price setting process as exogenously given, the only people who act strategically are informed traders. A *strategy* for an informed trader i has to specify a market order for all t and all realisations of the history H_t and his signal S. The market order can be a random variable if i plays a mixed strategy. A *strategy profile* specifies such market order rules for all traders. A *Perfect Bayesian Equilibrium (PBE)* is a strategy profile such that no trader has an incentive to deviate from her strategy given the others' strategies and the competitive price setting described by (1) and (3). These prices are using Bayesian updating in a way that is consistent with the strategy profile. Because information is perfectly correlated across traders in our setting, there is no need to specify i's beliefs about the other traders' types.

Before introducing formal notation for PBE, we derive necessary conditions that simplify things. First of all, it is useful to observe that market makers will never become as optimistic as a trader with good news, or as pessimistic as a trader with bad news. This is because they always attribute some of their observations to noise.

Lemma 1. *For all t and all realisations of H_t, the competitive prices satisfy $\pi_0 < A_t, B_t < \pi_1$, regardless of the strategies of informed traders.*

The second condition states that at $t = 2$, traders with good news will buy and traders with bad news will sell for sure. This follows by backward induction. Since their order at $t = 2$ is their last move in the game, traders do not consider anything except for the quoted prices relative to their posterior. The condition is then immediate from Lemma 1.

Lemma 2. *Take any PBE. At $t = 2$, all informed traders buy with probability 1 if $S = 1$, and sell with probability 1 if $S = 0$.*

The third condition states that at $t = 1$, traders with good news will never sell, and traders with bad news will never buy. One could imagine a situation where a trader with good news sells today in order to depress future prices and get an even better deal tomorrow. Chakraborty and Yilmaz (2004) show that with a monopolist trader, this happens when the time horizon is sufficiently long. The following proof shows that in our setting with two periods, the price manipulation incentive can never be strong enough for this to happen.

Lemma 3. *Take any PBE. At $t = 1$, all informed traders sell with probability 0 if $S = 1$, and buy with probability 0 if $S = 0$.*

Proof. We show the result for $S = 1$. The proof for $S = 0$ is essentially the same. Suppose that at $t = 1$, trader i sells with probability $\varepsilon > 0$, holds with probability α and buys with probability $1 - \alpha - \varepsilon$. Since he will buy at $t = 1$ by Lemma 2, his expected payoff is

$$\begin{aligned} \varepsilon \left[B_1 - E\left[A_2 | S = 1, a_1^i = \text{sell}\right]\right] \\ +\alpha \left[\pi_1 - E\left[A_2 | S = 1, a_1^i = \text{hold}\right]\right] \\ +(1 - \alpha - \varepsilon)\left[\pi_1 - A_1 + \pi_1 - E\left[A_2 | S = 1, a_1^i = \text{buy}\right]\right] \end{aligned}$$

Suppose he deviates by buying with probability $1 - \alpha$ and holding with probability α. The net expected profit from this deviation is ε times

$$2\pi_1 - (A_1 - B_1) - \left[E\left[A_2 | S = 1, a_1^i = \text{sell}\right] - E\left[A_2 | S = 1, a_1^i = \text{buy}\right]\right]$$

By Lemma 1, the second and third terms in this expression are strictly less than $\pi_1 - \pi_0$. Hence the deviation is profitable, contradicting PBE. □

It follows that any PBE can be described by a profile $(\alpha_i, \beta_i)_{i=1}^N$, where α_i is the probability that an informed trader i holds when $S = 1$, and β_i is the probability that informed trader i holds when $S = 0$.

Since agents are ex ante identical and receive the same information, it is natural to look primarily for symmetric PBE.[3] A symmetric PBE is a profile (α, β) such

[3] One might even be able to rule out non-symmetric PBE in some circumstances, but we have not yet achieved a formal proof.

that $(\alpha_i, \beta_i) = (\alpha, \beta)$ for all i is a PBE.

3 Characterisation of symmetric PBE

Let (α, β) be a symmetric strategy profile. Competitive prices will depend on α and β, as the strategies influence the information content of traders' past moves. Note, however, that ask prices never depend on β since players with bad news never buy, and similarly bid prices never depend on α. We make the dependence explicit by writing $A_t(\alpha)$ and $B_t(\beta)$.

We first consider the incentives of an informed agent i with good news ($S = 1$) at $t = 1$. Suppose that all such agents' action rules say to hold with probability α and buy with probability $1-\alpha$ at $t = 1$, and the market makers' price setting is based on these rules. If agent i holds with probability α, his expected payoff is

$$\begin{array}{r} \alpha\left[\pi_1 - E\left[A_2(\alpha)\,|S=1, a_1^i = \text{hold}\right]\right] \\ (1-\alpha)\left[\pi_1 - A_1(\alpha) + \pi_1 - E\left[A_2(\alpha)\,|S=1, a_1^i = \text{buy}\right]\right] \end{array}$$

Consider a deviation from i's action rule where he hold slightly more frequently, i.e. with probability $\alpha+\varepsilon$, for $\varepsilon > 0$. Let $G_N(\alpha)$ denote the marginal gain from such a deviation. We have

$$\begin{aligned} G_N(\alpha) &= E\left[A_2(\alpha)\,|S=1, a_1^i = \text{buy}\right] - E\left[A_2(\alpha)\,|S=1, a_1^i = \text{hold}\right] \\ &- \left[\pi_1 - A_1(\alpha)\right] \end{aligned} \quad (5)$$

Intuitively, the first term in G_N is the incentive to manipulate prices by holding more frequently and therefore lowering the market makers' posterior about the fundamental value. The second term is the incentive to get a profitable deal today. If $G_N(\alpha) > 0$, then informed traders with good news will want to deviate from α by holding more frequently as the price manipulation incentive dominates. If $G_N(\alpha) < 0$, they will want to buy more frequently, as the incentive to deal today is stronger. Note that the expectations in G_N do not depend on β, because agent i knows that nobody has bad news.

We can immediately observe the following: If $\alpha > \frac{1}{2}$, then a hold at $t = 1$ is better news for the market maker than a buy when combined with a buy at

$t = 2$, as is it is a stronger indication of $S = 1$. Hence, the price manipulation incentive will actually be negative, since by holding more frequently the informed trader sends a positive signal to the market maker. Similarly, for $\alpha = \frac{1}{2}$, the price manipulation incentive disappears altogether. This, combined with the incentive to deal today, implies that $G_N(\alpha) < 0$ for all $\alpha \geq \frac{1}{2}$.

Similarly, consider an informed agent with bad news ($S = 0$) at $t = 1$. Let $H_N(\beta)$ denote the marginal gain he would derive from holding slightly more frequently. We have

$$\begin{aligned} H_N(\beta) &= E\left[B_2(\beta)\,|S = 0, a_1^i = \text{hold}\right] - E\left[B_2(\beta)\,|S = 0, a_1^i = \text{sell}\right] \\ &- \left[B_1(\beta) - \pi_0\right] \end{aligned}$$

Repeating our previous argument for $H_N(\beta)$ yields the following:

Lemma 4. *$G_N(\alpha) < 0$ for all $\alpha \geq \frac{1}{2}$, and $H_N(\beta) < 0$ for all $\beta \geq \frac{1}{2}$.*

We know that (α, β) is a symmetric PBE if and only if neither type of informed agent has an incentive to deviate from it. In particular, we cannot have any informed agents holding with a greater probability than $\hat{p} \geq \frac{1}{2}$ in a symmetric PBE. This is because $G_N(\hat{p}) < 0$ and $H_N(\hat{p}) < 0$ imply that they would want to deviate by buying more frequently in the case of good news, or by selling more frequently in the case of bad news. This helps us to narrow down the description of PBE.

Lemma 5. *(α, β) is a symmetric PBE if and only if two conditions hold:*

1. Either $\alpha = 0$ and $G_N(0) \leq 0$, or $\alpha \in \left(0, \frac{1}{2}\right)$ and $G_N(\alpha) = 0$.
2. Either $\beta = 0$ and $H_N(0) \leq 0$, or $\beta \in \left(0, \frac{1}{2}\right)$ and $H_N(\beta) = 0$.

We can now state the following general existence result:

Proposition 1. *There always exists a symmetric PBE.*

Proof. First, we need to find an α satisfying condition 1 in Lemma 5. If $G_N(0) \leq 0$ then $\alpha = 0$ satisfies the condition. If $G_N(0) > 0$. Lemma 4 and the fact that $G_N(\alpha)$ is continuous in α imply that there exists an $\alpha \in \left(0, \frac{1}{2}\right)$ with $G_N(\alpha) = 0$, satisfying the condition. The argument for β is the same. □

The only thing that is lacking to extend the existence result of Proposition 1 to an "existence and uniqueness" result with a general characterisation of when price manipulation occurs would be the fact that $G_N(\alpha)$ is strictly decreasing in α for $\alpha \in \left(0, \frac{1}{2}\right)$, and that $H_N(\beta)$ is strictly decreasing in β for $\beta \in \left(0, \frac{1}{2}\right)$. We have been able to prove this for the case of one or two traders, which allows us to fully characterise the symmetric PBE. See Appendix A for these results. We also have a counterexample showing that $G_N(\alpha)$ can be increasing for α for large N. However, we have not encountered multiple equilibria in any of our simulations in subsequent sections. It appears that there may be a general uniqueness result not driven by monotonicity of $G_N(\alpha)$. We are currently working on establishing this, and a more thorough analysis will be in future drafts of the paper.

It helps intuition to derive a more explicit expression for the expectation terms in $G_N(\alpha)$. Since agents have the same information and strategies as him, i knows that H_2 will consist of buys and holds, and that $A_2(\alpha)$ will depend purely on the number of holds (rather than buys) at $t = 1$. For convenience, we denote this number by K, and simply write the second-period ask price as $\hat{A}(K, \alpha)$. The market makers will set

$$A_1(\alpha) = \pi \frac{\mu q (1-\alpha) + \frac{1-\mu}{3}}{\mu\gamma(\pi, q)(1-\alpha) + \frac{1-\mu}{3}} \tag{6}$$

$$\hat{A}(K, \alpha) = \pi \frac{\mu q \alpha^K (1-\alpha)^{N-K} + \frac{1-\mu}{3^{N+1}}}{\mu\gamma \alpha^K (1-\alpha)^{N-K} + \frac{1-\mu}{3^{N+1}}}, \text{ for } K \in \{0, 1, .., N\} \tag{7}$$

Note the following symmetry property: $\hat{A}(N-K, a) = \hat{A}(K, 1-\alpha)$.

From his perspective, the number of holds from other agents, denoted K_{-i}, is a binomial random variable with $N-1$ trials and success probability $1-\alpha$, and we have

$$\begin{aligned} P[K_{-i} = k | S = 1] &= \binom{N-1}{k} \alpha^k (1-\alpha)^{N-1-k} \\ &\equiv p(N-1, \alpha, k), \text{ for } k \in \{0, 1, ..., N-1\} \end{aligned} \tag{8}$$

where $\binom{N-1}{k} = \frac{(N-1)!}{k!(N-1-k)!}$ is the binomial coefficient. It is also useful to recall the following symmetry property of the binomial distribution: the binomial coefficient satisfies $\binom{N-1}{k} = \binom{N-1}{N-1-k}$, and hence we have

$p\,(N-1,\alpha,N-1-k) = p\,(N-1,1-\alpha,k)$.

Now let $g\,(\alpha) = \sum_{k=0}^{N-1} p\,(N-1,\alpha,k)\,\hat{A}\,(k,\alpha)$. The first two terms in $G\,(\alpha)$, as defined in Equation 5, can be written as

$$\begin{aligned}
E\left[A_2^{(\alpha)}\,(H_2)\,|S=1,a_1^i=\text{buy}\right] &= g\,(\alpha) \\
E\left[A_2^{(\alpha)}\,(H_2)\,|S=1,a_1^i=\text{hold}\right] &= \sum_{k=0}^{N-1} p\,(N-1,\alpha,k)\,\hat{A}\,(k+1,\alpha) \\
&= \sum_{k=0}^{N-1} p\,(N-1,\alpha,N-1-k)\,\hat{A}\,(N-k,\alpha) \\
&= \sum_{k=0}^{N-1} p\,(N-1,1-\alpha,k)\,\hat{A}\,(k,1-\alpha) = g\,(1-\alpha)
\end{aligned}$$

Similar expressions can be obtained to characterise $H_N\,(\beta)$ explicitly. In particular, it can be shown that $H_N\,(\beta) = G_N\,(\beta)\,|_{q=1-q}$.

4 Non-manipulative equilibria

Consider a symmetric PBE where in the case of good news, all N players buy with probability 1 at $t=1$ in equilibrium ($\alpha=0$). Then, traders take every opportunity to exploit their informational advantage. We will call this a *non-manipulative strategy.* We now examine how the incentives to play non-manipulative strategies in symmetric PBE are linked to the parameters of the model, in particular the amount of the information in the market μ and the number of traders N.

More information in the market (large μ) means that the market makers attribute less of their observations to noise when setting prices. This has two effects on the incentives to deviate from non-manipulation. Firstly, the impact of traders' actions today on tomorrow's prices increases, increasing the price manipulation incentive. Secondly, the ask price today increases, reducing the incentive to trade today. Both effects work against non-manipulative strategies.

Secondly, one might expect that large markets (large N) reduce the price manipulation incentive, as competition among traders reduces the price impact of an individual trader and hence his incentive to manipulate prices. However, more traders in the market mean that the information content of actions at

$t = 1$ is higher. This boosts the price manipulation incentive, while leaving the incentive to trade today unchanged. This intuition translates into the following result.

Proposition 2. *There exists a symmetric PBE where traders with good news play non-manipulative strategies if and only if* $\mu \leq \bar{\mu}_N(\pi, q)$, *where*

$$\bar{\mu}_N(\pi, q) = \frac{1}{1 + 3^{1+\frac{N}{2}}\gamma(\pi, q)} \tag{9}$$

'

Furthermore, there exists a symmetric PBE where traders with bad news play non-manipulative strategies if and only if $\mu \leq \bar{\mu}_N(\pi, 1-q)$.

Proof. Consider the case of good news. By Lemma 5, there exists a symmetric PBE with $\alpha = 0$ if and only if $G(0) \leq 0$. Using our characterisation of $G(\alpha)$ above, we obtain

$$G(0) = \pi \left[\frac{\mu q + \frac{1-\mu}{3^{N+1}}}{\mu\gamma(\pi, q) + \frac{1-\mu}{3^{N+1}}} + \frac{\mu q + \frac{1-\mu}{3}}{\mu\gamma(\pi, q) + \frac{1-\mu}{3}} - \frac{q}{\gamma(\pi, q)} - 1 \right]$$

This expression is strictly positive when $\mu = 1$, strictly negative when $\mu = 0$, and strictly increasing in μ. It follows that $G(0) \leq 0$ if and only if $\mu \leq \bar{\mu}(\pi, q)$, where $\bar{\mu}(\pi, q)$ is the unique number in $(0, 1)$ that solves $G(0)|_{\mu=\bar{\mu}(\pi,q)} = 0$. Solving the resulting quadratic equation in μ and eliminating the root outside of $(0, 1)$ yields the expression in (9).

There exists a symmetric PBE with $\beta = 0$ if and only if $H(0) = G(0)|_{q=1-q} \leq 0$, which is the case if and only if $\mu \leq \bar{\mu}(\pi, 1-q)$. □

In line with our intuition, non-manipulative behaviour is more likely when there is not a lot of information (low μ). Also, as expected, non-manipulative behaviour becomes less likely in large markets, since the information threshold $\bar{\mu}$ is decreasing in N. However, it is surprising quite how powerful this effect is: $\bar{\mu}$ converges to zero at the rate $3^{-\frac{N}{2}}$. Thus as the market size increases, we converge exponentially fast to a situation where non-manipulative strategies are impossible in equilibrium, regardless of the other parameters.

Figure 1 illustrates the condition in Proposition 2. A symmetric PBE where traders with good news play non-manipulative strategies exists if and only if the

Figure 1: The condition for non-manipulative behaviour

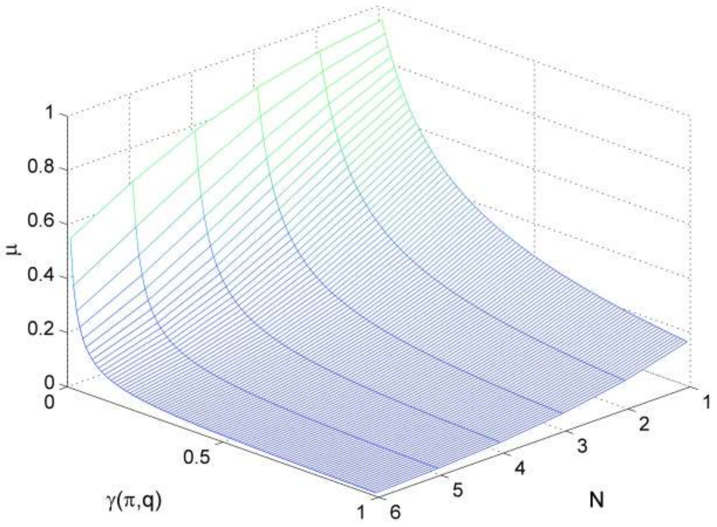

parameters of the model lie below the surface plotted. The graph emphasises the drastic effect of an increase in market size - as soon as N exceeds 2, there is only a very small portion of the parameter space in which non-manipulation equilibria exist.

5 Manipulative equilibria

We have established that as the number of traders N increases, we are unlikely to observe a situation where informed agents take every opportunity to trade on their information in equilibrium (i.e. non-manipulative strategies). It follows that there is often some degree of price manipulation by informed agents in equilibrium. The parameters of their strategies, α and β, intuitively capture the degree of price manipulation because they are the probabilities with which the agents forgo a profitable deal based on their informational advantage at $t = 1$. From now on, we will focus on the strategies of informed traders with good news, characterising the values of α that are observed in equilibrium. Parallel results would apply to bad news and β.

Hence, it is interesting to examine how the equilibrium values of α changes as the number of traders increases. As discussed in the introduction, the direction of the effect depends on the balance between the 'competition effect' of the reduced price impact of an individual trader when N increases and the strategic complementarity effect of an increase in N. We let α_N and denote the minimum value of α and that may be part of a symmetric PBE[4].

Let us examine whether there is scope for the strategic complementarity to dominate. We now show that for each N, and holding the parameters π and q fixed, there is an interval of value of μ for which the degree of price manipulation in equilibrium would be raised by adding another trader, or $\alpha_N < \alpha_{N+1}$. This interval can be decomposed into two sections. Firstly, the interval $(\bar{\mu}_{N+1}(\pi, q), \bar{\mu}_N(\pi, q)]$, i.e. the region where μ is small enough to allow non-manipulative behaviour of N traders ($\alpha_N = 0$), but not small enough to do so for $N+1$ traders ($\alpha_{N+1} > 0$). Hence, increasing the number of traders increases these measures of price manipulation on a subset of the parameter space with positive Lebesgue measure, and we can state the following:

Proposition 3. *For any N, there exists an open subset of the parameter space such that $\alpha_N < \alpha_{N+1}$. There also exists an open subset of the parameter space such that $\beta_N < \beta_{N+1}$.*

We have demonstrated that the strategic complementarity effect dominates for some parameter constellations. However, we have not been able to fully characterise when this is the case, or to quantify the effect. This is because as soon as $N > 1$, the function $G(\alpha)$ becomes highly non-linear in α and we can no longer explicitly solve the equation $G(\alpha_N) = 0$ to find α_N.[5] We therefore use simulations to further investigate the implications of Proposition 3. In particular, we hold the value of the parameters π and q fixed and examine how α_N changes with the market size N for different values of μ.

Figure 2 below presents the results for $\pi = \frac{1}{2}$ and $q = \frac{3}{4}$. We let μ be anywhere between zero and one, and look at the equilibrium outcomes for N between

[4] So far, we do not have a general result for uniqueness of the symmetric PBE. Thus we cannot speak of "the equilibrium value" of α. As mentioned in Section 3, we expect a uniqueness result to hold, and have not come across non-uniqueness in any of our simulations. This is why, for now, we focus on the minimum values of α, but refer to it as "the equilibrium value" in a slight abuse of terminology.

[5] From now on, we focus on the equilibrium strategies of informed traders with good news only, that is we examine α_N and ignore β_N. All of the analysis below could of course be conducted in a parallel fashion for β_N.

Figure 2: The degree of price manipulation in equilibrium as a function of μ

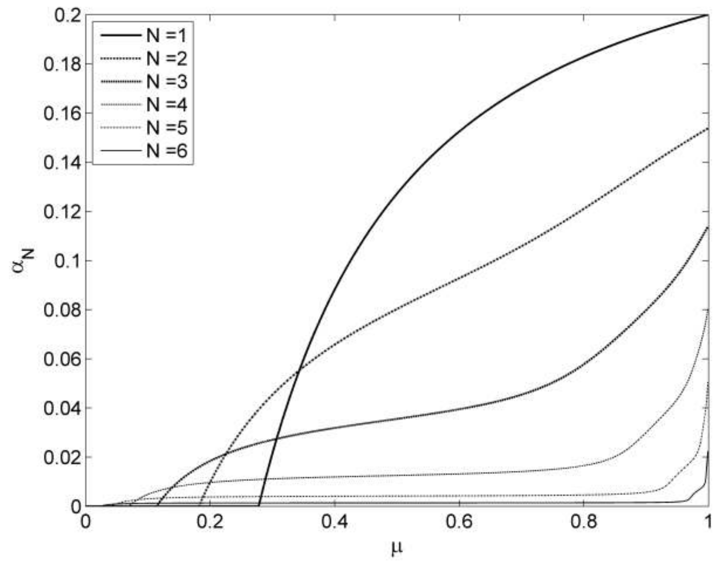

one and six. As implied by Proposition 2, the degree of price manipulation in equilibrium is zero for small values of μ, when $\mu \leq \bar{\mu}_N(\pi, q)$. For the region $\mu > \bar{\mu}_N(\pi, q)$, the degree of price manipulation appears to be strictly increasing in μ. This is intuitive - as μ gets larger, the incentive to manipulate prices increases because, other things being equal, the price impact of each trader increases. Also, the ask price today increases, thus reducing the incentive to deal today.

As for the impact of market size, the simulations show that for low values of μ, α_N may increase in N, as suggested by Proposition 3. Figure 3 'zooms in' on our simulated results for low values of μ to examine this issue. It appears that for each N, there is a significant interval of values of μ for which $\alpha_N < \alpha_{N+1}$.

However, the simulations also show that despite this effect, α_N quickly becomes small numerically as N increases, and in fact converges to zero as N gets large. The effect suggested by Proposition 3 is only numerically important for small N, and the 'competition effect' is, after all, quite powerful. We further analyse this aspect by plotting α_N as a function of N for different values of μ, as in Figure 4. It is interesting to note that the rate of convergence is considerably slower for values of μ close to 1, that is in markets with a lot of information.

Figure 3: The degree of price manipulation in equilibrium as a function of small μ

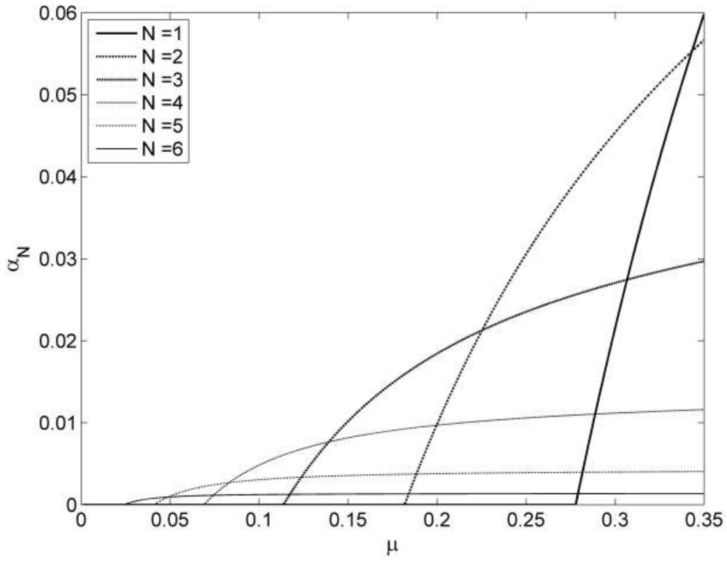

Figure 4: The degree of price manipulation in equilibrium as a function of N

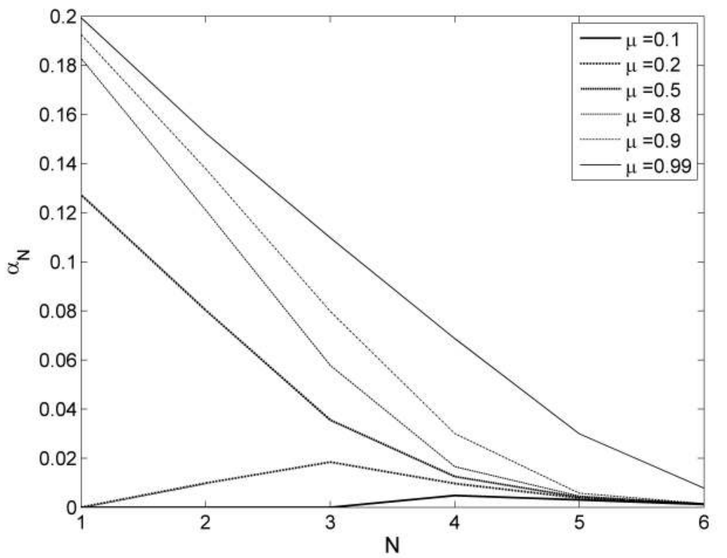

Im summary, two main results emerge about the relationship between the number of traders N and the degree of price manipulation in equilibrium α_N. Firstly, in small markets with a lot of noise (low μ), increasing the number of traders may actually increase α_N slightly, due to a strategic complementarity effect. Secondly, as N gets larger, this effect becomes less important, and α_N converges to zero. However, the convergence is slower than one might think if there is a lot of information in the market (high μ).

The fact that α_N converges to zero for large N stands in contrast to the conclusions of Section 4. There, we showed that non-manipulative behaviour rapidly becomes less likely for large N. This might have lead one to conclude that significant price manipulation persists for large N. The analysis in this section points out that in fact, we converge to equilibria that are *effectively* non-manipulative, where agents only forgo profitable deals with very small probabilities. The main result in Chakraborty and Yilmaz (2004) rules out non-manipulative behaviour, according to their slightly stronger definition of manipulation (see the introduction for details), in a very similar model to ours with a longer timer horizon. They demonstrate that their argument goes through for any N, in a manner that is very similar to our arguments in Section 4. Our analysis here casts some doubt on the power of this result for large N. It is to be suspected that if we were able to calculate the actual degree of price manipulation in equilibrium in their model, it would be found to also become insignificant for larger markets. Whether

6 Manipulation, prices, learning and volatility

Having analysed the degree of price manipulation in equilibrium, we would now like to know how this affects the behaviour of observed prices. Since in Glosten-Milgrom type models such as ours, prices are informationally efficient (in the sense that they reflect all public information), this analysis will also shed light on how price manipulation affects the speed of market learning.

We analyse the impact of price manipulation by considering a situation where good news have been received by informed traders ($S = 1$). Conditional on this event, we calculate the impact on the expectation and the volatility (i.e. the standard deviation) of transaction prices at $t = 1$ and $t = 2$.

Conditional on good news, everyone who conducts a transaction will buy at

$t = 1$, and the remaining agents will hold. Hence, the average transaction price at $t = 1$ is simply $A_1(\alpha_N)$ as defined in Equation (6). It is deterministic so that its volatility is always zero.

The average transaction price at $t = 2$, on the other hand, is a random variable conditional on good news. This randomness is introduced by agents playing mixed strategies in the first period due to the incentives to manipulate prices. Its expectation is defined as

$$\begin{aligned}\bar{A}_2(\alpha_N) &= E[A_2|S=1] \\ &= \sum_{k=0}^{N} p(N, \alpha_N, k)\, \hat{A}(k, \alpha_N)\end{aligned}$$

where $\hat{A}(k, \alpha_N)$ is the ask price conditional on k holds and $N-k$ buys as defined in Equation (7), and p is the binomial probability density function defined in Equation (8). The volatility of second-period transaction prices conditional on good news is defined as

$$\sigma_2(\alpha_N) = \sqrt{\sum_{k=0}^{N} p(N, \alpha_N, k)\left[\hat{A}(k, \alpha_N) - \bar{A}_2(\alpha_N)\right]^2}$$

Note that if equilibrium play is non-manipulative ($\alpha_N = 0$), which happens under the conditions derived in Section 4, the average second period transation price is deterministic. Then, its expectation and equal to $\hat{A}(0, 0)$ and its volatility is zero.

Again, since we have no analytical solution for α_N, we resort to simulated results to analyse price and volatility impact of manipulative behaviour. As in Section 5, we conduct the simulations by holding π and q fixed and varying μ and the market size N. All results below are based on $\pi = \frac{1}{2}$ and $q = \frac{3}{4}$.

We will analyse the price impact of manipulation incentives by comparing equilibrium prices to benchmark values. Useful benchmarks are the prices that would prevail under non-manipulative behaviour, i.e. $A_1(0)$ and $\bar{A}_2(0) = \hat{A}(0, 0)$. Recall that these prices would be equilibrium prices if we removed price manipulation incentives, for instance by replacing our traders with two generations of short-lived trader. Below, we report results for the impact of manipulative behaviour on prices in percentage terms. These figures are calculated as

Figure 5: The impact of manipulative behaviour on A_1 (in percentage terms)

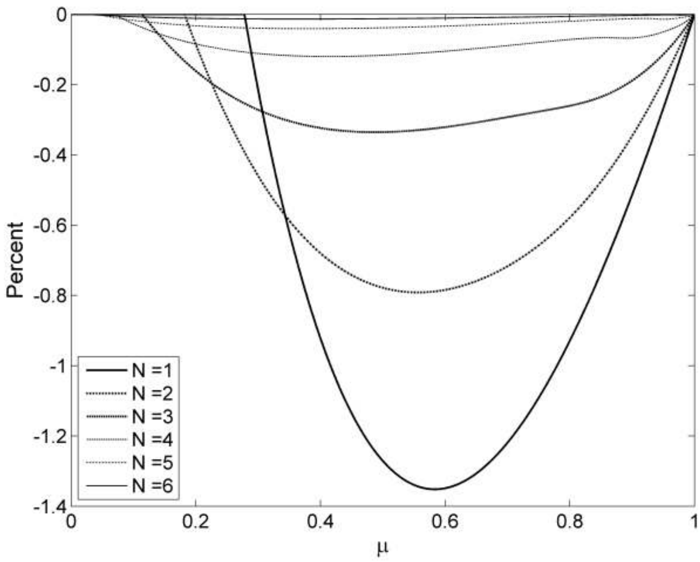

$\left(\frac{A_1(\alpha_N)}{A_1(0)} - 1\right) \times 100$ and $\left(\frac{\bar{A}_2(\alpha_N)}{\bar{A}_2(0)} - 1\right) \times 100$ respectively.

Figure 5 shows the impact on first-period transaction prices. The effect is moderate - price manipulation reduces A_1 by a maximum of about 1% - and fairly insignificant for markets with more than two traders.

The effect on average second period prices is more pronounced, as shown in Figure 6. The price manipulation incentive can lower the average second period price by over 2.5%. Proposition 3, we established that under certain circumstances, increasing the number of traders may increase the degree of price manipulation in equilibrium due to a strategic complementarity effect. This is reflected in the price. In particular, the strongest price impact of the manipulation incentive is not achieved for the case of a monopolistic trader, but for $N = 2$ and intermediate values of μ, around 0.4.

On the other hand, as already suggested by our simulation results in Section 5, manipulation incentives quickly become less frequent as N gets large, which is reflected in the price impact. For instance, price manipulation never has a price impact of more than half a percent when there are more than three traders in the market.

Figure 6: The impact of manipulative behaviour on $\bar{A}_2$ (in percentage terms)

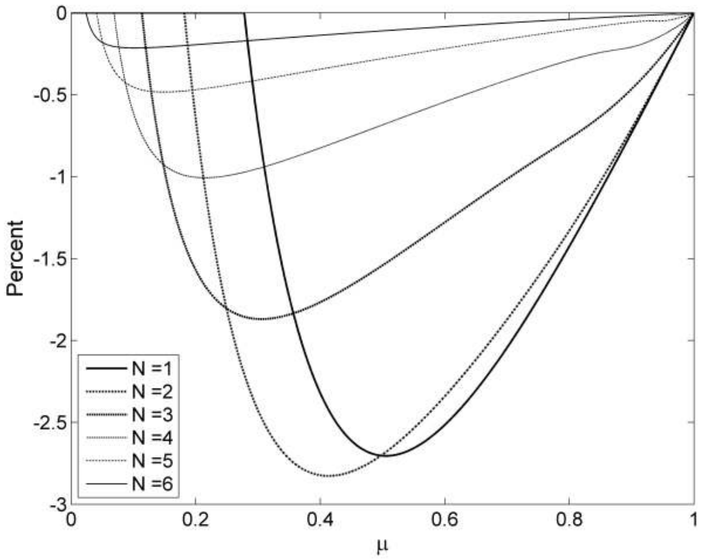

The effect of price manipulation in equilibrium on volatility is depicted in Figure 7. It appears that manipulative behaviour can introduce a standard deviation of almost 0.05 to second-period prices. Given that second-period prices lie between $\frac{1}{2}$ and $\frac{3}{4}$, this is a significant volatility component. When another trader is added to the market, the volatility impact of manipulative behaviour often increases sharply. For most values of μ the strongest volatility impact is felt in markets with two or three traders.

Note also that even for values of μ where the degree of price manipulation in equilibrium, α_N, and the average price impact of this behaviour are highest in markets with a monopolistic trader (e.g. $\mu = 0.4$, see Figures 2 and 6), the volatility impact is highest for larger markets.

This seems surprising, but can be explained intuitively. Consider a situation where we move from N to $N+1$ traders, and this lowers the degree of price manipulation in equilibrium from α_N to $\alpha_{N+1} < \alpha_N$. Even though manipulative behaviour becomes less frequent, we have introduced 'tail events'. For instance, if all $N+1$ traders hold, the second-period price will be lower than it could ever have been with N traders. To see this, note that firstly, holds are now worse news because traders with good news hold less frequently, and that secondly,

Figure 7: The impact of manipulative behaviour on volatility σ_2

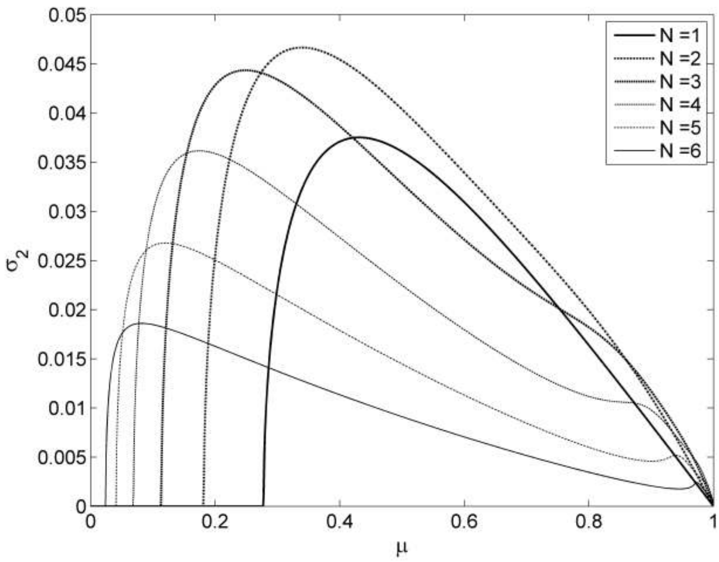

the actions of $N + 1$ traders contain more information than those of N traders, so that the price impact of bad news is more pronounced. Furthermore, by a similar logic, if all $N + 1$ traders buy, the second-period price will be higher than it could have been with N traders. Hence, moving from N to $N + 1$ traders increases the dispersion of second-period prices and may thus increase their volatility, even though manipulative behaviour becomes less common.

7 Preliminary Conclusion

Our results so far show that price manipulation incentives may be strengthened, not weakened, by the presence of more traders. In particular, the impact of price manipulation on volatility is most potent in oligopolistic markets. The argument that price manipulation is irrelevant due to competition between traders remains valid in large markets. However, in markets where a few traders have a significant joint market share, it should not be accepted unconditionally. Price manipulation incentives may well slow down learning and introduce extra uncertainty in financial markets in general, not only in the special case of one monopolistic trader.

A few important extensions to our analysis will be considered in future drafts of this paper. Firstly, many of the simulated results about the degree of price manipulation in equilibrium and the impact on prices and volatility in Sections 5 and 6 are likely to have analytical counterparts, which are currently being developed. In particular, we aim to include analytical descriptions of the asymptotic behaviour of equilibrium as N gets large and the rates of convergence in question. Secondly, we are currently working on checking the robustness of our result to slight changes in the assumptions about market setup. For instance, it is important to investigate how our results are affected by letting agents' signals imperfectly correlated, or by letting them move sequentially. We are unlikely to be able to derive analytical results as neat as ours under these extensions, but we aim to show whether they change our conclusions qualitatively, by means of numerical simulations if necessary. Finally, as discussed in Section 3, we are still working on a characterisation of when the symmetric PBE we analyse is unique.

References

ABREU, D., AND M. BRUNNERMEIER (2003): "Bubbles and Crashes," *Econometrica*, 71(1), 173–204.

ALLEN, F., AND D. GALE (1992): "Stock-Price Manipulation," *Review of Financial Studies*, 3, 503–529.

AVERY, C., AND P. ZEMSKY (1998): "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

BRUNNERMEIER, M. (2001): *Asset Pricing Under Asymmetric Information.* Oxford University Press.

CHAKRABORTY, A., AND B. YILMAZ (2004): "Informed manipulation," *Journal of Economic Theory*, 114, 132–152.

CHAMLEY, C. (2004): *Rational Herds.* Cambridge University Press.

Glosten, L., and P. Milgrom (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 14, 71–100.

Kindleberger, C., and R. Aliber (2011): *Manias, Panics and Crashes.* Palgrave, 6th edn.

Kyle, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

——— (2008): "How to Define Illegal Price Manipulation," *American Economic Review: Papers and Proceedings*, 98, 274–279.

Park, A., and H. Sabourian (2011): "Herding and contrarian behaviour in financial markets," *Econometrica*, 79, 973–1026.

Appendix

A Symmetric PBE with $N \in \{1, 2\}$

A.1 Properties of the functions $G(\alpha)$ and $B(\beta)$

Lemma. *For $N \in \{1, 2\}$, $G(\alpha)$ is strictly decreasing in α and $B(\beta)$ is strictly decreasing in β.*

Proof. We show this for $G(\alpha)$. The statement for $B(\beta)$ follows immediately from the fact that $B(\beta) = G(\alpha)|_{q=1-q}$.

Case (i): $N = 1$. Using the characterisation of $G(\alpha)$ from Section 3 gives

$$G(\alpha) = \pi \left[\frac{\mu q (1-\alpha) + \frac{1-\mu}{9}}{\mu\gamma(1-\alpha) + \frac{1-\mu}{9}} - \frac{\mu q \alpha + \frac{1-\mu}{9}}{\mu\gamma\alpha + \frac{1-\mu}{9}} - \frac{q}{\gamma} + \frac{\mu q (1-\alpha) + \frac{1-\mu}{3}}{\mu\gamma(1-\alpha) + \frac{1-\mu}{3}} \right]$$

which is decreasing in α since $q > \gamma$. We write γ instead of $\gamma(\pi, q)$ to save space.

Case (ii): $N = 2$. We know that

$$G'(\alpha) = g'(\alpha) + g'(1-\alpha) + A_1'(\alpha)$$

where

$$g(\alpha) \quad = \quad \alpha \hat{A}(1,\alpha) + (1-\alpha)\hat{A}(0,\alpha)$$

and

$$\hat{A}(K,\alpha) = \pi \frac{\mu q \alpha^K (1-\alpha)^{2-K} + \frac{1-\mu}{27}}{\mu \gamma \alpha^K (1-\alpha)^{2-K} + \frac{1-\mu}{27}}. k \in \{0,1,2\}$$

We also know that $A_1'(\alpha) < 0$ since $q > \gamma$, so

$$G'(\alpha) < g'(\alpha) + g'(1-\alpha)$$

We have

$$g'(\alpha) = \alpha \frac{\partial \hat{A}(1,\alpha)}{\partial \alpha} + (1-\alpha)\frac{\partial \hat{A}(0,\alpha)}{\partial \alpha} - \left[\hat{A}(0,\alpha) - \hat{A}(1,\alpha)\right]$$

Moreover, we can explicitly calculate

$$\begin{aligned} \hat{A}(0,\alpha) - \hat{A}(1,\alpha) \quad &= \quad \pi \left[\frac{\mu q (1-\alpha)^2 + \frac{1-\mu}{27}}{\mu \gamma (1-\alpha)^2 + \frac{1-\mu}{27}} - \frac{\mu q \alpha (1-\alpha) + \frac{1-\mu}{27}}{\mu \gamma \alpha (1-\alpha) + \frac{1-\mu}{27}} \right] \\ &= \pi \quad \frac{1-\mu}{27} \frac{\mu (q-\gamma)(1-\alpha)(1-2\alpha)}{\left[\mu\gamma(1-\alpha)^2 + \frac{1-\mu}{27}\right]\left[\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}\right]} \end{aligned}$$

and

$$\frac{\partial \hat{A}(1,\alpha)}{\partial \alpha} = \pi \frac{1-\mu}{27} \frac{\mu (q-\gamma)(1-2\alpha)}{\left[\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}\right]^2}$$

Furthermore, we can show that $\frac{\partial \hat{A}(0,\alpha)}{\partial \alpha} < 0$, so we have

$$\begin{aligned} g'(\alpha) \quad &< \pi \quad \frac{1-\mu}{27} \left\{ \frac{\mu(q-\gamma)\alpha(1-2\alpha)}{\left[\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}\right]^2} - \frac{\mu(q-\gamma)(1-\alpha)(1-2\alpha)}{\left[\mu\gamma(1-\alpha)^2 + \frac{1-\mu}{27}\right]\left[\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}\right]} \right\} \\ &\propto \quad (1-2\alpha) \left\{ \frac{\alpha}{\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}} - \frac{1-\alpha}{\mu\gamma(1-\alpha)^2 + \frac{1-\mu}{27}} \right\} \\ &= \quad \frac{(1-2\alpha)(2\alpha-1)}{\left[\mu\gamma(1-\alpha)^2 + \frac{1-\mu}{27}\right]\left[\mu\gamma\alpha(1-\alpha) + \frac{1-\mu}{27}\right]} \leq 0 \, \text{for all } \alpha \end{aligned}$$

This implies that $g'(1-\alpha) < 0$, and hence $G'(\alpha) < 0$ as required. □

A.2 Uniqueness and full characterisation of symmetric PBE

Combining Lemma A.1 with an argument analogous to Proposition 1 yields the following result, which will hopefully replace Proposition 1 in the body of the paper once we have achieved a general proof that $G(\alpha)$ is strictly decreasing.

Proposition A.1. *For $N \in \{1, 2\}$, there exists a unique symmetric PBE $(\alpha, \beta) \in \left[0, \frac{1}{2}\right)^2$, with the follwing properties:*

1. There is price manipulation by agents with good news ($\alpha > 0$) if and only if $\mu > \bar{\mu}(\gamma, N)$.
2. There is price manipulation by agents with bad news ($\beta > 0$) if and only if $\mu > \bar{\mu}(1-\gamma, N)$.

A.3 A counterexample for $N > 2$

Figure A.1 shows a counterexample indicating that $G(\alpha)$ is not necessarily decreasing in α for large N. The parameters chosen are $N = 20$, $\mu = 0.05$, $\pi = \frac{1}{2}$ and $q = \frac{3}{4}$.

B Further comparative statics

When analysing the degree of price manipulation in equilibrium, α_N, in Section 5 and its price and volatility impact in Section 6, we took the parameters π and q as given and conducted comparative statics with respect to μ and N. To complement this analysis, we now also conduct some comparative statics with respect to π and q. We now take μ and N as given. We only show the results for $\mu = \frac{1}{2}$ and $N = 2$, but they are qualitatively similar for different values.

Figure B.1 shows the impact of changing π and q on α_N. Two main effects emerge:

- For low π and high q, there is non-manipulative behaviour ($\alpha_N = 0$). The reason for this can be seen by examining the condition in Proposition 2 in Section 4: low π and high q imply a low $\gamma(\pi, q)$, which increases the range of value of μ for which non-manipulative behaviour may prevail.

Figure A.1: A counterexample: $G(\alpha)$ may be increasing

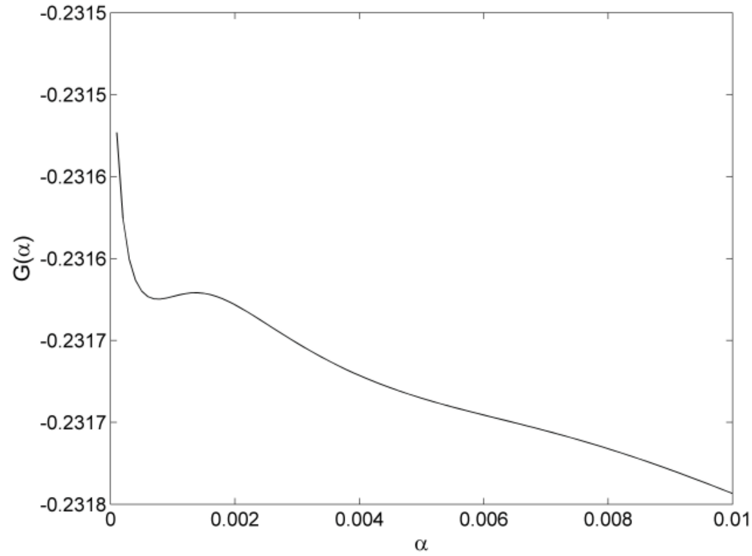

- When there is manipulative behaviour in equilibrium ($\alpha_N > 0$), the degree of manipulation is increasing in π.
- The degree of manipulation is decreasing in q for low values of π, and (slightly) increasing in q for high values of π.

Figure B.1: The degree of price manipulation as a function of π and q^a