6 Manipulation and Repeated Games in Futures Markets

Graciela Chichilnisky

This chapter analyzes the possibility of manipulation in futures markets, concentrating on the effects that manipulation may have on their informational efficiency. We use the concept of manipulation as it arises in the study of noncooperative games with imperfect information.

Forward and futures markets illustrate sharply many of the issues central to the economics of uncertainty and of imperfect information.[1] Clearly, future economic activity is an area in which conditions of uncertainty and of imperfect information arise quite naturally. With respect to uncertainty about future conditions, the existence of a full set of future markets or the equivalent is seen as a precondition for attaining allocative efficiency. One of the major roles of such markets is to allow agents to trade so as to allocate risks optimally among themselves, according to each agent's attitudes toward risk. In this view, futures markets exist because they allow traders with different risk positions toward the future to trade with each other for mutual gain (see, for example, Edwards 1982).

A second, different, role of futures markets is akin to that of a general financial market. In this role, the futures market is seen as an instrument for gathering and distributing information about future market conditions to other parts of the economy (Grossman 1977). This information is of importance for decision making about inventories, outputs and investment, as well as in financial transactions. The performance of futures markets in this sense is measured by their informational efficiency.

We are concerned here with a particular issue concerning informational efficiency, the manipulation of futures markets. This subject has long been of practical importance, but has not until now commanded attention in the literature. The issue of manipulation arises, for instance, in the study of what are institutionally known as squeezes or corners. In both cases, an implicit assumption is that some agents control certain strategic information and that they may use such control to influence the market to their advantage—for instance, through their impact on prices. We assume that agents are not fully informed about the characteristics of all other traders (such as their de-

A grant from the Center for the Study of Futures Markets, Columbia University, is gratefully acknowledged. I thank R. Anderson, R. Kihlstrom, A Kyle, and G.M. Heal for helpful discussions.

mands) and that each agent may use his private information to influence prices to his possible advantage. The context is therefore that of games with imperfect information, and we explore the possibility of manipulation when agents play in a noncooperative fashion—that is, through Nash equilibrium strategies. By manipulation, I refer to the strategic use of information and signals to obtain more advantageous outcomes. I shall illustrate certain examples of manipulation, such as market squeezes: the temporary aberration of the futures prices and spot prices for strategic advantage.

The first section establishes the concepts of games with imperfect information and of manipulation. A brief discussion of the literature is given. A class of games is then used to explore the extent to which the problem of manipulation is likely to arise in these markets. One theorem shows that manipulation arises quite generally, and with it, the informational efficiency in these markets may decrease. Using these games as examples we then set up the problem of manipulation in a repeated game context—that is, games where players are assumed to play repeatedly with each other through time, even ad infinitum (Heal 1976). In this latter case, the incidence of manipulation is greatly reduced. Futures markets become more efficient in their informational role.

We next examine the extent to which a futures market may be viewed as repeated games. This view depends on a number of features, including the degree of anonymity and of restrictions on entry. I argue that these two features are related, in the sense that more anonymity may ease entry. On the other hand, anonymity may prevent the futures market from behaving as a repeated game, thus making it more vulnerable to manipulation.[2]

The problem can be summarized as follows: disclosure that is, less anonymity) may prevent manipulation and therefore improve the informational efficiency of the market. On the other hand, disclosure (less anonymity) may restrict entry, and therefore produce an efficiency loss. There is, in this sense, a tradeoff between informational efficiency and free entry.

It is often argued that the ease of entry in futures markets is a significant improvement from the conditions prevailing in more traditional forward markets. The role of the clearing houses, as discussed in Edwards (9), is in part related to preserving as much anonymity as possible in futures markets. Anonymity and free entry appear to be rather important features of futures markets. It follows that the possibility of manipulation is higher in these markets because they do not easily satisfy the characteristics of repeated games.

The conclusions are that a certain amount of market manipulation can be expected in futures market because of their informational structure, and that manipulation will have some negative effects on the informational efficiency of these markets. Self-policing measures involving some form of

disclosure could decrease to a certain extent the incidence of manipulation. However, such measures carry a cost in terms of barriers to entry and the accompanying efficiency losses. It seems therefore that an overall approach to the problem is to seek an optimal tradeoff between the two types of efficiency losses: informational inefficiencies and restrictions to entry.

The Concept of Manipulation

The concept of manipulation has been studied now for a number of years (Chichilnisky and Heal 1982). It arises most naturally in the context of noncooperative games with imperfect information. This section summarizes the conceptual issues involved and describe briefly existing results.

Economic Games

A game is defined here by specifying four objects:

1. The strategies available to each player—that is, the strategy space S.
2. The space of outcomes, denoted X.
3. The payoff function (or game form) g, a function which assigns an outcome to the strategies played by the individuals.
4. Individual characteristics, such as preferences over outcomes, that determine the strategic behavior of the players.

The term "game with imperfect information" denotes a game in which the players are not fully informed about one or more of the aspects of the game. For instance, players may be aware only of some of the strategies available to them, so they do not know their strategy space S accurately. Another typical incidence of imperfect information is when each agent is not fully aware of the characteristics of the other agents. This type of imperfect information will be most relevant here and we discuss it in some detail.

One important role of future contracts is to provide price signals that can be used by the producers and distributors to allocate real resources. More specifically, futures prices collect and interpret the underlying economic information about conditions of supply and demand and so may influence storage and inventory decisions. In the following, we shall discuss how the issue of manipulation is linked with that of the efficiency of futures markets, and also the different concepts of efficiency that emerge.

The extent to which one can rely on futures prices conveying accurate information about the market's characteristics is relevant for the efficiency of futures markets (Edwards' section 5 and 6). Agents' characteristics, such as preferences, influence demand, and demand affects futures prices. Therefore, when agents' characteristics are unknown, each agent may give strategic signals to the market about these characteristics, in an attempt to shape the pricing structure to his advantage. For instance, a net sale may be considered a signal of an agent's preference. An agent may choose this signal strategically to influence prices according to his preferences. A strategy for each agent i is then a net sale s_i, which is taken as a signal for the agent's preference. This signal will affect market prices at the equilibrium. One can formulate precisely in this context the issue of manipulation. We say that a game with imperfect information is manipulable when for at least some player i, the outcome of the game that obtains when this player gives a signal $\tilde{s}_i$ that misrepresents his characteristic s_i (preference) is better (according to i) than the outcome that obtains if he gives a correct signal about his preference. That is, denote by

$$(s_1, \ldots, \hat{s}_i, \ldots, s_k)$$

a $k - 1$ tuple of strategies of all players but i, where $\wedge$ denotes that the corresponding strategy is deleted; $\succ_i$ denotes "preferred to" by the ith player; X is the outcome space, and the player's strategies are in S.

A game g is manipulable if for some strategy of player i, and some $k - 1$ tuple of strategies of all other players but i, denoted

$$(s_1, \ldots, \hat{s}_i, \ldots, s_k) \, \varepsilon S^{k-1}$$

the outcome

$$g(s_1, \ldots, \tilde{s}_i, \ldots, s_k) \succ_i g(s_1, \ldots, s_i, \ldots, s_k)$$

where s_i is the true characteristic of the ith player, and $\tilde{s}_i \neq s_i$.

This concept of manipulability formalizes the notion that it is individually optimal for some player to misrepresent his characteristics, at least in some cases. As already noted, informational efficiency requires the accurate transmission of information by prices. Therefore, if individual deception leads to different prices than those reflecting the true market conditions, it could translate into a loss of efficiency for the market as a whole. The issue of manipulation is therefore linked to that of market efficiency. This link, however, is not simple, and is discussed in more detail in the following sections. In particular, we shall define a class of games along the lines discussed here and study their manipulability in the last section.

We now give a brief overview of existing results on the manipulation of games that seem useful for the study of manipulation in futures markets. The first results in the theory of manipulation appeared in Gibbard (1973). A certain type of game is called "straightforward" when the individual has no incentive to misrepresent his characteristics in his choice of strategy. The informational structure of straightforward games is such that players do not communicate at all. We now discuss briefly the concept of game solution in relation to the degree of communication among players because it will help formulate the problem with precision.

Imperfect information may take several forms. An agent may be unaware of the other agents' characteristics, but he may be able to observe their strategic moves. This is different from a game where agents are unaware of each other's characteristics and are also unable to observe each other's moves.

The effects of different informational structures is seen more readily through the concept of solution or equilibrium. For example, in a game where each player knows nothing about the other's characteristics and is also unable to observe their strategic moves, the typical concept of a solution is that of dominant strategy equilibrium. In this concept, adopted by Gibbard, each player is playing his dominant strategy—that is, p_i for the ith player, which ensures him of the best possible outcome no matter what other players may be playing. Formally, s_i is a dominant strategy for i if for any $k - 1$ tuple

$$(s_1, \ldots, \hat{s}_i, \ldots, s_k) \, \varepsilon \, s^{k-1}$$

and for all strategies $\bar{s} \neq s_i$ in S, then

$$g(s_1, \ldots, s_i, \ldots, s_k) \succ_i g(s_1, \ldots, \bar{s}, \ldots, s_k)$$

A straightforward game is one in which giving the correct signal about one's characteristic is a dominant strategy for each player, and this gives rise to a dominant strategy equilibrium of the game. Gibbard's theorem can now be simply summarized, even though a few definitions are needed for stating it with precision. For a wide family of games, the only straightforward games are dictatorial. Dictatorial games are those in which the outcome is always identical to the preferred outcome stated by one of the players, called the "dictator." Dictatorial games do not provide an adequate representation of markets.

This result establishes that most nondictatorial games are manipulable, in the sense of not being straightforward. The phenomenon of manipulability appears therefore rather widespread. However, closer examination of Gibbard's result shows that the conditions of his theorem may be quite restrictive. His games generally have no dominant strategies. Therefore, in particular, correct signaling cannot be a dominant strategy equilibrium. Therefore, his games fail to be straightforward may be because they do not have any equilibrium. His result may appear to be mostly a statement about the stringency of the concept of dominant strategy equilibrium and of the assumption that there is absolutely no communication between the players. In addition, Gibbard assumes that no restriction exists on the players' a priori preferences.[3]

Several later articles viewed manipulation results in a wider, and perhaps more realistic perspective (Laffont and Maskin 1980; Chichilnisky and Heal 1981, 1982). We draw from this latter literature in this discussion. The first widening was to recognize that players do observe each other's strategic moves, even though they may ignore each other's true characteristics. Second, it is seldom the case that agents have all possible characteristics, so that it suffices to study market games where the players have characteristics within a subclass of all characteristics.

The first point, about the observability of each other's strategies, leads one immediately to a different concept of solution (or equilibrium) of the game. The concept generally used in games where individuals take into account each other strategic moves is that of a Nash equilibrium. A Nash equilibrium is defined as a vector of strategies

$$(s_1^*, \ldots, s_i^*, \ldots, s_k^*)$$

where strategy p_i^* is such that the ith player maximizes his utility given all other player's strategies. Formally:

$$\text{For all } i,\ g\,(s_1^*, \ldots, s_i^*, \ldots, s_k^*) = \max_{s_i \varepsilon S} \{u_i(g(s_1^*, \ldots, s_i^*, \ldots, s_k^*)\}$$

where u_i is a real valued utility function on outcomes representing the preference of the ith player.

A Nash equilibrium is a familiar concept in the study of market behavior; it is usually referred to as a noncooperative solution. The concept is used, for instance, for the study of markets whose agents have some degree of market power, such as monopolistic competitors. In this context it is called the Cournot solution or Cournot equilibrium. From now on we shall adopt this concept of a solution, which appears to be more realistic in the case of futures markets.

Efficiency and Manipulation

The examples in the last section made an implicit assumption about market behavior: that some agents' supply/demand behavior reflects on market prices. Obviously, in any general equilibrium model, market prices reflect the aggregate supply and demand, which is obtained by adding up individual's supply and demand functions. An individual's behavior therefore affects the equilibrium market prices. However, it is an assumption of the theory of competitive markets that each agent acts as if he has no influence at all on prices—that is, no market power. Our treatment of futures markets as noncooperative games with imperfect information therefore deviates from the standard competitive model in two respects. One is the lack of perfect information. The second aspect is that some of the players are aware that they may have some market power and may be able to influence prices. The concept of futures markets used here is in this sense closer to that used in the chapters by Anderson and Sundareson and by Kyle in this book.

We now turn to the issue of efficiency discussed in the last section. We explained that the manipulation of a futures market may be used in defining the efficiency of this market because the market prices in this case may not convey accurate information about market conditions. There may be another source of inefficiencies, this one related to the overall allocation of resources. If individuals play the market strategically as a noncooperative game and reach a Nash equilibrium solution, this solution need not be an efficient allocation of resources, even when information is perfect. It is well known that Nash equilibrium solutions do not always yield Pareto optimal allocations among the players. In this chapter, however, we concentrate on informational efficiency, which arises more frequently in the study of financial as well as futures markets.

Using the concept of a game introduced in the previous section, we give a formal example of the behavior of futures markets as noncooperative games with imperfect information. Again, we assume that each agent announces a net demand schedule, which is characterized by a number of parameters, and can therefore be viewed as a vector in euclidean space. This vector is a signal of his true net demand function emerging from the optimization of individual preferences. Each component of the vector may, of course, be either positive or negative, depending on whether the agent buys or sells the particular commodity indicated by that component.

We can assume without loss of generality that the initial net amount traded when the market opens is the vector q_0 in R^{nm}, where m denotes the number of delivery dates, and n the number of commodities. Opening futures prices p_0 are therefore described by a positive nm dimensional vector. A signal for agent i is a net demand schedule, a vector denoted Δq_i. It is convenient, but not essential, to assume that q_i also has dimension nm—that is, $q_i \varepsilon R^{nm}$. Because in futures markets no immediate payment is necessary at the moment of contract—that is, there are no budget constraints—in principle a signal can be any vector in R^{nm}, with some coordinates positive and others negative. In the final section we shall also refer to cases where the agents have budget constraints, which limit their signals to a subset of R^{nm}.

In its simplest and most general form, we conceive of the game as a function that assigns to individual net demand schedule signals a market price, which is a positive vector in R^{nm}—that is

$$g(\Delta q_1, \ldots, \Delta q_k) = p \varepsilon R^{(nm)}$$

Equivalently, we may consider the outcome as a change from initial to final prices—that is, we may rewrite the game in the form

$$g(\Delta q_1, \ldots, \Delta q_k) = \Delta p = p - p_0$$

This formalization is useful because we obtain more symmetry. The strategies of the players are vectors in R^{nm}, and the outcomes are vectors in R^{nm} as well. The game form or payoff function is therefore a function

$$g : (R^{nm})^k \rightarrow R^{nm}$$

where k is the number of players. In general, of course, the image of g will be a subset of R^{nm}.

We now focus on one class of games within this context, which is used later to explore the incidence of manipulation. The goal of each player is to attain a price change as close as possible to Δp^*_i, the ideal price change for this agent given his true (current or expected) market position. For example, assume that there are two periods and that each component of the price vector denotes the price for the same good in each of the two periods. Assume that in the first period the agent goes long for delivery of good a on the second period, and furthermore, that it is his private information that he does not wish to hold good a on or after the second period. Then if this agent can induce by strategic signaling a change in market prices that keeps future prices for good a at the first period, denoted $p_1(a)$, as low as possible, and second-period spot prices for a, denoted $p^2(a)$ as high as possible, he may be able to squeeze the market for delivery at the second date, provided he purchases enough in the first period for delivery at the second. His goal is then to obtain that change in price Δp^*_i that maximizes the ratio

$$\frac{p_1(a)}{p^2(a)}$$

For instance, if the agent's net position is long, his ideal price ratio in this market would be zero. In other cases—such as those with two or more delivery dates—one may consider the ideal price as representing instead futures prices at different delivery dates. If an agent holds a portfolio with different delivery dates, the ideal prices for this agent will in general be a vector whose components are positive. How an agent may influence prices to approximate his ideal price is described in the last section.

We now assume that prices are affected by the behavior of a subset of players who have market power P Prices will change in the same general direction of the excess demand vector of the players in P. More precisely, if Δq_i is the demand signal of the ith agent with market power, then Δp will be in the convex set of directions determining the signals Δq_i for all i in P. In particular, when all quantity signals are identical to each other—for example, to Δq (that is, everyone signals the same net demand)—then the change in prices will also be in the same direction. That is

$$\Delta p = \lambda \Delta q$$

for some positive number λ.

We can now describe the strategic behavior of the players. The optimal Nash strategy of the ith player, given that all other players in P are playing strategies Δq_j, for $j = 1, \ldots, k, j \neq i$, is that strategy $\overline{\Delta q_i}$ that yields a price change as close as possible to an ideal outcome Δp^*_i. Formally, the Nash strategy of the ith player is $\overline{\Delta q_i}$ if

$$g(\overline{\Delta q_1}, \ldots, \overline{\Delta q_i}, \ldots \overline{\Delta q_p}) - \Delta p^*_i = \min_{\Delta q \varepsilon R^{nm}} (g(\overline{\Delta q_1}, \ldots, \Delta q, \ldots, \overline{\Delta q_p}) - \Delta p^*_i)$$

where $P = \{1, \ldots, p\}$, and the minimum is taken with respect to the standard euclidean distance in R^{nm}.

A Nash equilibrium set of strategies $(\Delta q_1{}^*, \ldots, \Delta q^*_p)$ is one in which for each player i, the strategy Δq^*_i is optimal, given that player j is playing strategy Δq^*_j, for $j = 1, \ldots, p, j \neq 1$. There is manipulation only if the Nash equilibrium strategy of agent i, Δq^*_i, which he chooses strategically to attain the best outcome in the game g, is a misrepresentation of i's net demand Δq_i, obtained under competitive assumptions from utility maximization at the competitive equilibrium market prices. If some player gains by misrepresenting his market position—for example, if it influences prices to move in a different way than they would do if he was to represent his position accurately—the informational efficiency of futures prices in predicting subsequent spot prices will be diminished. In the last section, we show that in this type of market game manipulation will take place generally—that is, each player will in general obtain a more favorable price move by misrepresenting his position (theorem 2). We can further refine the result by showing that when the outcomes of the game are directions of price changes, then under these conditions there exists always one player who can attain whatever direction of price change he desires by manipulation (theorem 1).

The results of theorems 1 and 2 show that although we have dropped the restrictions of Gibbard's theorem and consider more plausible games where agents do take into consideration each other's strategic moves, the problem of manipulation is still present. The next section will study alternatives to the examples of games discussed here and explore the role of disclosure in the context of repeated games.

Disclosure and Repeated Games

We have discussed examples of the incidence of manipulation in futures markets viewed as noncooperative games with imperfect information. In this section we analyze the strategic behavior that arises when players play the game repeatedly, even ad infinitum. The incidence of manipulation is likely to decrease when the game is played by the same agents repeatedly because each player's strategic behavior is observed by the other players. Once manipulation is exposed and the player is identified, future signals from this player may be interpreted differently. In particular, it is possible that by playing the game repeatedly, the manipulative player will reveal his true market position through his strategic behavior. If this is the case, the longer-run informational efficiency of the futures market as a repeated game may be recovered, despite the possible incidence of manipulation in each one-shot game.

Several factors may stand on the way of the full disclosure of an individual's position through his strategic behavior. However, at least in certain examples one can give sufficient conditions to guarantee that an optimal strategy in a repeated game is to reveal one's true market position. Such examples will be seen to require some form of disclosure or loss of anonymity. However, disclosure or loss of anonymity may be associated with barriers to entry, which decrease the market's allocative efficiency. Therefore, the gains from informational efficiency of repeated games may be accompanied by efficiency losses from barriers to entry. A mixed policy to optimize this tradeoff may be called for.

Ideally, we would consider the games discussed in the last section as played repeatedly. However, this view would lead to games on infinite dimensional strategy spaces. Therefore, we study now a simpler example of one-shot games that will then be repeated indefinitely. This is analogous but different from a one-shot game studied by Akerloff (1970). An extension of our game to a supergame—that is, the game obtained by repeating this one-shot game ad infinitum—is obtained. A precedent is Heal (1976) who

extended Akerloff's game to a supergame and produced the first results in the area of incentives in repeated games. Because we repeat the game indefinitely, it is simpler to assume that only two strategies are available to each player, one representing truthful demand A and the second misrepresenting it for calculated strategic advantage B.

Our game differs from those of Akerloff and Heal in several ways. Heal requires that a "good quality" good obtained through trade have an intrinsic value for the player, a value that is the same for both players and is independent from what the other player's strategy is. Thus, his game has only four parameters: the value of obtaining a "good quality" good, the value of departing from a "good quality" good, and the same two values for "bad quality" goods. Here, we need instead eight parameters because there is no intrinsic value here to a truthful strategy. This value is determined by the market response, which depends of course on the other player's strategy.

We assume that there are two players, 1 and 2, and define eight parameters as follows. When players 1 and 2 play both strategy A, the outcome for 1 is α_{11} and for 2 is β_{11}; when 1 and 2 both play B the outcomes are α_{22} and β_{22}; when 1 plays A and 2 plays B, the outcomes are α_{12} and β_{12}, respectively, and finally when 1 plays B and 2 plays A, they are α and β_{21} respectively. We now define this game formally. The game form g is a function

$$g: \{A,B\}^2 \rightarrow \{(\alpha_{11}, \beta_{11})\ (\alpha_{12}, \beta_{12})\ (\alpha_{21}, \beta_{21})\ (\alpha_{22}, \beta_{22})\} \subset R^2$$

where $\{A, B\}$ is the set of strategies of each player, consisting of two strategies, A and B. The set of outcomes is contained in R^2. Each different set of strategies is assigned one outcome—for example, by construction

$$g(A, B) = (\alpha_{12}, \beta_{12})$$

We can represent the same game also in the more familiar matrix form:

$$\begin{array}{lll} & & \text{Player 2} \\ & & \begin{array}{cc}\text{strategy } A & \text{strategy } B\end{array} \\ \text{Player 1} & \begin{array}{r}\text{strategy } A \\ \\ \text{strategy } B\end{array} & \left(\begin{array}{cc} \alpha_{11}\ \beta_{11} & \alpha_{12}\ \beta_{12} \\ & \\ \alpha_{21}\ \beta_{21} & \alpha_{22}\ \beta_{22} \end{array}\right) = g \end{array}$$

By analogy with the previous game we may assume that the truthful outcome $(\alpha_{11}, \beta_{11})$ is Pareto-efficient—that is, that any other outcome which has a higher value for one of the players will necessarily have a lower value for the other. We shall also discuss cases where $(\alpha_{11}, \beta_{11})$ is not Pareto-efficient.

The next step is to find the non-cooperative solutions of this one-shot game. We are concerned with the cases where the game can be manipulated and wish to investigate how the repetition of this game may improve matters. The game can be manipulated when the outcome of playing strategy B (deceit) is an improvement over that of playing strategy A (stating one's truthful position). We may assume without loss of generality that the deceitful player is 1. Then if

$$\alpha_{21} > \alpha_{11} \text{ and } \alpha_{22} > \alpha_{12}$$

player 1 will always play strategy B (deceit). Player 2 will therefore always choose between β_{21} and β_{22} only. Under the assumption that $\beta_{22} > \beta_{21}$ (that is, it is preferable to respond to deceit with deceit), the only Nash equilibrium of this game is the pair of strategies (B, B) with payoff $(\alpha_{22}, \beta_{22})$. This will happen even though the truthful (A, A) strategy vector may yield a Pareto superior outcome—that is, even if $(\alpha_{11}, \beta_{11}) > (\alpha_{22}, \beta_{22})$, where $>$ is the standard vector order in euclidean space R^2. The Nash equilibrium—that is, noncooperative behavior of the agents—may lead to Pareto inferior outcomes if each player has an incentive to deceive the other in the one-shot game. Such a market would not be informationally efficient because at the equilibrium, we expect deceitful behavior of the agents.

Now assume the game is played repeatedly. Consider the following infinite strategy Γ for player 1: to real his correct position A in the first period, and in the tth period to play B if and only if player 2 has been deceitful in some previous period $s < t$. Define Ω to be the symmetric policy for player 2. We can now compute the discounted future payoff of this strategy for both player. The playoff of Γ to player 1 is 2 plays Ω is

$$\alpha_{11} \sum_{t=1}^{\infty} \Delta^t \tag{6.1}$$

where $0 < \Delta < 1$ is the discount factor. Similarly, the payoff to 2 of Ω if 1 plays T is

$$\beta_{11} \sum_{t=1}^{\infty} \Delta^t \tag{6.2}$$

Now, if player 2 plays strategy Ω, can player 1 benefit by departing from strategy A? Assume that from $t = 1$ to $t = T$, 1 plays A, and for $t \geq T + 1$, 1 plays B. Then if 2 follows strategy Ω, 2 will play B from $t = T + 2$ onward, and therefore, the best 1 can do is to play B from there onward also. Therefore, the highest payoff to 1 of departing from strategy T at time $T + 1$ is

$$\alpha_{11} \sum_{t=1}^{T} \Delta^t + \alpha_{21} \Delta^{T+1} + \alpha_{22} \sum_{t=T+2}^{\infty} \Delta^t \tag{6.3}$$

We can now compare the payoff to strategy Γ for player 1, to the payoff to this deviation from Γ. This is, we compare equations 6.1 and 6.3. It is easy to check that the payoff of equation 6.1 exceeds that of 6.3 if and only if

$$\Delta < \frac{\alpha_{11} - \alpha_{21}}{\alpha_{22} - \alpha_{21}}$$

Therefore, for sufficiently small discount rates Δ, the pair of strategies Γ, Ω is always a Nash equilibrium of the repeated game. This shows that for small discount rates—that is, when players value their future trades sufficiently—one possible solution to this game is that each player acts according to his true position without attempting to manipulate the outcome.

We therefore may attain informational efficiency if the discount rate is sufficiently small. Moreover, it can also be shown that for any discount rate, a Pareto-efficient solution of the game is an equilibrium of the supergame. Therefore, the truthful strategy (A, A) leads to a Pareto-efficient allocation. It follows that in such cases one may add Pareto efficiency of resource allocations to the informational efficiency of the solutions.

An interesting problem arises in those cases where the truthful strategy set (A, A) is informationally efficient but not Pareto-efficient. As discussed previously, this may arise in the Nash equilibrium of noncooperative market games even with perfect information, such as games of monopolistic competition. In such cases, it cannot be guaranteed that gaining more information about the market conditions (for example, through repeated games) will improve the outcome. More information may lead in some cases to all agents being worse off.

Finally, it should be pointed out that, in general, any Pareto-efficient allocation of a one-shot game will be a Nash equilibrium of the supergame. This result implies in particular that if the truthful strategy (A, A) is Pareto-efficient, it will always be a Nash equilibrium solution to the supergame.

Efficiency Gains and Losses from Disclosure

The previous section studied a one-shot game where the agents have an incentive to manipulate their signals to their advantage. It also showed sufficient conditions for this incentive to disappear when the game is repeated indefinitely. The intuitive reason is that when the game is repeated, the players build up reputations and may therefore internalize at least some of the losses that they may inflict on others in previous periods. The incentive to manipulate is therefore decreased. We exhibited two sufficient conditions for attaining a manipulation-free outcome. One is that the truthful strategies define a Pareto-efficient equilibrium. The other is that the agents have relatively low discount rates for the future. Clearly, the extent to which future trades matter will be reflected by more concern for one's current commercial reputation. These results lead us naturally to question the conditions under which a futures market can be considered a repeated game. This will be the first subject of this section. The second will be to explore the implications of this on efficiency.

One factor that emerged clearly in the discussion in the last section is that some form of strategic retaliation is necessary to prevent repeated manipulation. Clearly, such a strategy would require that the manipulating agents be identified. For example, if manipulation is followed by exit from the market, and perhaps reentry under a different brand name, manipulation may go unchecked and be repeated indefinitely.

To formalize this concept, one reformulates the repeated game defined previously to take exit into account. Depending on the returns outside of the game, one may be able to formulate precisely the optimal exit policy of a manipulative agent. For example, assume that the returns outside the game are x dollars per period. Consider the following strategy ξ for player 1. Player 1 plays the game straight for T periods, then it manipulates it on period T, and leaves the market on period $T+1$. Then the payoff to 1 of strategy ξ, under the assumption that player 2 will not manipulate unless 1 does (that is, strategy Ω) is

$$\alpha_{11} \sum_{1}^{T-1} \Delta^t + \alpha_{21}\Delta^T + x \sum_{t=T+1}^{\infty} \Delta^t \qquad (6.4)$$

We may now compare the payoff of strategy Γ for player 1, with the payoff of two other strategies: manipulating and staying in the market, and manipulating and exiting. Clearly it will be preferable for 1 to follow strategy ξ rather than the straight strategy T if and only if

$$\Delta > \frac{\alpha_{11} - \alpha_{21}}{x - \alpha_{21}} \qquad (6.5)$$

Notice that the choice of cheat and exit strategy becomes more attractive in two cases:

1. The higher is the rate of discount of the future payoffs.
2. The higher is the payoff x outside the market, and this is independent from the stopping time.

Obviously, x must be larger than α_{22} because otherwise player 1 would never contemplate leaving the market. Also x must be smaller than α_{11} for this player to want to play at all. Therefore strategy ξ will generally be preferred to manipulating forever and will also be preferable to playing straight with high rates of discount of future payoffs.

As a result, there is the concern that at any point of the game a player may manipulate and then exit. Unless a player can be fully identified in terms of his history of trades, one cannot expect the players to reveal truthfully their market positions so that newcomers, who have no market history, would normally be suspect. With full disclosure, a wedge is driven between oldtime players and newcomers, which effectively restricts entry.

When more sophisticated strategies are considered, it can be expected that a natural concept of entry fee may arise—that is, the cost associated with developing a good market reputation for the newcomers. The formalization of such a concept would seem useful to compute the efficiency losses associated with restricted entry arising from disclosure. Or, equivalently, it may measure the efficiency gains from anonymity, in the form of free entry. Therefore, a measure of the informational efficiency gained by disclosure (in which case we may have a repeated and manipulation-free game) and the efficiency losses due to restricted entry caused by full disclosure would seem required.

Results for One-Shot Games and Applications

In this section we prove results on the manipulation of games with imperfect information. At the end of the section we shall also discuss their possible applications for the analysis of futures markets.

Let us assume that there are $k \geq 2$ players with market power. We shall examine their Nash strategies and the corresponding outcomes of a non-cooperative game. The game is defined by a game form g, a strategy set S in R^n for each player, and an outcome set X in R^n. Each strategy is a vector representing a net demand schedule for n commodities. The game form is a continuous function g: $(R^n)^k \to X$, which assigns to each k-tuple of strategies an outcome that is a direction of price change D_p in R^n, or else no change at all—that is, the vector $(0, \ldots, 0)$.

We shall assume that each player knows g, X, and S, and that they observe each other's moves; player i's preferred direction of price change is denoted Dp^*_1. To provide a simple proof of our next result, we shall look at the special case $k = 2$, $n = 2$. The results in this section hold true for higher dimensional cases, but the proofs require more complex tools of algebraic topology.

THEOREM 1. *Consider a game g: $(R^2)^2 \to R^2$ defined as above. Assume that prices move in the direction of a convex combination of the changes in net demands of the agents. There then exists a player who is always able to secure, as a Nash equilibrium outcome, his preferred direction of price change, for any (nonzero) strategies the other is playing. To attain this outcome, this player will generally misrepresent his true (net) demand, but his strategic net demand vector need never have a higher absolute value than that of the other player.*

Proof. Because g maps strategies into directions of price changes, and such directions are in a one-to-one correspondence with points in the unit circle S^1 of R^s, we may consider g: $(R^2)^2 \rightarrow S^1$. We shall now restrict ourselves to nonzero strategies, so we may look at g: $(R^2 - \{0\})^2 \rightarrow S^1$

We study next the restriction of the map g to the set $(S^1)^2 \subset (R^2 - \{0\})^2$—that is

$$g : (S^1)^2 \rightarrow S^1$$

We may define the degree of g restricted to the diagonal set $D = \{(x_1, x_2) \in (S^1)^2 : x_1 = x_2\}$, because D is homeomorphic to the circles S^1 in R^2 (Chichilnisky 1981).

The condition that price changes in the direction of positive linear coordination demands implies that when all individual demands are co-linear, they move in the same direction. Thus g restricted to D is the identity map, so that $\deg g/D = 1$. Similarly we study the degree of g on each of the following subsets:

$$T_1 = \{(x, y): x = x_0, y \in S^1\}$$

and

$$T_2 = \{(x, y): y = y_0, x \in S^1\}$$

The degree of g restricted to T_1 and to T_2 is either zero or one, due to the convexity condition (Chichilnisky 1982b)—that is

$$0 \leq \deg g/T_1 \leq 1 \tag{6.5}$$

and

$$0 \leq \deg g/T_2 \leq 1 \tag{6.6}$$

Now, because D is homotopic to $T_1 U T_2$ within $(S^1)^2$, it follows that

$$\deg g/D = \deg g/T_1\ UT_2 = \deg g/T_1 + \deg g/T_2 \tag{6.7}$$

Since $\deg g/D = 1$, this implies by equation 6.5 and 6.6 that of the two degrees on the right side of equation 6.7, one must be zero, and the other one.

Assume without loss of generality that $\deg g/T_1 = 1$ and $\deg g/T_2 = 0$. This implies that when player 2 plays y_0 in S^1 for any outcome D_p^* in S^1 there exists an $x(y_0)$ in S^1 such that $g(x, y_0) = D_p^*$. By continuity of g, this is also true for any other y in S^1, $y \neq y_0$. This proves that the first player can always find a Nash strategy that ensures him of his preferred outcome, no matter what strategy player 2 plays. The argument is now easily extended to strategies in R^2.

For any y in R^2, consider the circle $S^1(y)$ centered in the origin, passing through y. Because g is defined on R^2, in particular g is defined on $S^1(y)$, and thus the preceding argument applies to this circle—that is, to the map

$$g/_{S^1(y)xS^1(y)} : S^1(y)xS^1(y) \rightarrow S^1$$

This proves that for the game g: $(R^2)^2 \rightarrow S^1$, the Nash response function of player 1 has always D_p^* as its outcome—that is, $g(x(y), y) = D_p^*$, $\forall D_p^*$ in R^2. Because the Nash equilibrium is in the intersection of both response functions of the players, the proof of the theorem is complete.

The following is an example of a game g as in theorem 1.

Example 1: A special case of Walras games of misrepresentation of preferences. This example draws on the literature on market manipulation (see Hurwicz 1972, 1979). We consider first a pure exchange market with two persons and two goods. Given an initial endowment W_i and a price p, each agent determines a utility-maximizing bundle $z^*(p)$. As the price p varies, the geometric locus of $z^*(p)$ in the commodity space constitutes the offer curve of this agent.

One may consider the game where each individual chooses strategically an offer curve to maximize his strictly convex preference subject to a budget constraint depending on W_i. Given a pair of such strategies, denoted h_1 and h_2, the outcome of the Walras game is defined by the determination of the market clearing prices for h_1 and h_2 and the subsequent selection of the corresponding Walras equilibrium allocation. In the case of multiple solutions, one is chosen.

Hurwicz has shown that the set of Nash allocations of the preceding game, which corresponds to equilibria in Nash strategies, coincides with the interior of the lens L^* constituted by the true offer curves. In general, therefore, manipulation of the market will take place, in the sense that the Nash solutions of the Walras game when agents play strategically is different from the set of Walrasian equilibrium market allocations.

We consider now a special case of the preceding game.

Given inital endowments W_i and preferences U_i, $i = 1, \ldots, k$ (k agents), let $\Sigma(W_i, U_i)$ be the set of Walrasian equilibria of the pure exchange economy described by $(W_i, U_i)i = 1, \ldots, k$. We assume that initial endowments (W_i) are given, and that the preferences $U_i, i = 1, \ldots, k$, may vary over a family of preferences parameterized by vectors in R^n (n goods).

This is a restricted domain assumption. For example, for $n = 3$, $k = 2$ let $U_1 = \min(a_1x,\ b_1y,\ c_1z)$ and $U_2 = \min(a_2x,\ b_2y,\ c_2z)$, so that U_1 is fully described by the three = dimensional vector (a_1, b_1, c_1) and U_2 by (a_2, b_2, c_2). Alternatively consider a faimily of Cobb-Douglas utilities $\{U_i = (x^{\alpha}, y^{\beta}, z^{\gamma})\}$ each utility U_i indexed by a vector in euclidean space, namely $(\alpha, \beta, \gamma) \in R^3$.

We shall assume that there exists a continuous map ϕ from utilities to equilibria

$$\{U_i\} \xrightarrow{\phi} \sum(U_i, W_i)$$

where ϕ assigns a Walrasian equilibria to each utility U_i in a continuous fashion. Because the utilities U_i's are assumed to be characterized by vectors in R^n, then the map ϕ can be written as

$$\phi : (R^n)^k \rightarrow \sum(.;W_i),$$

where $\phi\ (r_1^1, \ldots, r_n^1, \ldots, r_l^k)$ is a Walrasian equilibria of the pure exchange economy $(r_1^1, \ldots, r_n^1, \ldots, r_n^k; \{W_i\})$ with initial endowments $\{W_i\}$ and preferences $\{U_i\}$ represented by the nk vector of parameters $(r_1^1, \ldots, r_n^k)$.

If we now consider the equilibrium price p^* supporting the Walrasian equilibria allocation $\phi(r_1^1, \ldots, r_n^k; \{W_i\})$, then we obtain from ϕ a continuous map

$$g : (R^n)^k \rightarrow R^n$$

assigning to $(r_1^1, \ldots, r_n^k)$ the equilibria price of $\phi(r_1^1, \ldots, r_n^k, \{W_i\})$. This map satisfies the conditions statement of theorem 1, defining a game form $g: (R^n)^k \rightarrow R^n$. The strategy of the ith player is therefore an n-dimensional vector $(r_1^i, \ldots, r_n^i)$ in R^n representing his preference, or corresponding demand schedule. Each individual vector will represent a variation from an initial vector of parameters $(r_1^0, \ldots, r_n^0)$. Individual strategies are variations over a given preference or initial demand schedule. We now make the following additional assumption:

Regularity assumption: The equilibrium price p^* varies continuously as a function of individual demands in the direction of the convex combination of changes in individual demands. This condition can be described intuitively by saying that the equilibrium price moves in a certain direction whenever individual utilities change so as to assign higher utilities to commodity bundles in that direction. This condition can be weakened significantly, for instance to request that the map of from $(R^n)^k$ into R^n has

degree 1 over certain subsets. In view of our two assumptions, the game form as defined by the Walras game satisfies all the conditions of theorem 1, and therefore the results of theorem 1 apply to this example.

We may also refer to cases where, because of constraints, the players may not play all possible net demand vectors as strategies. In the two-dimensional case, we may consider therefore that the strategies open to each player are restricted to a box in R^2, denoted Z^2. The manipulation of such games was studied in Chichilnisky and Heal (1982), and we quote here those results.

THEOREM 2. *Let g be a regular*[4] *game with strategies in Z^2 for each player, and outcomes in a convex set of R^2 (that is, the price space). Then g is nonmanipulable in Nash equilibrium only if g is locally simple—that is, locally a constant or locally dictatorial. Furthermore, locally constant or dictatorial games are nowhere dense in the family of continuous game forms g: $Z^2 \rightarrow R^2$. Therefore, generically, games g: $Z^2 \rightarrow R^2$ are manipulable.*
For a proof see theorem 2 and proposition 5 of Chichilnisky and Heal (1982).

We now give a corollary of theorem 1 that will be used in the following application to futures markets:

COROLLARY 1. *Let g: $(R^2)^2 \rightarrow R^2$ be a game as in theorem 1. Then there always exists a player that can ensure that the price of one of the goods will move in the opposite direction of his net demand vector for this good, at least in some ranges of his demand.*
Proof. This follows from the proof of theorem 1. The fact that $\deg/T_2 = 0$ implies, together with the convex hull condition, that the set of values of g on the set T_2, that is

$$\{g(x_0, y) : y \in S^1\ (x_0)\}$$

does not cover T_2. That is, as the net demand vector y of the second agent describes clockwise the circle $S^1\ (x_0)$, the outcome must move counter-clockwise at least for some values of y. Therefore, as net demand of player 2 increases for one good, the price change moves in the opposite direction. This completes the proof.

We now explore an application of theorem 1 and its corollary 1 to a particular example of manipulation of futures markets, related to what is sometimes called a market squeeze (for a discussion and definitions, see for example Kyle, chapter 5, this volume).

Example 2: Market squeezes and the competitive fringe. For this example we must specify in more detail the institutional framework of the problem. We shall assume that there are two types of agents, those with market power and those without it. The latter are called the "competitive fringe." They are distinguished in operational terms by the fact that when operations are contractual but not physical (that is, no physical goods are exchanged, only contracts), the prices are determined by the market behavior of the players with market power. However, if as delivery date arrives physical deliveries take place, then the price changes are influenced by the physical volumes of demands and supplies of all players, including those without market power, until physical markets clear. We shall assume, as usual, that price changes move in the same general direction as aggregate excess demands.

In our example, there is one good a and two periods. We shall consider two prices: p_1 (a) denotes the futures price of a at period 1, and $p^2(a)$ represents the spot price of a at period 2. Obviously, with perfect information and no manipulation these two prices should be equal but for storage costs. As we will see, however, where there exist agents with market power in period 1, it will be possible (under certain conditions) to drive a wedge between these two prices to the advantage of the manipulative agent.

We shall consider a case where the first period is very close to the delivery date (or second period), so that the clearing house is not able to close the wedge through its periodic monitoring operations. Assume that there are two players with market power denoted 1 and 2 and a competitive fringe of undetermined size. Assume that the direction of price change is as before in the convex hull of player 1 and player 2's net futures contract demand for good a (to be delivered at date 2). Then corollary 2 establishes that at least for one agent, say player 1, it will be possible in some cases to increase its demand for a (to be delivered at date 2) and go sufficiently long without at the same time increasing, or even while decreasing, the futures price at which he contracts in period 1.

An intuitive explanation of this case could be as follows. If in previous periods agent 1 had traded with an agent with market power denoted 2, and 2 went sufficiently short, then in period 1, the second agent could prevent the futures price of a from rising—for example, while agent 1 goes long by buying only from the competitive fringe. Because we assumed that until physical trade takes place, the competitive fringe does not affect market prices, futures prices for a remains low, even as player 1 goes sufficiently long that his demand exceeds physical supplies in period 2. The manipulation is now completed. As period 2 arrives, if player 1 purchased more than the total physical quantities available, then obviously the spot price of good a will rise in period 2. This increase will give a net gain to player 1 if he accepts monetary compensation for the lack of delivery. This gain, of course, will only be meaningful if player 1 did not actually buy futures in good a because he needed good a in period 2; contrary to what he expressed about his demand for a in period 1, he does not actually need to consume a in period 2, so he can materialize the gain of the price wedge he produced through manipulation.

This result has two key elements. First, as provided by theorem 1 and corollary 1, in period 1 one agent may increase his futures demand for good a and go sufficiently long without increasing his futures price in contracts at period 1. The competitive fringe that went short in the aggregate in period 1 for delivery in period 2 will affect spot prices in the second period, because then delivery is enforced, so that the other physical scarcity of quantities traded affects spot prices.

A closer look at agent 1's strategy suggests that this agent may do well to buy first from those agents with the most market power. If they go short, they may help prevent increases in futures prices, thus allowing player 1 to continue to buy from the competitive fringe at lower prices and increase his long position significantly just up to the date where delivery must take place.

Notes

1. Forward markets do not mark to market as futures markets do.
2. For an institutional example of this tradeoff between ease of entry and manipulation, see "Antitrust Study of U.S. Bond Trading." *The New York Times*, April 4, 1983.
3. That is, a priori each player may have any possible preference among different commodities. The game is supposed to be straightforward with respect to any arbitrarily given set of players, each of which may have any possible preference.
4. A regular game is one whose game form $g : R^{nk} \rightarrow R^n$ satisfies generic transversality conditions; see Chichilnisky and Heal (1981).

References

Akerlof, G. (1970). "The Market for Lemons," *Quarterly Journal of Economics*.

Chichilnisky, G. (1982a). "Incentive Compatible Games: A Characterization of Strategies and Domains," Working Paper 139, Department of Economics, Columbia University.

———. (1982b). "Social Aggregation Rules and Continuity", *Quarterly Journal of Economics*.

———. (1981). "The Topological Equivalence of the Pareto Condition and the Existence of a Dictator," *Journal of Mathematical Economics*, 1981.

——— and Heal, G. (1982). "A Necessary and Sufficient Condition for Straightforwardness." Working Paper 138, Department of Economics, Columbia University.

———. (1981). "Nash-Implementable Games." Working Paper, University of Essex.