

On the Gains and Losses of Speculation in Equilibrium Markets

Tuomas Sandholm

Department of Computer Science
Washington University
One Brookings Drive
St. Louis, MO 63130-4899
sandholm@cs.wustl.edu

Fredrik Ygge

EnerSearch and
Department of Computer Science (IDE)
University of Karlskrona/Ronneby
372 25 Ronneby
Fredrik.Ygge@enersearch.se

Abstract

In computational markets utilizing algorithms that establish a market equilibrium (general equilibrium), *competitive behavior* is usually assumed: each agent makes its demand (supply) decisions so as to maximize its utility (profit) assuming that it has no impact on market prices. However, there is a potential gain from *strategic behavior* (via speculating about others) because an agent does affect the market prices, which affect the supply/demand decisions of others, which again affect the market prices that the agent faces.

This paper presents a method for computing the maximal advantage of speculative behavior in equilibrium markets. Our analysis is valid for a wide variety of known market protocols. We also construct demand revelation strategies that guarantee that an agent can drive the market to an equilibrium where the agent's maximal advantage from speculation materializes.

Our study of a particular market shows that as the number of agents increases, gains from speculation decrease—often turning negligible already at moderate numbers of agents. The study also shows that under uncertainty regarding others, competitive acting is often close to optimal, while speculation can make the agent significantly worse off—even if the agent's beliefs are just slightly biased. Finally, protocol dependent game theoretic issues related to multiple agents countering speculating are discussed.

1 Introduction

General equilibrium theory, a microeconomic market framework, has recently been successfully adapted for and used in computational multiagent systems in many application domains, see e.g. [6; 16; 17; 9; 18; 2]. It provides a distributed method for efficiently allocating goods and resources among agents. Such a market can have two types of agents, *producers* and *consumers*. It has a finite number of *commodities*. The amount of each commodity is unrestricted, and each commodity is usually assumed arbitrarily divisible. Different elements

within a commodity are not distinguishable, while elements from different commodities are. Each consumer has a *utility function* which encodes its preferences over different *consumption bundles*, i.e. over vectors. Each element of the vector describes how much of a given commodity the agent consumes. Each consumer also has an initial *endowment* of the different commodities. The producers can use some commodities to produce others. The *production vector* of a producer describes how much of each commodity the agent produces. Net usage of a commodity is denoted by a negative number. A producer's capability of turning inputs into outputs is characterized by its *production possibilities set*, which is the set of feasible production vectors. The producer's profits are divided among the consumers according to predetermined proportions which need not be equal (one can think of the consumers owning stocks of the producers).

The market is said to be in *general equilibrium* in terms of the prices on commodities, consumers' consumption decisions, and producers' production decisions if

- I markets clear: for each commodity, production plus endowments equals consumption, and
- II each consumer consumes a bundle of commodities such that the agent could not afford another bundle of higher utility given its initial endowments, the current prices, and the profits it receives from producers, and
- III each producer uses the feasible production vector that maximizes its profits given the prices.

If the production possibilities sets are convex, and the consumers' preferences are continuous, non-decreasing and locally insatiable, then at least one such equilibrium exists [8; 5; 15].¹ A sufficient condition for *uniqueness* of such an equilibrium is that the demand for each good is nondecreasing in the prices of the other goods.

¹If agents act competitively (as opposed to strategically via speculation), each general equilibrium is Pareto efficient: no agent can be made better off—with any methodology for making decisions—without making some other agent worse off. Such competitive general equilibria are also stable in the sense of the *core* solution concept of coalition formation games: no subgroup of agents is motivated to pull out of the general equilibrium and form their own market.

The analysis in this paper is based on the assumption that a protocol is used that establishes a market price for each good such that supply meets demand, and that reallocation is performed *after* these prices have been established. There are many alternative algorithms that can be used to find such a general equilibrium. Clearly, if no such equilibrium exists, no algorithm can find it. In this paper we analyze the gains and losses of strategic behavior via speculation. We do this by analyzing the equilibrium. If an equilibrium does not exist, the agents will not achieve a resource reallocation, and in that case, the gains and losses of speculation are not well defined. The equilibrium-based analysis makes our results protocol independent—as long as the agents exchange goods only after an equilibrium has been reached. This allows our results to hold for most market algorithms that have been used to find an equilibrium. Some of these algorithms are now discussed.

To reach a general equilibrium, the *price tâtonnement process* is usually used. This is an iterative mechanism, and the trades, production, and consumption are assumed to occur only after the process has terminated. At each iteration, the *auctioneer* sets a vector of prices. Then all agents have to declare a vector of how much they are willing to buy and sell of each commodity at the current prices. Based on this information, the auctioneer updates the price vector for the next iteration. Under certain technical conditions, this process is guaranteed to converge to a general equilibrium [8]. Within computational multiagent systems, Wellman has developed a general equilibrium based software system called WALRAS. As example domains, he has used WALRAS in flow routing in a network [16], and in configuration design [17]. Wellman discusses how the latter application fails to meet the assumptions for existence of a general equilibrium. For example, because some design parameters are discrete, production possibilities sets cannot be convex. Mullen and Wellman have applied WALRAS to a distributed information network example [9]. The iterative market process of WALRAS differs from tâtonnement. Specifically, WALRAS uses asynchronous declarations by the agents, and the agents bid demand functions (of price) as opposed to just quantities. This process converges to a general equilibrium [2]. As in tâtonnement, trades in WALRAS only occur after the market process has terminated.

In addition to the price-based market mechanisms, resource-based mechanisms exist for reaching the general equilibrium. In such a resource-based protocol (*cf. quantity tâtonnement* [8]), the auctioneer sets the allocation of commodities to agents at each iteration, and agents report how much more or less they are willing to pay for each commodity. The auctioneer then takes these declarations into account in changing the allocation in the next iteration: agents that increase their willingness to pay get more, and others get less than on the previous

iteration. The algorithm terminates when an equilibrium is reached. Kurose and Simha have developed a market mechanism—applied to file allocation—where at each iteration, the agents compute their marginal worths for resources [6]. Based on these worths, resources are reallocated at every iteration. This differs from most other equilibrium approaches where trades only occur after all iterations have been completed. In Kurose and Simha's approach, the solution gets better at every iteration, and is guaranteed to finally converge to the optimum. Recently, resource-based market mechanisms have been used and thoroughly studied in electricity distribution [18]. A sufficient condition for the applicability of our analysis to resource-based approaches is that reallocation is not performed until an equilibrium has been reached.

Classically in equilibrium markets, the agents are assumed to act *competitively*: they treat prices as exogenous. This means that each agent makes and reveals its demand (supply) decisions truthfully so as to maximize its utility (profit) given the market prices—assuming that it has no impact on those prices. The idea behind this *price-taking assumption* is that the market is so large that no single agent's actions affect the prices. However, this is paradoxical since the agents' declarations completely determine the prices. The price-taking assumption becomes valid as the number of agents approaches infinity: with infinitely many agents (of comparable size), each agent is best off acting competitively since it will not affect the prices.

However, in markets with a finite number of agents, an agent can act strategically, and potentially achieve higher utility by over/under representing [7, pp. 220–223], [4]. In doing so, the agent has to speculate how its misrepresentation affects the market prices, which are simultaneously affected by how other agents respond to the prices which changed due to the first agent's strategic actions. This paper addresses the question of how much an agent can gain or lose by such strategic behavior in such a complex setting of interrelationships.

The theory in this paper stems from well established principles in microeconomics—such as general equilibrium theory and the theory of optimal price setting in oligopoly and monopoly markets—and is extended for situations where agents have estimation errors in their beliefs regarding the behavior of other agents. First, a general, protocol independent study of the potential gains and losses from speculation is presented in Section 2. This is concretized for the case where only one agent is speculating under perfect information in Section 3, and under biased beliefs about others in Section 4. Simultaneous speculation by multiple agents is game theoretically discussed in Section 5. Finally, Section 6 concludes. Note that all numerical utility values in this paper have been multiplied by 100 for readability. In all, we believe that the methodology presented

in this paper is important for builders of computational markets where the agents represent self-interested real world parties that can tailor their agents so as to take advantage of the other agents in the system.

2 A Method for Analyzing the Potential Gains from Speculation

The goal of a self-interested consumer is to find the consumption bundle that maximizes its utility. To find the optimal bundle when acting in an equilibrium market, the consumer must speculate how other agents respond to prices. This is because its demand decisions affect the prices, which affect the demand and supply decisions of others, which again affect the prices that the consumer faces. Using the model of other agents, the consumer computes its optimal demand decisions. Note that other agents might also be speculating (in the same way or some other, suboptimal way). That is included in the agent's model of the other agents.

The goal of a self-interested producer is to find the production vector that maximizes its profits.² Again, this requires a model of how others react to prices because the producer's production decisions affect the prices, which affect the demand and supply decisions of others, which again affect the prices that the producer faces.

Using standard notation from microeconomics, let there be n commodities, and let the price vector be $\mathbf{p} = [p_1, p_2, \dots, p_n]$, where p_i is the price for good i . Let the resulting allocation for consumers be $\mathbf{x}_i(\mathbf{p}) = [x_{i1}(\mathbf{p}), x_{i2}(\mathbf{p}), \dots, x_{in}(\mathbf{p})]^T$, where x_{ij} is consumer i 's allocation of good j . Let the initial endowments be $\mathbf{e}_i = [e_{i1}, e_{i2}, \dots, e_{in}]^T$, where e_{ij} is consumer i 's endowment of good j . Now, consumer i 's excess (net) demand of good j is $z_{ij}(\mathbf{p}) = x_{ij}(\mathbf{p}) - e_{ij}$. Let $u_i(\mathbf{x}_i)$ be consumer i 's utility as a function of its allocation, and let θ_{ih} be the fraction of producer h that consumer i owns. The producers' profits are divided among consumers according to these shares. However, the consumers are assumed to have no say-so in the producers' production decisions.

Furthermore, let $\mathbf{y}_i(\mathbf{p}) = [y_{i1}(\mathbf{p}), y_{i2}(\mathbf{p}), \dots, y_{in}(\mathbf{p})]^T$ be the production vector of producer i , where y_{ij} represents the amount produced, and a negative number means that the good is an input. Let Y_i be the production possibilities set of producer i . The profit of producer i is $\mathbf{p} \cdot \mathbf{y}_i(\mathbf{p})$, where $\mathbf{y}_i \in Y_i$. For presentation uniformity with the case of the consumer, we define the excess demand of producer h to be $z_{hj}(\mathbf{p}) = -y_{hj}(\mathbf{p})$.

Let there be k agents in addition to the speculating agent that we investigate. The excess demand of these k agents for good j is

$$z_j^k(\mathbf{p}) = \sum_{i=1}^k z_{ij}(\mathbf{p}). \quad (1)$$

²This makes the standard assumption that the producer is able to alter its production plan costlessly during the search for equilibrium.

For the purposes of this section, we do not make any restricting assumptions about how these k agents make their supply/demand decisions (which determine the excess demands). In particular, we do not assume that agents act competitively. The speculating agent that we investigate uses its information about $z_j^k(\mathbf{p})$ as the basis of its strategic behavior as is now described.

The total excess demand with the speculating agent included is

$$z_j(\mathbf{p}) = z_j^k(\mathbf{p}) + z_{sj}(\mathbf{p}). \quad (2)$$

Once the market has reached a general equilibrium, supply meets demand, i.e. $z_j(\mathbf{p}) = 0$ for every good j .³ Substituting this into (2) gives

$$z_{sj}(\mathbf{p}) + z_j^k(\mathbf{p}) = 0. \quad (3)$$

2.1 Case A: Speculating Consumer

A solution to the following maximization problem gives the highest utility that a speculating consumer theoretically can obtain.

$$\begin{aligned} \max_{\mathbf{p}} u_s(\mathbf{x}_s(\mathbf{p})) \quad & \text{s.t.} \\ \mathbf{x}_{sj}(\mathbf{p}) & \geq 0 \quad (\text{consumer does not produce}) \\ \mathbf{x}_{sj}(\mathbf{p}) & = e_{sj} - z_j^k(\mathbf{p}) \quad (\text{supply meets demand}) \\ \mathbf{p} \cdot \mathbf{z}_s(\mathbf{p}) & = \sum_{h \in \text{producers}} \theta_{sh} \mathbf{p} \cdot \mathbf{y}_h(\mathbf{p}) \quad (\text{budget constraint}) \end{aligned} \quad (4)$$

provided that the equilibrium is unique and the market protocol finds it (this is discussed further in Section 3.2). The last equality assumes that either the speculating consumer's utility does not decrease as the amount of any good is increased in the consumer's consumption bundle, or the consumer can freely dispose of any good. Otherwise, " $=$ " should be changed to " $\leq$ ".

2.2 Case B: Speculating Producer

Similarly to the case of a speculating consumer, a solution to the following maximization problem gives the highest profit that a speculating producer can obtain.

$$\begin{aligned} \max_{\mathbf{p}} \mathbf{p} \cdot \mathbf{y}_s(\mathbf{p}) \quad & \text{s.t.} \\ \mathbf{y}_s(\mathbf{p}) & \in Y_s \quad (\text{feasible production plan}) \\ \mathbf{y}_{si} & = z_i^k(\mathbf{p}) \quad (\text{supply meets demand}) \end{aligned} \quad (5)$$

provided that the equilibrium is unique and the market protocol finds it. The last equality turns into $y_{si} \geq z_i^k(\mathbf{p})$ if free disposal [8] for both inputs and outputs is possible for each commodity.

We call the solution to the applicable optimization problem above (depending on whether the speculator is a producer or a consumer) $\mathbf{p}^*$.

³This holds even if agents are strategic—assuming that the particular market algorithm finds an equilibrium. The equilibrium reflects how the speculating agent is acting strategically, and how the other agents have reacted to the new price vector that came about due to the strategic agent's actions.

3 Strategic Behavior with Perfect Information

In the previous section we obtained a method for determining the highest utility that an agent theoretically can obtain. This section shows how this method can be used in practice.

3.1 A Simple Example

To demonstrate the method, we show how it applies to a specific case. We choose a setting with no producers. In such *pure exchange markets*, the consumers just reallocate their initial endowments among themselves [15; 5]. We choose a set of agents that is similar to the one described by Hu and Wellman [3]. Specifically, we let every agent—except for the speculating one that we investigate—be a competitive agent with constant elasticity of substitution, i.e. a utility function of the form

$$u_i(\mathbf{x}) = \left(\sum_{j=1}^n \alpha_{ij} x_{ij}^\rho \right)^{\frac{1}{\rho}} \quad (6)$$

where we have chosen $\alpha_{ij} = 1$ and $\rho = \frac{1}{2}$. Since these agents act competitively, and the speculating agent is assumed to have perfect information, the analysis of this example is protocol independent as long as the resources are reallocated after equilibrium (3) has been reached.

For simplicity and readability, we use only two goods ($n = 2$). The endowments are the same for all the competitive agents and they are 2 for good 1 and 1 for good 2. We let the speculating agent have the utility function

$$u_s(\mathbf{x}) = \sum_{j=1}^n \ln(x_{sj}), \quad (7)$$

and an endowment of 1 for both goods.

Most of this section is devoted to computing the gain from strategic behavior. For the casual reader, the main results of this section are condensed in Table 1 and Figure 1 below.

We get, e.g. from [17] and the definition of excess demand, that the excess demand of the competitive agents is

$$z_1^k(\mathbf{p}) = k \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right). \quad (8)$$

Because prices are only relative, we can set one of the prices arbitrarily, e.g. p_n can be set to 1, i.e. $p_2 = 1$. From the budget constraint ($\mathbf{p} \cdot \mathbf{z}(\mathbf{p}) = 0$) we then get $z_n(\mathbf{p}) = -\sum_{j=1}^{n-1} p_j z_j(\mathbf{p})$, i.e. $z_2^k(p_1) = -p_1 z_1^k(p_1)$. Using these and (7) and (8), we get

$$u_s(\mathbf{x}(p_1)) = \ln \left(1 - k \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right) \right) + \ln \left(1 + k \left(\frac{2p_1 + 1}{p_1 + 1} - 2p_1 \right) \right). \quad (9)$$

From (7) we see that x_{sj} must be greater than zero. If the speculating agent chooses to minimize p_1 , it should

sell as much of x_1 as possible and thus, as seen from the expression for x_{s1} in (9), and the requirement that $x_{s1} > 0$, we have $p_1 > p_1^{\min} = -\frac{1}{2(2k+1)} + \sqrt{\frac{k}{2k+1} + \left(\frac{1}{2(2k+1)}\right)^2}$. Analogous reasoning for x_{s2} shows that $p_1 < p_1^{\max} = \frac{1}{4k} + \sqrt{\frac{1}{2} \left(1 - \frac{1}{k}\right) + \left(\frac{1}{4k}\right)^2}$. If k approaches infinity, both $p_1^{\min}$ and $p_1^{\max}$ approach $\sqrt{\frac{1}{2}} \approx 0.707$. Therefore, with an infinite number of agents, the speculator cannot afford to affect the price in any way.

The first derivative of u_s with respect to p_1 is

$$\frac{\partial u_s}{\partial p_1} = \frac{-1}{1 - k \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right)} k \left(2 \frac{1}{p_1(p_1 + 1)} - \frac{2p_1 + 1}{p_1^2(p_1 + 1)} - \frac{2p_1 + 1}{p_1(p_1 + 1)^2} \right) + \frac{1}{1 - k \left(\frac{2p_1 + 1}{p_1 + 1} - 2p_1 \right)} k \left(\frac{2}{p_1 + 1} - \frac{2p_1 + 1}{(p_1 + 1)^2} - 2 \right). \quad (10)$$

It turns out that $\lim_{p_1 \rightarrow p_1^{\min}} \frac{\partial u_s}{\partial p_1} > 0$ and $\lim_{p_1 \rightarrow p_1^{\max}} \frac{\partial u_s}{\partial p_1} < 0$, and that the solution to $\frac{\partial u_s}{\partial p_1} = 0$ is unique in the interval $p_1^{\min} < p_1 < p_1^{\max}$. Therefore, the optimum, p_1^* , is obtained by solving $\frac{\partial u_s}{\partial p_1} = 0$.

The results of optimal strategic behavior are compared to the results of competitive behavior by the same agent. When the agent acts competitively, the excess demand (with u_s) is $z_{s1} = \frac{p_1 + 1}{2p_1} - 1$. Setting the excess demand to zero gives $k \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right) + \frac{p_1 + 1}{2p_1} - 1 = 0$. Solving for the competitive price gives $p_1^c = \sqrt{\frac{2k+1}{4k+1}}$. The results are shown in Table 1.⁴

k	$p_1^{\min}$	$p_1^{\max}$	p_1^*	p_1^c	$u_s(p_1^*)$	$u_s(p_1^c)$
1	0.4343	1.281	0.7601	0.7746	1.746	1.626
2	0.5403	1.000	0.7402	0.7454	2.202	2.152
5	0.6303	0.8262	0.7227	0.7237	2.614	2.602
10	0.6667	0.7670	0.7154	0.7157	2.788	2.784
20	0.6863	0.7372	0.7114	0.7115	2.884	2.884
30	0.6931	0.7272	0.7100	0.7160	2.918	2.918
100	0.7029	0.7131	0.7080	0.7080	2.967	2.967

Table 1: Acting strategically vs. acting competitively. k is the number of agents acting competitively. $p_1^{\min}$ is the price for x_1 in the market when the speculating agent sells as much x_1 as possible. $p_1^{\max}$ is the price for x_1 in the market when the speculating agent sells as much x_2 as possible. p_1^* is the market price as a consequence of strategic acting. p_1^c is the market price as a consequence of competitive acting by the same agent. The values u_s are the corresponding utilities for the agent under observation.

In Figure 1, the utility is plotted for the situations where the agent acts strategically and where it acts competitively. As expected (see e.g. [11]), the larger the number of agents, the smaller the gain from strategic behavior, and the less reason not to act competitively. In this example, already when the number of competitive

⁴Since utility is ordinal, one should be careful when discussing degrees of improvement.

agents is around five, the gain from strategic behavior is negligible.




Figure 1: Comparison of strategic and competitive utility, $u_s(p_1^*)$ and $u_s(p_c^*)$. The horizontal axis shows a number of interesting values rather than a specific scale. We see that the larger the number of agents, the smaller the gain from strategic behavior.

3.2 Reaching Equilibrium

Obviously an agent's best strategy is to declare an excess demand function such that the market will converge to the prices that are optimal for the speculating agent. More formally, when perfect information is available, an agent's best strategy—even if the other agents are not acting competitively, and some of them may be producers—is to declare an excess demand function with the property

$$z_{sj}^*(p^*) = -z_j^k(p^*), \quad (11)$$

for each good j , and which has a form such that the particular algorithm for searching for the market equilibrium converges to p^* . Depending on the setting, the difficulty of finding such a function varies. As shown below, finding it for the simple example from above where all other agents' supply/demand functions can be taken as fixed (to do this, they do not have to be competitive in general) is normally not hard. However, finding it may be more difficult if the other agents try to act strategically via speculation as well. These game theoretic issues will be addressed in Section 5.

Having computed the optimal speculative solution, $u_s^*(p_1^*)$, we would like to describe the strategic behavior leading to this solution under any particular market protocol used. For example, if p_1 is established via an algorithm whose only requirement for finding the equilibrium is $\partial z_1(p_1)/\partial p_1 < 0$,⁵ and we have that $\partial z_1^k(p_1)/\partial p_1 < 0$ (i.e. demand decreases as price increases—as is the case in our example above), we see that if

$$z_{s1}^*(p_1^*) = -z_1^k(p_1^*) \text{ and } \frac{\partial z_{s1}^*(p_1)}{\partial p_1} \leq 0, \quad (12)$$

⁵An example of such an algorithm is binary search. The WALRAS market framework uses this implementation.

there is a single solution for $p_1 = p_1^*$, and that solution will be found by the algorithm.⁶ It turns out that simple demand revelation strategies exist for the speculator which guarantee that an equilibrium will be reached where the speculator's maximal gain from speculation (derived earlier in this paper) materializes. The following linear function defines one such strategy—i.e. it fulfills (12):

$$z_{s1}(p_1) = \bar{p}_1^* - \bar{z}_1^k(p_1^*) - p_1, \quad (13)$$

where $\bar{z}_1^k(p_1^*)$ and $\bar{p}_1^*$ are the speculator's (perfect) estimates of $z_1^k(p_1^*)$ and p_1^* respectively. Another viable strategy is defined by the constant function

$$z_{s1}(p_1) = -\bar{z}_1^k(p_1^*). \quad (14)$$

4 Dealing with Uncertainty: Strategic Behavior under Biased Beliefs about Others

This section extends the discussion to include the impact of uncertainty on the speculating agent's strategy.

4.1 The General Setting

Above, when an excess demand function was chosen based on perfect information, the exact form of the function was unimportant as long as it fulfilled (12). However, if the speculating agent cannot estimate z_1^k perfectly, its outcome will depend on the function chosen, i.e. the choice of an excess demand function will depend on the probability distribution that characterizes the uncertainty in the speculator's beliefs. Rather than analyzing how the excess demand is chosen for different uncertainty situations and how the speculator should try to reduce the uncertainty by learning about its environment in an efficient way, we describe how, for a specific choice of z_{sj} , a specific error in the estimation of the competitive agents' behavior affects the outcome of the market.

Independently of how the z_{sj} function is chosen, the possible market outcomes can be determined by solving (3). If no solution exists, of course, no algorithm can find it, and if multiple equilibria exist, a protocol dependent analysis is required to find out which one will be chosen.

4.2 The Example Revisited

Returning to the above example, we first assume that z_{s1} has been chosen as describe by (13), but that now there is an error in the estimation. We define this error

⁶The above reasoning extends easily to a market with more than two goods. If p is established via an algorithm whose only requirements for finding the equilibrium are $\partial z_i(p)/\partial p_i < 0$ and $\partial z_i(p)/\partial p_j \geq 0$, $i \neq j$ (e.g. WALRAS), and we have that $\partial z_i^k(p)/\partial p_i < 0$ and $\partial z_i^k(p)/\partial p_j \geq 0$, $i \neq j$, then (12) should be generalized to $z_{si}^*(p^*) = -z_i^k(p^*)$ and $\partial z_{si}^*(p)/\partial p_i \leq 0$ and $\partial z_{si}^*(p)/\partial p_j \geq 0$, $i \neq j$.

learn more about other agents' excess demand/supply functions. If many agents are involved in such probing, it seems that time becomes an important factor. Some agents might reveal progressively more of their competitive demands in order to speed up the convergence (as it might be urgent for them to get the resources traded), while others might extend the probing in order to maximize their benefit from the trade.⁹

Once the other agents' strategies are known (e.g. from a game theoretic equilibrium analysis), the methods of this paper can be used to analyze the speculating agent's strategy alternatives. The methods are protocol independent, and they can be used to estimate the potential gains from speculation in any particular setting, as well as to determine how far from the optimal strategy a particular strategy is—as long as the other agents' strategies can be fixed conceptually. This does not mean that they need to be known with certainty.

The methods of this paper can also be used when game theoretic analysis fails. Especially when speculation is based on expected actions of other agents—instead of a game theoretic equilibrium analysis of strategies—the theory of speculation under biased beliefs is highly applicable. In addition, the example market used in this paper suggests that if a market is of at least moderate size, the gain from speculation is so small, and the risk of being significantly worse off due to estimation error is so great, that one can expect that each agent will act close to its competitive behavior.

6 Conclusions

We presented a method for computing the maximal advantage of speculative strategic behavior in general equilibrium based markets. It is computed from the other agents' supply/demand functions (classic competitive behavior by the other agents is a special case of this). The method enables one to analyze how much an agent could gain or lose by speculating in a particular system, and it is also useful when evaluating different strategies since it allows one to determine how close to the optimal strategy they are. Our analysis is valid for a wide variety of known market protocols. We also constructed demand revelation strategies that guarantee that an agent can drive the market to an equilibrium where the agent's maximal advantage from speculation materializes.

Our study of a particular market shows that as the number of agents increases, the gains from speculation decrease—often turning negligible already at moderate numbers of agents. The study also shows that under

⁹Some work has addressed non-competitive behavior in WALRAS [3], although there was only one speculating agent in the experiments, and this agent was limited to simple linear price prediction about how its actions affect the prices. Further analysis is required to determine whether its optimal strategy can be captured in this model. This need not be the case because the optimal strategy may involve some more "aggressive" behavior, e.g. the probing described above.

uncertainty regarding other agents, competitive acting is often close to optimal, while speculation can make the agent significantly worse off—even if the agent's beliefs are just slightly biased.

We believe that computational agents representing self-interested real world parties will deviate from competitive behavior in practice only if the potential gain from speculation is sufficiently large compared to the cost of the computation or information gathering actions required for speculation. We also believe that an agent will not speculate when that incurs potentially large losses when estimation errors are present. Merging the above results with these beliefs suggests that speculation problems diminish as the market grows or the agents' information about others becomes less certain.

Finally, we discussed the protocol dependent game theoretic issues related to multiple agents counterspeculating. There are many open issues in this area and the difficulties should not be underestimated.

Acknowledgments

The authors thank Arne Andersson, Junling Hu, John Nachbar, Eric Schenk, Fernando Tohmé, Michael Wellman, and Curt Wells for interesting discussions and comments. Fredrik also thanks Hans Akkermans, Rune Gustavsson and Hans Ottosson for all their support.

References

- [1] S. Barbera and M. O. Jackson. Strategy-proof exchange. *Econometrica*, 63(1):51–87, 1995.
- [2] J. Q. Cheng and M. P. Wellman. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, 1997. To appear.
- [3] J. Hu and M. P. Wellman. Self-fulfilling bias in multiagent learning. In *Proc. of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 118–125, Keihanna Plaza, Kyoto, Japan, Dec. 1996.
- [4] L. Hurwicz. On informationally decentralized systems. In C. McGuire and R. Radner, editors, *Decision and Organization*, chapter 14, pages 297–336. University of Minnesota Press, 1986. Second edition.
- [5] D. Kreps. *A Course in Microeconomic Theory*. Princeton U. Press, 1990.
- [6] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705–717, 1989.
- [7] E. Malinvaud. *Lectures on Microeconomic Theory*. North-Holland, 1985.
- [8] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [9] T. Mullen and M. P. Wellman. A simple computational market for network information services. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pages 283–289, San Francisco, CA, June 1995.
- [10] J. Nash. Equilibrium points in n -person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950.
- [11] D. J. Roberts and A. Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica*, 44(1):115–127, 1976.
- [12] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [13] T. W. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 299–306, Keihanna Plaza, Kyoto, Japan, Dec. 1996.
- [14] T. W. Sandholm and V. R. Lesser. Advantages of a leveled commitment contracting protocol. In *AAAI-96*, pages 126–133, Portland, OR, Aug. 1996. Extended version: UMass Comp. Sci. technical report 95-72.
- [15] H. R. Varian. *Microeconomic analysis*. New York: W. W. Norton, 1992.
- [16] M. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.
- [17] M. Wellman. A computational market model for distributed configuration design. In *AAAI-94*, pages 401–407, Seattle, WA, July 1994.
- [18] F. Ygge and J. Akkermans. Power load management as a computational market. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 393–400, Keihanna Plaza, Kyoto, Japan, Dec. 1996.

through⁷

$$\overline{p_1^*} - \overline{z_1^k(p_1^*)} = (1 + e)(p_1^* - z_1^k(p_1^*)). \quad (15)$$

Provided that the speculator stops learning about $z_1^k(p_1)$ at some point in time, and hereby fixes $z_{s1}(p_1)$, any market protocol whose only requirement for finding the equilibrium is that $\partial z_1(p_1)/\partial p_1 < 0$ (e.g. the binary search used in WALRAS) is guaranteed to converge independently of how large the estimation error is, since (12) is satisfied. Inserting the expressions from (8), (13), and (15) in (3) gives

$$(1 + e) \left(p_1^* - k \left(\frac{2p_1^* + 1}{p_1^*(p_1^* + 1)} - 2 \right) \right) - p_1 + k \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right) = 0. \quad (16)$$

This equation was solved for different errors e , and different numbers of competitive agents k . The corresponding utilities—normalized to the utility that would result from competitive behavior—are plotted in Figure 2.




Figure 2: The utility as a function of the estimation error. k is the number of competitive agents.

As seen from Figure 2, when the number of agents is small, the gain from speculating is large; even when there is an error involved, it often pays off to speculate. But as the number of agents increases, acting competitively is close to acting strategically with perfect information, and it is worse to speculate, even with surprisingly small errors. The break even points for 1, 2, 5, 20, and 100 agents are at an estimation error of approximately $\pm 5\%$, $\pm 3\%$, $\pm 1\%$, $\pm 0.4\%$, and $\pm 0.05\%$ respectively. Looking the other way around, with 100 competitive agents and a $\pm 5\%$ error (the break even error for 1 competitive agent), the loss compared to acting competitively is substantial.

5 Strategic Behavior by Multiple Agents

In our simple example, only one agent was speculating and the others were acting competitively. Even though

⁷If the speculating agent can learn about the other agents and change its excess demand during the market process, the error might decrease during this process. The error described here is the error remaining when the process terminates.

(1)–(5) and (11)–(15) are valid even in the case where every agent is speculating, determining market outcomes in a protocol independent way is not possible in such settings in general. The reason is that an agent cannot treat others' strategies (e.g. policies for revealing excess supply/demand functions) as fixed because the others would like to tailor their strategies to the specific strategy that the agent chooses. The strategies are in *Nash equilibrium* [10; 8; 5] if each agent's strategy is its best response to the others' strategies. This can be viewed as a necessary condition for system stability in settings where all agents act strategically.⁸ In sequential protocols, one can also strengthen the Nash equilibrium solution concept in multiple ways by requiring that the strategies stay in equilibrium at every step of the game [8; 5]. Unlike our analysis, the Nash equilibrium outcome is specific to the market protocol. Important factors impacting the outcome are the order in which bids are submitted (see e.g. Stackleberg vs. Cournot models [8]), whether the bids are sealed or open [13], whether the protocol is iterative (the agents can change their excess demand between iterations) or not, whether the agents can decommit from their agreements by paying a penalty [14], etc.

In some games, no Nash equilibrium exists in pure (non-randomized) strategies. The following simple example illustrates this. Let there be two consumer agents, A and B, that engage in a market where they reveal their excess demand functions simultaneously and in a single round. Agent A can choose between two strategies (A1 and A2), and B can choose between B1 and B2. Provided that A knows that B will choose B1, A will choose A2, and A1 if B chooses B2. Provided that B knows that A will choose A2, B will choose B2, and B1 if A chooses A1. Now, from every possible pair of strategies, one agent would be motivated to deviate to another strategy, i.e. no Nash equilibrium exists. In general, existence and uniqueness of a general equilibrium (where agents act competitively) for a market does not imply existence and uniqueness of a Nash equilibrium (or an individually rational non-cooperative equilibrium [7]).

We argue that for each protocol proposed for implementation of equilibrium markets including self-interested computational agents, a thorough game theoretic analysis should be attempted [14; 13; 12]. However, as discussed above, some games lack a Nash equilibrium, or it may not be unique. In addition, some protocols are hard to analyze game theoretically. For example, in WALRAS, the agents might change their demand functions during the computation of the equilibrium. Then some agents may deliberately send false bids to generate more iterations of the market process in order to

⁸A stronger condition of stability is to require incentive compatibility, i.e. that an agent's strategy is optimal (for the agent) no matter what strategies others choose. Market protocols have been studied using this solution concept in [1].