

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCS-99-26

1999-01-01

Constructing Speculative Demand Functions in Equilibrium Markets

Tuomas Sandholm and Fredrik Ygge

In computational markets utilizing algorithms that establish a 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. Determining the optimal strategy when the speculator has perfect knowledge about the other agents is a well known problem which has been studied... [Read complete abstract on page 2.](#)

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Sandholm, Tuomas and Ygge, Fredrik, "Constructing Speculative Demand Functions in Equilibrium Markets" Report Number: WUCS-99-26 (1999). *All Computer Science and Engineering Research*. https://openscholarship.wustl.edu/cse_research/496

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

Constructing Speculative Demand Functions in Equilibrium Markets

Tuomas Sandholm and Fredrik Ygge

Complete Abstract:

In computational markets utilizing algorithms that establish a 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. Determining the optimal strategy when the speculator has perfect knowledge about the other agents is a well known problem which has been studied in oligopoly theory in economics. We describe the computation of such a strategy, and focus on an issue that has received little attention in economics, but which is of fundamental importance in computational markets: the revelation of demand strategies that drive the market to the desired equilibrium. The more realistic setting where the speculator has imperfect information about the other agents is more delicate. We demonstrate how speculation under biased beliefs about the other agents can result in considerable losses if traditional oligopoly strategies for perfect information are used. Furthermore, we show how the optimal demand is computed from probability distributions on the other agents' supply/demand functions. We also theoretically show when an optimal revealed demand function can be constructed independently of the probability distributions. Some pragmatics of choosing a demand function in the case of imperfect information (particularly useful for construction of computational agents for equilibrium markets) are given, and we show - with some empirical support - that it can be relatively easy to construct demand functions that results in a gain from speculation, even when estimation errors are large. Finally, game theoretic issues related to multiple agents counterspeculating simultaneously are discussed.

**Constructing Speculative Demand Functions
in Equilibrium Markets**

Tuomas Sandholm and Fredrik Ygge

WUCS-99-26

October 1999

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St. Louis MO 63130**

Constructing Speculative Demand Functions in Equilibrium Markets

Tuomas Sandholm *
Department of Computer Science
Washington University
St. Louis, MO 63130-4899, USA
www.cs.wustl.edu/~sandholm

SANDHOLM@CS.WUSTL.EDU

Fredrik Ygge *
EnerSearch AB and Uppsala University
Chalmers Science Park
S-412 88 Gothenburg, Sweden
www.enersearch.se/ygge

YGGE@ENERSEARCH.SE

Abstract

In computational markets utilizing algorithms that establish a 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. Determining the optimal strategy when the speculator has perfect knowledge about the other agents is a well known problem which has been studied in oligopoly theory in economics. We describe the computation of such a strategy, and focus on an issue that has received little attention in economics, but which is of fundamental importance in computational markets: the revelation of demand strategies that drive the market to the desired equilibrium.

The more realistic setting where the speculator has imperfect information about the other agents is more delicate. We demonstrate how speculation under biased beliefs about the other agents can result in considerable losses if traditional oligopoly strategies for perfect information are used. Furthermore, we show how the optimal demand is computed from probability distributions on the other agents' supply/demand functions. We also theoretically show when an optimal revealed demand function can be constructed independently of the probability distributions. Some pragmatics of choosing a demand function in the case of imperfect information (particularly useful for construction of computational agents for equilibrium markets) are given, and we show—with some empirical support—that it can be relatively easy to construct demand functions that results in a gain from speculation, even when estimation errors are large. Finally, game theoretic issues related to multiple agents counterspeculating simultaneously are discussed.

*. This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, Grant IIS-9800994, NUTEK's Promodis program, and EnerSearch's owners: ABB, ECN, Iberdrola, IBM, PreussenElektra, and Sydkraft. A short early version containing some of the results of this article appeared at the International Joint Conference on Artificial Intelligence (Sandholm & Ygge, 1997).




Figure 1: *High-level view of an equilibrium market mechanism.*

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. (Wellman, 1993, 1994; Mullen & Wellman, 1995; Ygge & Akkermans, 1996; Yamaki, Wellman, & Ishida, 1996; Ygge, 1998; Cheng & Wellman, 1998; Kurose & Simha, 1989). It provides a distributed method for efficiently allocating goods and resources among agents. The methodology allows for striking efficient tradeoffs across multiple goods and multiple agents.

1.1 The principles of equilibrium markets

The commodities in an equilibrium market can be physical, e.g. coffee and meat, or they can be more abstract, e.g. parameters of an airplane design (Wellman, 1994), flows in a traffic network (Wellman, 1993), electricity in a power network (Ygge & Akkermans, 1996), or mirror sites on the Internet (Mullen & Wellman, 1995).

The basic form of an equilibrium market mechanism is shown in Figure 1. First the auctioneer announces a new auction, possibly accompanied by some information for example about expected prices. Then agents respond with demand functions, telling what *change* in allocation they desire at different price alternatives (the demand may be negative, and hence represent supply). Based on the functions submitted by the agents, the auctioneer tries to establish a market clearing price (or rather a *set of prices*) such that demand equals supply for each commodity. Optionally the auctioneer may need to iteratively request further demand and supply information from the agents before equilibrium can be established. When an equilibrium price has been found, the commodities are reallocated among the agents as described by their demands and the market clearing prices.

1.2 Agent behavior in equilibrium markets

In an equilibrium market, the optimal behavior of a self-interested agent is to submit a demand function, which maximizes its utility after the trade has taken place. If an agent has negligible effect on prices (and does not form a coalition with other agents of such size that the coalition has any significant effect on prices), then clearly the optimal behavior of the agent is to give a demand function, which maximizes the agent's utility taken prices as given. This agent behavior is referred to as *price-taking* or *competitive* (Varian, 1992, p. 25). So far, competitive behavior has been a standard assumption in computational equilibrium

markets (Wellman, 1993, 1994; Mullen & Wellman, 1995; Ygge & Akkermans, 1996; Yamaki et al., 1996; Cheng & Wellman, 1998).

If an agent's behavior *does* affect prices, the agent's optimal behavior is no longer to act competitively (Hurwicz, 1972; Roberts & Postlewaite, 1976). The assumption that an agent's behavior does not affect the prices is unrealistic in many practical applications of equilibrium markets.

The case where an agent has perfect information about the revealed demand functions of the other agents is well studied in economics, see e.g. (Mas-Colell, Whinston, & Green, 1995, Chapter 12). However, the convergence aspects discussed in this paper have received little or no attention in economics. The more realistic case, where there is an uncertainty about the demand function of the other agents, has more recently been in focus in economics (Klemperer & Meyer, 1989) and is not as well understood.

1.3 Electricity markets – a prototypical application area

Electricity exchanges are an application area in which equilibrium markets are often used, and for which they are quite well suited. For example, the main electricity market in the Nordic countries, NordPool, see www.nordpool.no, is based on a mechanism in which all participating agents submit demand functions for the respective hours, and a market clearing price is established for each hour.

Because there are significant constraints in a power grid (having to do with transmission losses as well as the fact that the current that enters any given node has to equal the current that leaves that node), nodal pricing is usually used. That is, the prices are different at different nodes of the network. Therefore, some players often have very significant market power in specific parts of the network because they constitute a considerable portion of the aggregate supply or demand at specific nodes (Borenstein & Bushnell, 1998; Ilic, 1999). Such parties may want to use supply/demand strategies that capitalize on that market power instead of using strategies that take prices as exogeneously given. Such strategic acting is made more likely by the fact that the monetary stakes are large in electricity markets while the computation capability that is required to speculate profitably is modest and inexpensive. Put together, agents in electricity markets may be very advanced and use complex strategies for utilizing their market power. This article provides an algorithmic blueprint for constructing agents that profitably act strategically on behalf of the real-world parties that they represent.

In electricity markets, each agent normally has some uncertainty about the demand functions of the other agents. For example, although an agent with market power may have very good estimates of the costs of production and utility of consumption of the other agents, it may be hard to predict the behavior of these agents.¹ For example, the other agents might form different coalitions (see further the example in Section 4.2.3). In this article we show how an agent with market power can reduce the risk of potential losses because of speculation in the presence of uncertainty in such important settings by means of relatively simple parameterized demand functions (Section 4.2.5). Furthermore, we give

1. The reason that it might be realistic to assume that the agent has good estimates of true costs and utilities is that it often has knowledge about, for example, what other generators there are in a certain area, and the characteristics of standard generators are commonly known.

a realistic example where the uncertainty can be completely suppressed by means of our Theorem 4.1.

1.4 Contributions of the article

The general equilibrium framework has been ubiquitously used in electronic markets between computational agents in a variety of applications, many of which are related to the allocation of computational resources such as CPU-time, storage in distributed operating systems, bandwidth, *etc.* (Wellman, 1993, 1994; Mullen & Wellman, 1995; Ygge & Akkermans, 1996; Yamaki et al., 1996; Cheng & Wellman, 1998). Until now, most of that work has assumed that agents act as price-takers. However, it is clear that agents are not in general motivated to act according to this assumption, but to act strategically.

Quite early on, economists studied how much an agent can gain by deviating from price-taking behavior (Hurwicz, 1972; Roberts & Postlewaite, 1976). More recently this issue has been investigated also in the presence of uncertainty (Klemperer & Meyer, 1989). Thus, determining the gain from strategic behavior in these different settings is not new, though it has received relatively limited attention in the construction of equilibrium markets.

This article focuses on the construction of speculative demand functions. That is, it deals with how an agent can benefit from deviating from price-taking behavior in different market settings. Though some contributions (e.g. Theorem 4.1) are interesting from a pure economics perspective, the focus is on dealing with issues that have a practical value for constructing computational agents for equilibrium markets. The main high-level questions that we address are the following:

- What kind of strategic demand functions can an agent use to make the market algorithm *converge* to a solution that maximizes the speculator's gain, and looks like a real general equilibrium to the other market participants and the auctioneer?
- What methods can an agent use to *construct* such a strategic demand function algorithmically?

The answers vary depending on the number of goods, the other participants' demand functions, the speculator's information, *etc.*

Compared to the existing literature on oligopoly *etc.* in economics, which aims at describing and explaining real world phenomena among human market participants, this article investigates the use of equilibrium mechanisms for actually constructing markets (particularly computational markets for multiagent systems). The main implication of this is that whereas the work in economics pays little attention to how equilibrium is reached from the perspective of speculation, this issue is of central importance for our purposes. The speculator needs to be careful not to preclude the market from converging under the given market algorithm that is used for finding the equilibrium. We determine demand revelation strategies that drive the market to an equilibrium that maximizes the speculator's gain and looks like a real general equilibrium to the other agents. This holds for most market algorithms that can be used for finding an equilibrium. This convergence question is one type of interesting interplay between speculation strategies and computation.

Another type of interesting interplay between speculation strategies and computation arises from the agent trying to compute its optimal strategy. We show how to do this both

for complete and for incomplete information. We also look at the size of the losses that an agent can suffer if it uses biased information as if that information were exact. More interestingly, we present Theorem 4.1 which states sufficient conditions on the uncertainty of the demand of others such that the speculator can construct an optimal demand function independently of the probability distribution of this uncertainty. In connection with this, Theorem 4.2 gives sufficient conditions for enabling the construction of a demand function that fulfills the convergence requirements of most equilibrium market algorithms.

Some pragmatics of choosing a demand function under imperfect information are also given, suitable for cases where the optimal demand function cannot be computed (for computational reasons and/or when there only is some rough "gut feeling" about risks and the demands of the other agents). We show—with some empirical support—that it can be relatively easy to construct demand functions that are significantly better than traditional strategic demands from oligopoly theory, and which result in a gain from speculation, even when estimation errors are rather large.

In all, we believe that the methodology presented in this article is important for building software agents that act in equilibrium markets. It is also useful for builders of the computational markets themselves where the participating agents represent self-interested real world parties who can tailor their agents so as to take advantage of the other agents in the system.

1.5 Organization of the article

The article is organized as follows. First, a background—introducing the formalism and the economic concepts required for the rest of the paper—is given Section 2. The main sections of the article—Section 3 and Section 4—deal with the construction of demand functions for the case of perfect and imperfect information respectively. Relevant game theory related to simultaneous speculation by multiple agents is discussed in Section 5. Section 6 presents conclusions. The computation of the optimal demand with perfect information is demonstrated on a specific example in Appendix A, and different market algorithms (mechanisms) for general equilibrium markets are discussed in Appendix B.

2. Background

This section presents the notation for equilibrium markets and introduces the microeconomic theory that is needed to understand the rest of the article. Those readers that are very familiar with equilibrium markets and microeconomics may safely skip this section after having consulted the notational conventions captured in Table 1 of Section 2.5.

2.1 Formal definition of an equilibrium market

An equilibrium market has a finite number of *commodity goods* $g \in [1, 2, \dots, k]$. We make the standard assumptions that the amount of each commodity is unrestricted, and each commodity is arbitrarily divisible, i.e. continuous as opposed to discrete. Different elements within a commodity are not distinguishable, but different commodities are distinguishable from each other. The market also has prices $\mathbf{p} = [p_1, p_2, \dots, p_k]$, where $p_g \in \mathfrak{R}$ is the price of commodity g .

Without loss of generality, one commodity (let its index be the highest one, k) is regarded as a numeraire, the unit price of which is one, i.e. $p_k = 1$. This way the prices of all other commodities can be expressed in terms of this commodity. For example, we can use US dollars as this commodity and then say that the price of something else is, e.g., \$7 per unit. When participating in an equilibrium market, each agent submits (parts of and/or samples of) a net demand function, denoted $\underline{z}_i(\mathbf{p})$, $\underline{z}_i : \mathfrak{R}^{k-1} \mapsto \mathfrak{R}^{k-1}$, where i denotes an agent.² Thus, the net demand function describes the demand for $k-1$ commodities at $k-1$ prices (the remaining price being fixed). For example, $\underline{z}_i([1, 2, 1]) = [-1, 3, 0]$ means that at the prices 1, 2, and 1 for commodities 1, 2, and 3 respectively, agent i wants to sell 1 unit of commodity 1 and buy 3 units of commodity 2 for 5 units of commodity 4 ($-1 \cdot 1 + 2 \cdot 3 = 5$). Maybe the agent wants to sell one gallon of gas and buy three gallons of milk for \$5. This example describes a consumer exchanging some commodities for others, but \underline{z} can be used for denoting production as well (a produced unit is a negative demand of one).

The process of submitting parts of the demand function may be iterated if the market clearing price vector is outside the region captured by the submitted demand functions. One such process is the basic price tâtonnement process (Appendix B) in which point, $\underline{z}_{ig}(p_g)$, $\underline{z}_{ig} : \mathfrak{R} \mapsto \mathfrak{R}$, for the respective commodities, g , are sent to an auctioneer. Each of those demands is based on the current prices of the other commodities. That is, if those other prices change, a set of new demand functions may need to be submitted.

Once the auctioneer has established a market clearing price, say \mathbf{p}^* , defined by

$$\sum_i \underline{z}_{ig}(\mathbf{p}^*) = 0, \quad g = [1..k-1], \quad (1)$$

each agent will receive (or deliver) $\underline{z}_{ig}(p_g^*)$ of each of the $k-1$ commodities, g , and $\sum_{g=0}^{k-1} -\underline{z}_{ig}(p_g^*)p_g^*$ of the k :th commodity.

The analysis in this article is based on the assumption that a mechanism is used that establishes a market price for each commodity such that supply meets demand, and that reallocation is performed *after* these prices have been established, see Figure 1. There are many market algorithms that can be used to find such an equilibrium. Clearly, if no such equilibrium exists, no algorithm can find it. In this article 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 independent of the market mechanism—as long as the agents exchange commodities 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. An overview of useful market algorithms for finding equilibria, the algorithms' convergence properties, and other computational aspects is presented in Appendix B.

2. In a market with k commodities, it is possible to regard the net demand function as a mapping $\mathfrak{R}^k \mapsto \mathfrak{R}^k$. However, if one price is fixed (which can always be done without loss of generality) and the demand of the k th commodity is given by the budget constraint (discussed later), treating $k-1$ dimensions is always sufficient. Throughout the article, whenever $k-1$ commodities are considered, the k th commodity is assumed to be implicitly given by the above rules, with $p_k = 1$.

2.2 Basic general equilibrium theory

A market in the general equilibrium framework can have two types of agents, *consumers* and *producers*. Each consumer i has a *utility function* $u_i(\mathbf{x}_i)$ which encodes its preferences over different *consumption bundles* $\mathbf{x}_i = [x_{i1}, x_{i2}, \dots, x_{ik}]^T$, where $x_{ig} \in \mathfrak{R}_+$ is consumer i 's allocation of commodity g . Each consumer i also has an initial *endowment* $\mathbf{e}_i = [e_{i1}, e_{i2}, \dots, e_{ik}]^T$, where $e_{ig} \in \mathfrak{R}_+$ is his endowment of commodity g . The excess (net) demand of consumer i for commodity g is

$$z_{ig}(\mathbf{p}) = x_{ig}(\mathbf{p}) - e_{ig}. \quad (2)$$

In the context of general equilibrium, $x_{ig}(\mathbf{p})$ is agent i 's optimal choice/allocation at the given prices \mathbf{p} (cf. Equation (5) below), and $z_{ig}(\mathbf{p})$ hence denotes the agent's optimal *change* in allocation.

The producers—if any—can use some commodities to produce others. Let the vector $\mathbf{y}_j = [y_{j1}, y_{j2}, \dots, y_{jk}]^T$ be the *production vector*, where y_{jg} is the amount of commodity g that producer j 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* Y_j , which is the set of feasible production vectors. The *profit* of producer j is $\mathbf{p} \cdot \mathbf{y}_j$, where $\mathbf{y}_j \in Y_j$. 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). Let θ_{ij} be the fraction of producer j that consumer i owns. The producers' profits are divided among consumers according to these shares. The consumers are assumed to have no say-so in the producers' production decisions.³ For presentation uniformity with the case of the consumer, we define the excess demand of producer j (its optimal change in allocation at given prices, cf. Equation (6) below) to be

$$z_{jg}(\mathbf{p}) = -y_{jg}(\mathbf{p}). \quad (3)$$

Prices may change, and the agents may change their consumption and production plans, but actual production and consumption only occur once the market has reached an equilibrium. The market $(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$ is said to be in a *general (Walrasian) equilibrium* if

I markets clear: for each commodity, production plus endowments equals consumption.

Formally,

$$\sum_i \mathbf{x}_i^* = \sum_i \mathbf{e}_i + \sum_j \mathbf{y}_j^*, \text{ and} \quad (4)$$

II each consumer, i , consumes a bundle of commodities such that she could not afford another bundle of higher utility given her initial endowments, the current prices, and the profits she receives from producers. Formally,

$$\mathbf{x}_i^* = \arg \max_{\mathbf{x}_i \in \mathfrak{R}_+^k \mid \mathbf{p}^* \cdot \mathbf{x}_i \leq \mathbf{p}^* \cdot \mathbf{e}_i + \sum_j \theta_{ij} \mathbf{p}^* \cdot \mathbf{y}_j} u_i(\mathbf{x}_i), \text{ and} \quad (5)$$

3. Typically share holders will have a say so in the management of the company which has (long term) effects on the production decision. However, they do not participate in the day-to-day operation of the production which affects the short term equilibrium (Takayama, 1985) that we study in this article.

III each producer, j , uses the feasible production vector that maximizes his profits given the prices:

$$y_j^* = \arg \max_{y_j \in Y_j} \mathbf{p}^* \cdot y_j \quad (6)$$

General equilibrium solutions have some very desirable properties. First of all, if agents act competitively as price takers, each general equilibrium is *Pareto efficient*, i.e. no agent can be made better off without making some other agent worse off (Mas-Colell et al., 1995). This means that there is no possible search methodology for finding solutions to the agents' problem such that every agent is better off than in the general equilibrium.

The solution is also stable against collusion. Each general equilibrium with no producers is stable in the sense of the *core* solution concept of coalition formation games: no subgroup of consumers can increase their utilities by pulling out of the equilibrium and forming their own market (Mas-Colell et al., 1995).⁴

Unfortunately, in some settings no general equilibrium exists. As an example, suppose there is a market in which a producer has two machines, each with different marginal costs and some minimal base production. If only the least expensive machine is switched on, it might be the case that the price that the demand side is willing to pay is higher than the marginal cost for the second machine. However, the result of switching both machines on can be that there is too much supply and therefore market price drops below the marginal cost for the second machine. Hence, there is no price such that supply equals demand.

However, sufficient conditions for existence of an equilibrium can be stated on the excess demand functions $z_i(\mathbf{p})$:

Proposition 2.1 (Existence of equilibrium) *Let*

$$S^{k-1} = \left\{ \mathbf{p} \in \mathfrak{R}_+^k : \sum_{g=1}^k p_g = 1 \right\}.$$

If \mathbf{z} ($= \sum_i z_i$) : $S^{k-1} \mapsto \mathfrak{R}^k$ is a continuous function that satisfies Walras' law, $\mathbf{p} \cdot \mathbf{z}(\mathbf{p}) \equiv 0$, then there exists some \mathbf{p}^ in S^{k-1} which is a Walrasian equilibrium (Varian, 1992, pp. 319–322).*

The conditions of Proposition 2.1 are violated if demand is discontinuous. This can occur for example in network bandwidth allocation when an agent's utility jumps as the threshold for being able to participate in a video conference is reached. A second example of this is a scenario where a producer is able to produce either commodity A or commodity B at the same cost. As soon as the market price of A exceeds the market price of B , she would therefore produce A , and vice versa. Then, as soon as the price of A is only slightly higher than the price of B , the producer will only sell A . This would cause the price of A to fall below that of B , and the provider would prefer to switch to producing B .

Even if an equilibrium exists, it might not be unique. It is easy to construct examples that have multiple equilibria, see e.g. (Kehoe, 1991, pp. 2066). However, there is an easily understood sufficient condition for uniqueness:

4. The situation becomes more complex if producers are present: for example, if a set of consumers colludes, and they own part of a producer via the shares, what can the coalition produce? This issue is discussed further for example in (Ellickson, 1993).

Proposition 2.2 (Uniqueness of equilibrium under gross substitutes) *An equilibrium is unique if the aggregate demand for each commodity is nondecreasing in the prices of the other commodities (Mas-Colell et al., 1995).*

For example, as the price of spaghetti increases, consumers have to convert to satisfying their hunger with less expensive foods. It follows that the demand of tagliatelle increases. Thus, spaghetti and tagliatelle are gross substitutes.

On the other hand, the conditions of this proposition are not always met. For example, as the price of bread increases, the demand of butter decreases. Such *complementarities* are also common in production, where the producers often need all of the inputs to create the outputs.

The following proposition is somewhat more general and covers Proposition 2.2 (Mas-Colell et al., 1995, p. 623).

Proposition 2.3 *An equilibrium is unique if $z(\check{p}) = 0$ and $\check{p} \cdot z(p) > 0$ for all p not proportional to \check{p} .⁵*

Proof. Assume that \underline{p} is a market clearing price (which is not proportional to \check{p}). Then $\check{p} \cdot z(\underline{p}) = 0$, which violates the precondition $\check{p} \cdot z(p) > 0$. Hence the proposition holds. \square

The basic general equilibrium framework does not account for *externalities*. In *consumption externalities*, one agent's consumption affects another agent's utility. In *production externalities*, one agent's production possibilities set is directly affected by another agent's actions.

Some mechanisms to attack externality problems include taxes and viewing some of the externality issues as commodities themselves (Varian, 1992, Chapter 24).

2.3 Acting competitively vs. acting strategically

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 (each of small 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 its demand (Malinvaud, 1985, pp. 220-223), (Hurwicz, 1972). In doing so, the agent has to speculate how its

5. This is under the assumption that both z and p are given in k dimensions and that no price is considered fixed. If one price is fixed (which can always be done without loss of generality), the "proportional to" changes to "equal to", which is a stronger claim. If k dimensions are used and no price is fixed, then all price vectors proportional to \check{p} result in the same equilibrium in terms of allocations.

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.

Under the assumption that each agent acts competitively, the outcome of the market mechanism described in Section 2.1 is equivalent to the general equilibrium (cf. Section 2.2). However, in markets with a finite number of agents, competitive behavior need not be optimal. Then there may be a difference between the *revealed* demand, \underline{z} (of Section 2.1), and the competitive demand, z (of Equation (2) and Equation (3)), since there can be a profit from revealing a demand that deviates from the competitive demand. The outcome of the market of Section 2.1 is then not a (competitive) general equilibrium.⁶

In accordance with existing literature on equilibrium markets, e.g. (Cheng & Wellman, 1998), we allow ourselves to call the revealed demand the *net excess demand*. We, however, denote it by \underline{z} , rather than just z , since z is often used to denote competitive demand.

One should carefully keep in mind that the different properties of net demand and general equilibrium found in the economics literature do not always apply to equilibrium markets. For example, some of the desirable properties, such as Pareto efficiency and coalitional stability, are lost if agents act strategically instead of competitively. Often, the less the agents deviate from competitive behavior, the closer the market outcome of an equilibrium market will be to the competitive general equilibrium. Investigating the agents' motivation for deviating from competitive behavior is one of the main themes of this article.

On the positive side we have that much of the theory in economics focuses directly on the demand functions. For example, Propositions 2.1, 2.2, and 2.3 were based only on properties of the demand functions and do not assume competitive behavior. It follows that they hold for \underline{z} as well as for z . Therefore, the theorems can directly be used for telling whether or not an equilibrium exists and is unique in a given collection of demand functions. The propositions may also be used by auctioneers as requirements on submitted demand functions. For example, in the main Nordic electricity exchange, NordPool, the demand functions submitted for the respective future hours must be continuous (NordPool, 1998), and therefore the equilibrium is guaranteed to exist.⁷

2.4 A method for analyzing the potential gains from speculation

The goal of a self-interested consumer is to find a 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. The other agents may also be speculating (in the same or some other way). That is included in the agent's model of the other agents.

6. Strategic behavior by an agent can also cause the market allocation to be outside of the core, i.e. not coalitionally stable.

7. In NordPool, the demand functions for different hours are sent separately and no iterations are performed after the prices for the respective hours have been established, i.e. $\delta z_{ig}(\mathbf{p})/\delta p_h = 0$, $g \neq h$ is imposed by the mechanism. Due to the high correlation between financial (long term) and spot (short term) market prices, this market works relatively well anyway, even if $\delta z_{ig}(\mathbf{p})/\delta p_h \neq 0$, $g \neq h$ for most agents (Wolak, 1998; Ygge, 1999). Under the assumption that there exists a price so high that there is an excess supply and that there is a price so low that there is an excess supply, the equilibrium hence exists.

Let there be n agents in addition to the speculating agent, s , that we investigate. The excess demand of these n agents for commodity g is

$$z_g^n(\mathbf{p}) = \sum_{i=1}^n z_{ig}(\mathbf{p}) \quad (7)$$

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

The total excess demand with the speculating agent included is

$$z_g(\mathbf{p}) = z_g^n(\mathbf{p}) + z_{sg}(\mathbf{p}) \quad (8)$$

Once the market has reached an equilibrium, excess supply meets demand, i.e. $z_g(\mathbf{p}) = 0$ for every commodity g .⁸ Substituting this into Equation (8) gives

$$z_{sg}(\mathbf{p}) + z_g^n(\mathbf{p}) = 0 \quad (9)$$

2.4.1 CASE A: SPECULATING CONSUMER

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

$$\begin{aligned} \max_{\mathbf{p}} u_s(\mathbf{x}_s(\mathbf{p})) \quad \text{s.t.} & \quad (10) \\ x_{sg}(\mathbf{p}) \geq 0 & \quad (\text{consumer does not produce}) \\ x_{sg}(\mathbf{p}) = e_{sg} - z_g^n(\mathbf{p}) & \quad (\text{supply meets demand}) \\ \mathbf{p} \cdot \mathbf{z}_s(\mathbf{p}) \leq \sum_{h \in \text{producers}} \theta_{sh} \mathbf{p} \cdot \mathbf{y}_h(\mathbf{p}) & \quad (\text{budget constraint}) \end{aligned}$$

This is obtained provided that the equilibrium is unique and the market mechanism finds it (this is discussed further in Section 3).

2.4.2 CASE B: SPECULATING PRODUCER

The goal of a self-interested producer is to find the production vector that maximizes its profits.⁹ 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. A solution to the following maximization problem gives the highest profit that a speculating producer can obtain.

$$\begin{aligned} \max_{\mathbf{p}} \mathbf{p} \cdot \mathbf{y}_s(\mathbf{p}) \quad \text{s.t.} & \quad (11) \\ \mathbf{y}_s(\mathbf{p}) \in Y_s & \quad (\text{feasible production plan}) \\ y_{sg} = z_g^n(\mathbf{p}) & \quad (\text{supply meets demand}) \end{aligned}$$

8. This holds even if agents are strategically—assuming that the 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.

9. This makes the standard assumption that the producer is able to alter its production plan costlessly during the search for equilibrium. If the computational cost of replanning the production is non-negligible, this may not be the case.

This is obtained provided that the equilibrium is unique and the market mechanism finds it. The last equality turns into \geq if free disposal for both inputs and outputs is possible.

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

2.5 Notation

The notation used in this article is summarized in Table 1.

symbol	interpretation
i and g (indices)	These indices denote an agent and a commodity, respectively.
k	The number of commodities in the market.
n	The number of agents in the market.
$\mathbf{p} = [p_1, p_2, \dots, p_k]$	A price vector, where p_g is the price of commodity g .
$u(\mathbf{x})$	An agent's utility for allocation \mathbf{x} .
$\mathbf{e} = [e_1, e_2, \dots, e_k]$	The endowment (initial allocation).
$\mathbf{x} = [x_1, x_2, \dots, x_k]$	The allocation after the trade.
$\mathbf{z}(\mathbf{p}) = [z_1(\mathbf{p}), z_2(\mathbf{p}), \dots, z_k(\mathbf{p})]$	An agent's <i>competitive</i> demand.
$\underline{\mathbf{z}}(\mathbf{p}) = [\underline{z}_1(\mathbf{p}), \underline{z}_2(\mathbf{p}), \dots, \underline{z}_k(\mathbf{p})]$	An agent's <i>revealed</i> demand.

Table 1: *Notation of this article.*

3. Constructing a speculative demand function under perfect information

This section shows what strategies a speculator with perfect information should use so as to drive the market to a solution that maximizes his gain and looks like a (competitive) general equilibrium to the other agents. As defined by Equation (10) or Equation (11), the determination of the speculator's highest possible profit involved a maximization over price vectors \mathbf{p} . However, the fact that the speculator does not directly control the price vector—because prices are affected by others' excess supply and demand decisions as well—makes optimal speculation more difficult. In particular, the speculator is only in control of his revealed excess demand (or supply). Therefore, the speculator would like to choose his $\underline{\mathbf{z}}_s$ so as to drive the market to his desired price vector \mathbf{p}^* . In other words, 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 him. 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 are producers—is to declare an excess demand function with the property

$$\underline{z}_{sg}(\mathbf{p}^*) = -z_g^n(\mathbf{p}^*) \quad (12)$$

for each commodity g , and which has a form such that the particular algorithm for searching for the market equilibrium converges to \mathbf{p}^* .

If the market algorithm were guaranteed to find the equilibrium if one exists, the speculator could simply use a strategy, $\underline{z}_s(\mathbf{p})$, that satisfies

$$\begin{aligned} \underline{z}_{sg}(\mathbf{p}^*) &= -\underline{z}_g^n(\mathbf{p}^*) \text{ and} \\ \underline{z}_{sg}(\mathbf{p}) &\neq -\underline{z}_g^n(\mathbf{p}) \text{ if } \mathbf{p} \neq \mathbf{p}^* \end{aligned} \tag{13}$$

to get his most desired outcome, \mathbf{p}^* . There are several centralized algorithms with such guarantees of finding an (approximate) equilibrium. The first was Scarf's algorithm (Scarf, 1967), and several improvements have been made since, see (Kehoe, 1991; Ellickson, 1993).

However, most distributed equilibrium market algorithms—such as the tâtonnement schemes discussed in Appendix B—use iterative schemes for approaching the equilibrium. Under such schemes some speculative excess demand (or supply) functions $\underline{z}_s(\mathbf{p})$ lead the algorithm to converge to \mathbf{p}^* while others might cause the algorithm to not converge even if $\underline{z}_s(\mathbf{p})$ satisfies Equation (12). Now we will discuss convergence of the market to the speculator's desired solution. In particular, we want to analyze what conditions the speculator's excess demand (or supply) function needs to satisfy so that the market converges. For pedagogic reasons, the topic is covered from more specific settings to more general ones.

3.1 The case of two commodities

Having computed the optimal speculative solution, p_1^* , we would like to describe the strategic behavior leading to this solution under any particular market mechanism used. Before continuing we need to introduce the notation $\frac{\delta f(\mathbf{x})}{\delta x_i}$. For now, let us define it as follows—later in the article (in Theorem 4.2) we will generalize this definition. We write $\frac{\delta f(\mathbf{x})}{\delta x_i} \leq 0$ if $f(\mathbf{x})$ is non-increasing in x_i , and we write $\frac{\delta f(\mathbf{x})}{\delta x_i} \geq 0$ if $f(\mathbf{x})$ is non-decreasing in x_i .

If p_1 is established via an algorithm whose only requirement for finding the equilibrium is $\delta \underline{z}_1(p_1)/\delta p_1 \leq 0$, and we have that $\delta \underline{z}_1^n(p_1)/\delta p_1 \leq 0$ (i.e. excess demand in the market without the speculator is non-increasing with increasing price), we see that if

$$\underline{z}_{s1}(p_1^*) = -\underline{z}_1^n(p_1^*) \text{ and } \frac{\delta \underline{z}_{s1}(p_1)}{\delta p_1} \leq 0, \tag{14}$$

then there is a single solution for $p_1 = p_1^*$, and that solution will be found by the algorithm. An example market algorithm with this convergence requirement is binary search, which is carried out on one of the prices (the other one is fixed, e.g. to 1, without loss of generality). The WALRAS market simulation uses that algorithm (Cheng & Wellman, 1998).

It turns out that simple excess demand revelation strategies exist for the speculator which guarantee that an equilibrium will be reached where the speculator's maximal gain from speculation—that was derived in Section 2.4—materializes. For example, it turns out that two classic bid types from microeconomics satisfy Equation (14): the Cournot quantity bid and the Bertrand price bid (Varian, 1992). The Cournot bid is given by

$$\underline{z}_{s1}(p_1) = -\overline{\underline{z}_1^n(p_1^*)}, \tag{15}$$

where $\overline{z_1^n(p_1^*)}$ is the speculator's (perfect) estimate of $z_1^n(p_1^*)$. The Bertrand bid can be approximated by

$$z_{s1}(p_1) = -\overline{z_1^n(p_1^*)} + \begin{cases} C, & p_1 < \overline{p_1^*} - \epsilon \\ C - (p_1 - \overline{p_1^*} - \epsilon)\frac{C}{\epsilon}, & \overline{p_1^*} - \epsilon \leq p_1 \leq \overline{p_1^*} + \epsilon \\ -C, & p_1 > \overline{p_1^*} + \epsilon, \end{cases} \quad (16)$$

where $\overline{p_1^*}$ is the speculator's perfect estimate of p_1^* , C is a large positive constant, and ϵ is a small positive constant. That is, by selecting a sufficiently large C and a sufficiently small ϵ , the market clearing price will be arbitrarily close to $\overline{p_1^*}$.

3.2 Multiple commodities under gross substitutes

The reasoning of how to drive a market to the desired equilibrium extends easily to a market with more than two commodities. If \mathbf{p} is established via an algorithm whose only requirements for finding the equilibrium are $\delta z_i(\mathbf{p})/\delta p_i < 0$ and $\delta z_i(\mathbf{p})/\delta p_j \geq 0$, $i \neq j$ (e.g. WALRAS or the basic price tâtonnement algorithm which is reviewed in Appendix B), and we have that $\delta z_i^n(\mathbf{p})/\delta p_i < 0$ and $\delta z_i^n(\mathbf{p})/\delta p_j \geq 0$, $i \neq j$, then Equation (14) should be generalized to

$$z_{sg}(\mathbf{p}^*) = -z_g^n(\mathbf{p}^*) \text{ and } \frac{\delta z_{sg}(\mathbf{p})}{\delta p_g} \leq 0 \text{ and } \frac{\delta z_{sg}(\mathbf{p})}{\delta p_h} \geq 0, g \neq h \quad (17)$$

Again, the following simple speculation strategies satisfy these conditions (as in the two commodity case, $g \neq k$):

$$z_{sg}(p_g) = -\overline{z_g^n(p_g^*)}, \quad (18)$$

or

$$z_{sg}(p_g) = \begin{cases} C, & p_g < \overline{p_g^*} - \epsilon \\ C - (p_g - \epsilon)\frac{C}{\epsilon}, & \overline{p_g^*} - \epsilon \leq p_g \leq \overline{p_g^*} + \epsilon \\ -C, & p_g > \overline{p_g^*} + \epsilon. \end{cases} \quad (19)$$

This means that optimal speculation is computationally trivial if the speculator knows the others' excess supply and demand decisions, and if the market would have satisfied the convergence condition (Equation (17)) had the speculator not been present. To construct its optimal excess demand function, the speculator does not need to know the others' complete excess demand functions. The speculator only needs to know the others' aggregate excess supply and demand decisions at the particular price point, \mathbf{p}^* , to which the agent wants to drive the market, see Equation (10) and Equation (11).

3.3 The general case

Many algorithms for finding an equilibrium, for example basic price tâtonnement (Proposition B.1), are guaranteed to find an equilibrium, $\check{\mathbf{p}}$, if

$$\underline{z}(\check{\mathbf{p}}) = \mathbf{0} \text{ and } \check{\mathbf{p}} \cdot \underline{z}(\mathbf{p}) > 0 \text{ for all } \mathbf{p} \text{ not proportional to } \check{\mathbf{p}}. \quad (20)$$

This is a strictly weaker requirement than gross substitutes (Mas-Colell et al., 1995). The proportionality condition collapses to an equality condition when one of the prices is fixed, e.g. $p_k = 1$ as before, which does not lose generality since prices are only relative anyway.

Say that Equation (20) is a sufficient convergence condition for the algorithm in question. Now, the speculator can drive the market to the price point \mathbf{p}^* that maximizes his gain by using a strategy $\underline{z}_s(\mathbf{p})$ that satisfies

$$\underline{z}_s(\mathbf{p}^*) = -\underline{z}^n(\mathbf{p}^*) \text{ and } \mathbf{p}^* \cdot [\underline{z}_s(\mathbf{p}) + \underline{z}^n(\mathbf{p})] > 0 \text{ for all } \mathbf{p} \neq \mathbf{p}^*. \quad (21)$$

Next we derive a simpler (but somewhat more restrictive) condition that suffices to guarantee that Equation (21) is satisfied.

Proposition 3.1 *Assume that the algorithm converges if Equation (20) holds. Assume also that*

$$\underline{z}^n(\mathbf{p}^{n*}) = \mathbf{0} \text{ and } \mathbf{p}^{n*} \cdot \underline{z}^n(\mathbf{p}) > 0 \text{ for all } \mathbf{p} \neq \mathbf{p}^{n*}. \quad (22)$$

In other words, if the speculator would not participate, there would be some market clearing price point, \mathbf{p}^{n} , and the market would reach it. Now, the speculator can drive the market to the price point \mathbf{p}^* that maximizes his gain by using a strategy $\underline{z}_s(\mathbf{p})$ that satisfies*

$$\underline{z}_s(\mathbf{p}^*) = -\underline{z}^n(\mathbf{p}^*) \text{ and } \mathbf{p}^* \cdot \underline{z}_s(\mathbf{p}) > (\mathbf{p}^{n*} - \mathbf{p}^*) \cdot \underline{z}^n(\mathbf{p}) \text{ for all } \mathbf{p} \neq \mathbf{p}^*. \quad (23)$$

Proof. The equality in Equation (21) is trivially satisfied by the equality in Equation (23). What remains to be shown is that the inequality in Equation (23) satisfies the inequality in Equation (21):

$$\begin{aligned} & \mathbf{p}^* \cdot \underline{z}_s(\mathbf{p}) > (\mathbf{p}^{n*} - \mathbf{p}^*) \cdot \underline{z}^n(\mathbf{p}) \\ \Rightarrow & \mathbf{p}^* \cdot \underline{z}_s(\mathbf{p}) > (\mathbf{p}^{n*} - \mathbf{p}^*) \cdot \underline{z}^n(\mathbf{p}) - \mathbf{p}^{n*} \cdot \underline{z}^n(\mathbf{p}) \\ \Leftrightarrow & \mathbf{p}^* \cdot \underline{z}_s(\mathbf{p}) > -\mathbf{p}^* \cdot \underline{z}^n(\mathbf{p}) \\ \Leftrightarrow & \mathbf{p}^* \cdot [\underline{z}_s(\mathbf{p}) + \underline{z}^n(\mathbf{p})] > 0 \end{aligned}$$

This completes the proof. \square

The simplicity of Conditions (21) or (23) makes optimal speculation easy when the speculator knows the aggregate excess demand function of the others, and once the speculator has determined to which price point, \mathbf{p}^* , he wants to drive the market.

However, the extremely simple speculation strategies which were adequate in the case of gross substitutes (Equation (15) and (16)), do not in general satisfy the inequalities in Conditions (21) and (23). This is not surprising: when there is more structure in the others' aggregate excess demand, such as gross substitutes, the speculator can capitalize on that structure by using simpler strategies.

In general, a market can have multiple equilibria, and the speculator would like to make sure that the market converges to one of the ones that maximize his gain, not to one of the other ones. In markets that satisfy gross substitutes with the speculator included, this is not a concern since such markets can have at most one equilibrium (Proposition 2.2). However, in more general settings, the speculator should make sure that his strategy, $\underline{z}_s(\mathbf{p})$, satisfies $\underline{z}_s(\mathbf{p}) = -\underline{z}^n(\mathbf{p})$ only at the price vector(s) to which he wants the market to converge.

4. Constructing a speculative demand function under imperfect information

This section extends the discussion to include the impact of biased beliefs and uncertainty on the speculating agent's strategy. 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 Condition (21) or (23). However, if the speculating agent cannot estimate \underline{z}_1^n perfectly, its outcome will depend on the function chosen. In what follows, we analyze the choice of a demand function under biased beliefs or uncertainty.

4.1 Convergence to a market equilibrium

Independently of how the \underline{z}_{sg} function is chosen, the possible market outcomes can be determined by solving Equation (9). If no solution exists, no algorithm can find it, and if multiple equilibria exist, an analysis that is specific to the market algorithm is required to find out which one will be reached.

The convergence criteria under imperfect information are, of course, the same as under perfect information, i.e. the criteria described in Sections 3.1 – 3.3 are still valid. In the case of monotonically decreasing demand and gross substitution, i.e. Sections 3.1 and 3.2, uncertainty causes no problems in terms of convergence: if the demand of the speculator is monotonically decreasing (and if the commodities are gross substitutes in the multi-commodity case) based on imperfect information about the other agents, they will of course still be monotonically decreasing (and gross substitutes) regardless of the estimation error about the others' demand. Thus, under the assumption that the other agents fulfill the requirements of monotonically decreasing demand and gross substitution and that the speculator fulfills Equation (17), the market equilibrium will be established as in the case of perfect information.

For the general multi-commodity case of Section 3.3, however, things are more complicated. Imperfect information can cause the speculator to create a demand that leads to nonconvergence of the market. Specifically, to guarantee convergence of the market to his desired equilibrium using Condition (21), the speculator would have to choose $\underline{z}_s(\mathbf{p})$ so that $\underline{z}_s(\mathbf{p}^*) = -\underline{z}^n(\mathbf{p}^*)$ and $\mathbf{p}^* \cdot [\underline{z}_s(\mathbf{p}) + \underline{z}^n(\mathbf{p})] > 0$ for all $\mathbf{p} \neq \mathbf{p}^*$, which depends on the others' aggregate demand function, $\underline{z}^n(\mathbf{p})$. The following example shows that if the speculator uses biased information to construct his speculative demand, the market may not converge even if an equilibrium exists under the constructed speculation, and the market would have converged had the speculator not participated in it.

Example. Let $k = 3$, and let $p_3 = 1$. Say that the market algorithm used is basic price tâtonnement. This means that the market searches for market clearing prices using the following price adjustment rule (this is discussed further in Appendix B):

$$\frac{dp_g}{dt} = \underline{z}_g(\mathbf{p}(t)). \quad (24)$$

Say that the other agents' aggregate demand without the speculator is

$$\begin{aligned} \underline{z}_1^n(\mathbf{p}(t)) &= (p_2 - p_2^{n*}) - (p_1 - p_1^{n*}) \\ \underline{z}_2^n(\mathbf{p}(t)) &= -(p_1 - p_1^{n*}) - (p_2 - p_2^{n*}), \end{aligned} \quad (25)$$

where \mathbf{p}^{n*} is the market clearing price vector if the speculator did not participate. Then the price trajectory of the market without the speculator is given by

$$\begin{aligned}\frac{dp_1}{dt} &= (p_2 - p_2^{n*}) - (p_1 - p_1^{n*}) \\ \frac{dp_2}{dt} &= -(p_1 - p_1^{n*}) - (p_2 - p_2^{n*}).\end{aligned}\quad (26)$$

Moving to a coordinate system where $x = p_2 - p_2^{n*}$ and $y = p_1 - p_1^{n*}$ this can be written as

$$\begin{aligned}\frac{dy}{dt} &= x + ay \\ \frac{dx}{dt} &= -y + ax.\end{aligned}\quad (27)$$

where $a = -1$. This pair of differential equations can be solved generally as follows (Simmons, 1972). First, eliminate dt :

$$\frac{dy}{dx} = \frac{x + ay}{ax - y}.\quad (28)$$

This is solved by introducing polar coordinates r and θ defined by $x = r \cos \theta$ and $y = r \sin \theta$. Since

$$r^2 = x^2 + y^2 \quad \text{and} \quad \theta = \tan^{-1} \frac{y}{x},$$

we see that

$$r \frac{dr}{dx} = x + y \frac{dy}{dx} \quad \text{and} \quad r^2 \frac{d\theta}{dx} = x \frac{dy}{dx} - y.$$

Using these equations, Equation (28) can be written as

$$\frac{dr}{d\theta} = ar,$$

so

$$r = ce^{a\theta}.\quad (29)$$

If $a < 0$ this represents a spiral that converges counterclockwise to $x = 0, y = 0$, i.e. $p_2 = p_2^{n*}, p_1 = p_1^{n*}$. If $a > 0$ this represents a spiral that diverges clockwise from that point. If $a = 0$, this represents a circle centered at that point. The constant, c , depends on the starting point, i.e. what initial prices the price tâtonnement algorithm uses.

So, since $a = -1$ in the market without the speculator, the market would converge to $p_1 = p_1^{n*}, p_2 = p_2^{n*}$. Now, the speculator wants to drive the market to an equilibrium with prices \mathbf{p}^* . One way he can do this is by constructing total aggregate demands that lead to a converging spiral according to Equation (29). For example, the speculator can set the total aggregate demands as

$$\begin{aligned}\bar{z}_1(\mathbf{p}(t)) &= (p_2 - p_2^*) - \frac{1}{2}(p_1 - p_1^*) \\ \bar{z}_2(\mathbf{p}(t)) &= -(p_1 - p_1^*) - \frac{1}{2}(p_2 - p_2^*).\end{aligned}\quad (30)$$

This would make the market converge to $p_1 = p_1^*, p_2 = p_2^*$ according to Equation (29) since $a = -\frac{1}{2} < 0$.

However, say that the speculator has the following biased beliefs about the others' aggregate demand:

$$\begin{aligned}\bar{z}_1^n(\mathbf{p}(t)) &= (p_2 - p_2^{n*}) - 2(p_1 - p_1^{n*}) \\ \bar{z}_2^n(\mathbf{p}(t)) &= -(p_1 - p_1^{n*}) - 2(p_2 - p_2^{n*}).\end{aligned}\quad (31)$$

So, the speculator still believes that the market would converge to $p_2 = p_2^{n*}$, $p_1 = p_1^{n*}$ by Equation (29) if he would not participate, but the spiral would be different than the actual one since the believed a is -2 while actually $a = -1$.

The speculator constructs his demand function using the target aggregate demand (Equations (30)), and his beliefs of the others' aggregate demand (Equations (31)) as follows:

$$\begin{aligned}z_{s1}(\mathbf{p}(t)) &= \bar{z}_1(\mathbf{p}(t)) - \bar{z}_1^n(\mathbf{p}(t)) = -p_2^* + p_2^{n*} + 1.5p_1 + 0.5p_1^* - 2p_1^{n*} \\ z_{s2}(\mathbf{p}(t)) &= \bar{z}_2(\mathbf{p}(t)) - \bar{z}_2^n(\mathbf{p}(t)) = p_1^* - p_1^{n*} + 1.5p_2 + 0.5p_2^* - 2p_2^{n*}.\end{aligned}\quad (32)$$

As discussed above, the speculator designed the target demand so that this speculation would make the market converge to \mathbf{p}^* .

However, the actual total demand with the speculator is now

$$\begin{aligned}z_1(\mathbf{p}(t)) &= z_{s1}(\mathbf{p}(t)) + z_1^n(\mathbf{p}(t)) = p_2 - p_2^* + 0.5p_1 + 0.5p_1^* - p_1^{n*} \\ z_2(\mathbf{p}(t)) &= z_{s2}(\mathbf{p}(t)) + z_2^n(\mathbf{p}(t)) = -p_1 + p_1^* + 0.5p_2 + 0.5p_2^* - p_2^{n*}.\end{aligned}\quad (33)$$

This pair of equations can be rewritten as

$$\begin{aligned}z_1(\mathbf{p}(t)) &= (p_2 - A) + \frac{1}{2}(p_1 - B) \\ z_2(\mathbf{p}(t)) &= -(p_1 - B) + \frac{1}{2}(p_2 - A),\end{aligned}\quad (34)$$

where

$$\begin{aligned}B &= \frac{3p_1^* + 4p_2^* + 2p_1^{n*} - 4p_2^{n*}}{5} \\ A &= 2B - 2p_1^* - p_2^* + 2p_2^{n*}.\end{aligned}\quad (35)$$

Therefore, the market will diverge in a spiral since $a = \frac{1}{2} > 0$. An equilibrium exists now at $p_1 = B$, $p_2 = A$, but it is unstable: the basic price tâtonnement algorithm will not find it unless the algorithm happens to be started exactly at that setting of prices.

Figure 2 illustrates the setting numerically when the market would converge to $p_1^{n*} = 3$, $p_2^{n*} = 2$ if the speculator would not participate, and the speculator's desired solution is $p_1^* = 4$, $p_2^* = 5$.

In batch mode market algorithms where each agent submits an entire demand function up front, such divergence problems can easily occur. On the other hand, in iterative algorithms, the agents post only parts of the demand curve to the auctioneer at every iteration, and at the beginning of each iteration, the auctioneer posts new prices for the agents to observe. In the iterative algorithms, the speculator can observe the divergence, which signals to him that his estimate of the others' aggregate demand function is incorrect. Then the speculator can try to gather additional information of the others to get a better estimate, and try to avoid nonconvergence.




Figure 2: *Left: Convergence of the prices to $p_1^* = 4$, $p_2^* = 5$ in basic price tâtonnement as the speculator designed. The curve represents particular starting prices—defined by $\theta = 0, c = 400$ —that are far from \mathbf{P}^* . Right: The speculator’s strategy causes the market to diverge because the speculator had biased information of the others’ aggregate demand function. The curve represents particular starting prices—defined by $\theta = 0, c = 0.001$ —that are close to the new equilibrium $p_1 = B$, $p_2 = A$.*

4.2 Deriving the demand function

Imperfect information about the other agents can be of many forms. Here we will investigate two different forms: the case were the speculator has biased belief (i.e. is not aware that it has imperfect information), and the case were the speculator has probability distributions of the relevant parameters of the other agents.

Now we demonstration of the reasoning via a simple example. Let the market have no producers. In such *pure exchange markets*, the consumers just reallocate their initial endowments among themselves. The set of agents is similar to the one described by Hu & Wellman (1996, 1998). Specifically, we let every agent—except for the speculating one that we investigate—be a competitive agent with constant elasticity of substitution, i.e. an agent having a utility function of the form

$$u_i(\mathbf{x}) = \left(\sum_{g=1}^n \alpha_{ig} x_{ig}^\rho \right)^{\frac{1}{\rho}}, \quad (36)$$

where we have chosen $\alpha_{ig} = 1$ 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 mechanism independent as long as the resources are reallocated after equilibrium (Equation (9)) has been reached.

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

$$u_s(\mathbf{x}) = \sum_{g=1}^n \ln(x_{sg}), \quad (37)$$

and an endowment of 1 for both commodities. Those readers that are relatively unfamiliar with economic oligopoly theory are encouraged to study the derivation of an optimal strategy for this example under perfect information in Appendix A. An understanding of the perfect information case significantly simplifies the understanding of the case of imperfect information, which we will present next.

We now use an error/uncertainty model with two free variables, ϵ_α and ϵ_ρ . The speculating agent's estimate of the utility function of the competitive agents is (cf. Equation (36))¹⁰

$$\bar{u}_i(\mathbf{x}) = \left((1 + \epsilon_\alpha)x_1^{0.5+\epsilon_\rho} + x_2^{0.5+\epsilon_\rho} \right)^{\frac{1}{0.5+\epsilon_\rho}}. \quad (38)$$

4.2.1 BIASED BELIEFS AND NAIVE DEMAND FUNCTIONS

In this section we show that strategies that are optimal in the setting with perfect information, e.g. the classical Cournot and Bertrand models, can be devastating in the setting with biased beliefs, even when the beliefs are only slightly biased.

Note that $\delta z_1(p_1)/\delta p_1 \leq 0$ holds for this example if the Cournot and Bertrand bids are used, independent of the size of the error. Hence, provided that the speculator stops learning about $z_1^n(p_1)$ at some point in time, and hereby fixes $z_{s1}(p_1)$, any market algorithm whose only requirement for finding an equilibrium is that $\delta z_1(p_1)/\delta p_1 \leq 0$, is guaranteed to converge.¹¹ An example of such an algorithm is binary search, which is also used in WALRAS (cf. Appendix B).

The results of using the Cournot bid in our example market are illustrated in Figure 3 and Figure 4. With two competitive agents, the speculator is worse off by speculating than by acting competitively whenever $|\epsilon_\alpha| > 4\%$.

Next, the outcome of Bertrand bids in the presence of biased information is shown in Figures 5 and 6. It is clear that with this type of bid the situation is even worse for the speculator. With two competitive agents, the speculator is worse off by speculating whenever $|\epsilon_\alpha| > 0.15\%$. With a hundred competitive agents, the potential losses from speculation as a result of biased information are drastic already when $|\epsilon_\alpha| > 0.005\%$.

Generally we can say—with some empirical support from the above simulations—that when the speculator's market share is significant, the gain from speculation is relatively large even with some estimation error. But as the speculator's market share decreases, the gain from speculation decreases and the potential losses increase significantly, even if the information is only slightly biased.

4.2.2 FINDING THE OPTIMAL DEMAND FUNCTION IN THE GENERAL CASE

Clearly the demand functions of the previous sections (Cournot and Bertrand bids) are rather naive and they are unreasonable in the presence of imperfect information, cf. (Klemperer & Meyer, 1989). However, determining the optimal demand function for certain

10. If 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.

11. However, the speculator can of course have estimation errors that are so large that the speculator cannot afford to pay for the bids that it makes, i.e. the speculator is limited by its budget constraint. For example, if the agent tries to set a price that is far from the competitive price using a Bertrand bid with a large number of competitive agents, the cost in terms of z_k may be enormous.




Figure 3: *Gains and losses from speculation under biased information with a Cournot bid as a function of ϵ_α and ϵ_ρ for $n = 2$.*




Figure 4: *Gains and losses from speculation under biased information with a Cournot bid as a function of ϵ_α for $\epsilon_\rho = 0$ and $n \in [1, 2, 5, 10, 100]$. The higher the value of n , the lower the possible gain and the greater the loss when beliefs are biased. Hence the top curve corresponds to $n = 1$ and the bottom curve corresponds to $n = 100$.*




Figure 5: *Gains and losses from speculation under biased information with a Bertrand bid as a function of ϵ_α and ϵ_ρ for $n = 2$.*

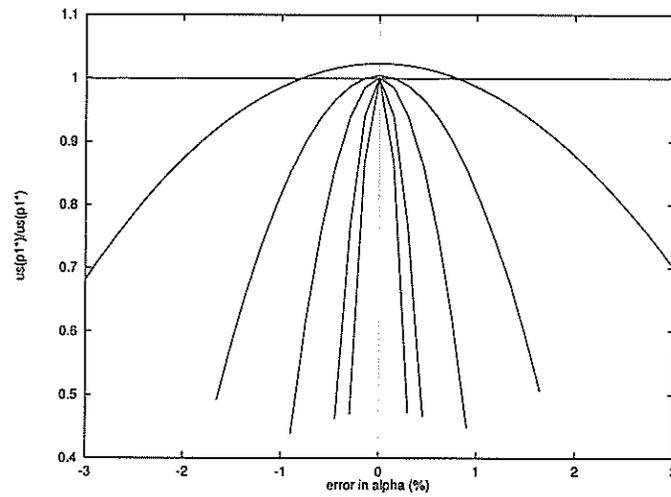


Figure 6: *Gains and losses from speculation under biased information with a Bertrand bid as a function of ϵ_α for $\epsilon_\rho = 0$ and $n \in [1, 2, 5, 10, 20, 30]$. The higher n is, the lower is the possible gain, and the greater is the loss when beliefs are biased. Hence the top curve corresponds to $n = 1$ and the bottom curve corresponds to $n = 30$.*

probability distributions is very hard in the general case. Already with our very simple example, finding the optimal demand function (based on maximization of expected utility) requires the solution of

$$\max_{z_{s1}(p_1)} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \ln \left(1 - n \left(\frac{\alpha_1^{\frac{1}{1-\rho}} (2p_1+1)}{p_1^{\frac{1}{1-\rho}} (\alpha_1^{\frac{1}{1-\rho}} p_1^{\frac{\rho}{\rho-1}} + 1)} - 2 \right) \right) + \ln \left(1 + np_1 \left(\frac{\alpha_1^{\frac{1}{1-\rho}} (2p_1+1)}{p_1^{\frac{1}{1-\rho}} (\alpha_1^{\frac{1}{1-\rho}} p_1^{\frac{\rho}{\rho-1}} + 1)} - 2 \right) \right) f_{\alpha}(\epsilon_{\alpha}) f_{\rho}(\epsilon_{\rho}) d\epsilon_{\alpha} d\epsilon_{\rho} \quad (39)$$

$$\text{such that } z_{s1}(p_1) + z_1^n(p_1) = 0,$$

where $\alpha_1 = 1 + \epsilon_{\alpha}$, $\rho = 0.5 + \epsilon_{\rho}$, $f_{\alpha}(\epsilon_{\alpha})$ is the probability distribution of ϵ_{α} , and $f_{\rho}(\epsilon_{\rho})$ is the corresponding distribution for ϵ_{ρ} .¹²

Solving the above problem can be very complicated, especially with more complicated demand functions of the other agents, a more complicated utility function for the speculator, realistic probability distributions, and larger numbers of commodities (i.e. $k > 2$). In this example we only modeled uncertainty about the utility of the other agents, but assumed perfect knowledge about their behavior, i.e. that they act competitively. This need not be the case in general. We will not discuss the general case further here since it is mainly a matter of different aspects of numerical analysis, but investigate a particular setting in which an optimal demand function can be relatively easily established, and then discuss some more pragmatic approaches to speculation.

4.2.3 CASES WHERE THE OPTIMAL DEMAND FUNCTION IS INDEPENDENT OF THE PROBABILITY DISTRIBUTION

What makes the solution of Equation (39) so complicated is that when determining the optimal $z_{s1}(p_1)$ for each relevant price¹³, p_1 , many different points in the uncertainty space must be considered. To give a very simple example, assume that in the above error model two possible errors, $\langle \epsilon_{\alpha}^1, \epsilon_{\rho}^1 \rangle$ and $\langle \epsilon_{\alpha}^2, \epsilon_{\rho}^2 \rangle$ result in the same optimal price, p_1^* , but that $z_1^n(p_1^*, \epsilon_{\alpha}^1, \epsilon_{\rho}^1) \neq z_1^n(p_1^*, \epsilon_{\alpha}^2, \epsilon_{\rho}^2)$. Then $z_{s1}(p_1)$ must be some type of compromise based on the probabilities of $\langle \epsilon_{\alpha}^1, \epsilon_{\rho}^1 \rangle$ and $\langle \epsilon_{\alpha}^2, \epsilon_{\rho}^2 \rangle$. However, if the setting is such that no "compromises" are required, the problem becomes significantly easier. The following theorem captures this reasoning formally.

Theorem 4.1 *Let a and a' be variable assignments denoting certain values of uncertain parameters of the other agents. If $\mathbf{p}^*(a) = \mathbf{p}^*(a') \Rightarrow \mathbf{z}^n(\mathbf{p}^*(a), a) = \mathbf{z}^n(\mathbf{p}^*(a), a')$, then it is possible to construct an optimal demand function independently of the probability of a vs. the probability of a' .*

Proof. Constructive proof: Let the demand of the speculating agent for each price, \mathbf{p} , be computed as follows. It need not be the case that for every \mathbf{p} , there exists an a such that $\mathbf{p}^*(a) = \mathbf{p}$. For all \mathbf{p} for which no such a exists, $\mathbf{z}_s(\mathbf{p})$ can be set to any value such that

12. $\int_{-\infty}^{\infty} f_{\alpha}(\epsilon_{\alpha}) d\epsilon_{\alpha} = \int_{-\infty}^{\infty} f_{\rho}(\epsilon_{\rho}) d\epsilon_{\rho} = 1$

13. Here, "each relevant price" means every price at which a demand is requested by the auctioneer. This is not limited to the simple one dimensional case discussed in the example.

the market algorithm converges to \mathbf{p} , cf. Figure 8. For every \mathbf{p} for which there exists an a such that $\mathbf{p}^*(a) = \mathbf{p}$, choose $\mathbf{z}_s(\mathbf{p}) = -\mathbf{z}^n(\mathbf{p}, f(\mathbf{p}))$, where $\mathbf{z}^n(\mathbf{p}, f(\mathbf{p}))$ is the demand of the other agents at price \mathbf{p} and variable assignment $f(\mathbf{p})$, and $f(\mathbf{p})$ is the variable assignment for which $\mathbf{p}^*(f) = \mathbf{p}$. From the precondition $\mathbf{p}^*(a) = \mathbf{p}^*(a') \Leftrightarrow \mathbf{z}^n(\mathbf{p}^*(a), a) = \mathbf{z}^n(\mathbf{p}^*(a), a')$ we then have that it is possible to construct an $f(\mathbf{p})$ and a $\mathbf{z}_s(\mathbf{p})$ which is optimal for every possible market clearing price. \square

We now demonstrate how this theorem can be applied to the above example. In our analysis here we limit ourselves to $\epsilon_\alpha > -1$ and $-0.5 < \epsilon_\rho < 0.5$ since going outside these borders would make the estimate of the utility function of the competitive agents (Equation (38)) non quasi-concave and that would have a big, yet rather uninteresting, impact on the analysis. Going back to our example, it turns out that $\epsilon_\alpha = 0$ or $\epsilon_\rho = 0$ fulfills $\mathbf{p}^*(a) = \mathbf{p}^*(a') \Leftrightarrow \mathbf{z}^n(\mathbf{p}^*(a), a) = \mathbf{z}^n(\mathbf{p}^*(a), a')$. (See Figure 7 for the case where $\epsilon_\rho = 0$). Therefore, if the speculator knows either $\{\epsilon_\alpha = 0, \epsilon_\rho \in (-0.5, 0.5)\}$ or $\{\epsilon_\alpha \in (-1, \infty), \epsilon_\rho = 0\}$, it can use a very simple algorithm (Algorithm 1) to determine $\mathbf{z}_{s1}(p_1)$.¹⁴




Figure 7: The optimal demand when $\epsilon_\rho = 0$ is plotted together with $-\mathbf{z}^n(p, \epsilon_\alpha)$ for a number of different ϵ_α values.

Despite the somewhat naive nature of Algorithm 1 (including binary search in two dimensions), it is relatively efficient and the computation of \mathbf{z}_{s1} for a specific price is a matter of a few milliseconds with a regular 300MHz uniprocessor computer. If required, the computations can be sped up significantly by the use of a Newton-Raphson method (see e.g. (Press, Teukolsky, Vetterling, & Flannery, 1994)) as that would not only decrease the number of iterations of each search, but also easily take advantage of previous estimations of p_1^* when entering the second repeat-loop of Algorithm 1. We ran Algorithm 1 on a large number of different prices for $\epsilon_\alpha = 0$ and for $\epsilon_\rho = 0$. The optimal curves of these two cases are plotted in Figure 8.

14. See also the C++ code available on-line (Section 6).

//Let ϵ represent ϵ_α or ϵ_ρ depending on which one is set to zero.

ϵ_{max} = initial max guess; ϵ_{min} = initial min guess

Repeat

$$\epsilon = \frac{\epsilon_{max} - \epsilon_{min}}{2}$$

p_1^{max} = initial max guess; p_1^{min} = initial min guess

Repeat

$$p_1^* = \frac{p_1^{max} - p_1^{min}}{2}$$

if $\frac{\partial u_s(p_1^*)}{\partial p_1}$ is sufficiently above zero¹⁵

$$p_1^{min} = p_1^*$$

if $\frac{\partial u_s(p_1^*)}{\partial p_1}$ is sufficiently below zero

$$p_1^{max} = p_1^*$$

Until $\frac{\partial u_s(p_1^*)}{\partial p_1}$ is sufficiently close to zero

if $p_1^* - p_1$ is sufficiently above zero

$$\epsilon_{max} = \epsilon$$

if $p_1^* - p_1$ is sufficiently below zero

$$\epsilon_{min} = \epsilon$$

Until p_1^* is sufficiently close to p_1

$$z_{s1}(p_1) = -z_1^n(p_1, \epsilon)$$

Algorithm 1: FINDING THE OPTIMAL DEMAND, z_{s1} , AT A SPECIFIC PRICE p_1 WHEN $\epsilon_\alpha = 0$ OR $\epsilon_\rho = 0$. NOTE THAT THE PARTIAL DERIVATIVE ($\frac{\partial u_s(p_1^*)}{\partial p_1}$) NEED NOT EXIST: A NUMERICAL APPROXIMATION OF THE CHANGE IN UTILITY PER CHANGE IN PRICE AND A CONTINUOUS u IS SUFFICIENT.

In the example above, there was uncertainty only about the utility functions of the other agents (cf. Equation (38)). However, a and a' in Theorem 4.1 can just as well represent uncertainty about the *strategies* of the other agents, as will be shown in the following example. In this example we investigate speculation by a producer in a simple electricity market with only two commodities ($k = 2$); electricity of a specific time period and money. The producer has a production cost of $\frac{1}{10}(-z_1)^2$ (in terms of the money z_2 ,¹⁶ where z is a consumed amount of the power. The producer has the following beliefs about the market:

- There are five identical consumers and one producer in the market.
- The utility of each consumer is $u(x) = 100 - (x_1 - 10)^2 + x_2$, and the endowment of each consumer is zero.

16. The actual cost and utility functions of this example have been chosen for their simplicity. The main important and realistic feature is that marginal cost increases in the region under investigation, and that marginal utility decreases in the same region.




Figure 8: *The optimal demand curves for $\epsilon_\alpha = 0$ and for $\epsilon_\rho = 0$ (error in ρ and error in α respectively). The reason that the demand with error in ρ covers relatively few points is that in the error range investigated ($-0.5 < \epsilon_\rho < 0.5$), p_1^* can only be within the plotted area. The demand outside this region can be set to any value such that the market algorithm converges.*

- The consumers plan to collude. They will bid so as to optimize against the producer's competitive behavior. In their speculative demand, the consumers will base their bid on the $\gamma - \beta$ approach (described below in the Section 4.2.5), with $\beta = 0$.
- With a certain probability, each consumer will deviate from the collusive agreement and instead bid its competitive demand, in order obtain a higher utility than if keeping the agreement.

From the above data we would like to construct the producer's optimal demand (supply) function. As we will see, we can rely on Theorem 4.1 and construct an optimal demand function independent of the probability that a consumer will deviate from the collusive agreement. By using the method of Section 2, we find that the optimal outcome for the coalition is a price of 5 (compared to the competitive equilibrium which yields a price of 6.67). By using the $\gamma - \beta$ parameter method of Section 4.2.5, each consumer's revealed demand for obtaining this price is $z(p) = \frac{2}{3}(10 - \frac{p}{2})$ (the competitive demand would be $z(p) = 10 - \frac{p}{2}$). Thus, the demand facing the producer is: $z(p) = j \left(\frac{2}{3}(10 - \frac{p}{2}) \right) + (5 - j)(10 - \frac{p}{2})$, where j can have any integer value from zero to five—corresponding to the number of consumers that deviate from the collusive agreement. As can be seen directly from Figure 9, Theorem 4.1 is applicable.

The optimal price and the corresponding supply for the different values of j are (the syntax of the data is $\langle j, \text{price}, \text{demand} \rangle$): $\langle 0, 17.14, 7.14 \rangle$, $\langle 1, 17, 7 \rangle$, $\langle 2, 16.84, 6.84 \rangle$, $\langle 3, 16.67, 6.67 \rangle$, $\langle 4, 16.47, 6.47 \rangle$, and $\langle 5, 16.25, 6.25 \rangle$. These points (capturing all possible values of j) define the optimal demand curve for the producer. This optimal demand is




Figure 9: *The optimal demand (supply) function for the producer in a case of uncertainty about strategies of the consumers in the electricity trade example of Section 4.2.3.*

plotted in Figure 9. Hence, irrespective of whether the consumers decide to break the collusion or not, the producer will always have an optimal response to their choice.

4.2.4 CASES WHERE THE OPTIMAL DEMAND FUNCTION IS DECREASING WITH PRICE

Theorem 4.1 described the conditions under which a demand function can be constructed independently of the probability distribution of the error. However, in some markets there may be restrictions on the shape of the demand functions that an agent is allowed to submit to the auctioneer (such as in the NordPool electricity market, as was discussed earlier in the article). This means that there might be cases where it is possible to construct a demand function that is optimal independent of the probability distribution of the error, but this function cannot be used due to the restrictions. As described in previous sections, one typical restriction is that demand must be non-increasing in price. We now give sufficient conditions for the optimal demand function to be decreasing in price.

Theorem 4.2 *If*

1. u_s is quasi-concave¹⁷, and
2. $\frac{\delta z_1^n(p)}{\delta p_1} < 0$ for all possible variable assignments, and
3. $p_1^c(a_1) < p_1^c(a_2)$ (this is just a matter of naming), and

17. Having a quasi-concave utility function is equivalent to having convex preferences. An agent's preferences are convex if all iso-utility curves (i.e. the set of different allocations having the same utility) are convex, see e.g. (Takayama, 1985, pp. 181–182).

$$4. \text{sign}(\underline{z}_1^n(p_1^c(a_1), a_1))^{18} = \text{sign}(\underline{z}_1^n(p_1^c(a_2), a_2)) \Rightarrow -\frac{\delta \underline{z}_1^n(p_1, a_1)}{\delta p_1} \geq -\frac{\delta \underline{z}_1^n(p_1, a_2)}{\delta p_1},$$

$$\min(p_1^*(a_1), p_1^*(a_2)) \leq p_1 \leq p_1^c(a_2),^{19}$$

then $\underline{z}_{s1}^*(p_1^*(a_1), a_1) > \underline{z}_{s1}^*(p_1^*(a_2), a_2)$ and $p_1^*(a_1) < p_1^*(a_2)$.

Proof. From preconditions (1) and (2) we have $\underline{z}_{s1}(p_1^c, a_i) = 0 \Rightarrow \underline{z}_{s1}(p_1^*, a_i) = 0$; $\underline{z}_{s1}(p_1^c, a_i) < 0 \Rightarrow \underline{z}_{s1}(p_1^*, a_i) < 0$; and $\underline{z}_{s1}(p_1^c, a_i) > 0 \Rightarrow \underline{z}_{s1}(p_1^*, a_i) > 0$. Hence, all cases except $\underline{z}_1^n(p_1^c(a_1), a_1) > 0 \wedge \underline{z}_1^n(p_1^c(a_2), a_2) > 0$ and $\underline{z}_1^n(p_1^c(a_1), a_1) < 0 \wedge \underline{z}_1^n(p_1^c(a_2), a_2) < 0$ are trivial. We show the proof for the latter case. The proof for the former case is analogous.

From preconditions (1), (2), and (3) we have $\text{MRS}(p_1^c(a_1), a_1) < \text{MRS}(p_1^c(a_2), a_2)$.²⁰ For this case, the gain from speculation is obtained by lowering \underline{z}_{s1} . (In other words, monopoly (and few customer market) prices are lower than competitive prices, because of lower revealed demand.) Hence, as can be seen from the MRS discussion above, $p_1^*(a_1) > p_1^*(a_2)$ requires that the effect on prices by lowering \underline{z}_{s1} must be greater with a_2 than with a_1 in the interval $\min(p_1^*(a_1), p_1^*(a_2)) \leq p \leq p_1^c(a_2)$. From precondition (4) we have that this cannot be the case. Now (2), (3) and (4) directly imply $\underline{z}_{s1}^*(p_1^*(a_1)) > \underline{z}_{s1}^*(p_1^*(a_2))$. \square

4.2.5 PRAGMATICS OF CHOOSING A DEMAND FUNCTION

We believe that the theory presented above is useful for analyzing the possible gains from speculation in equilibrium markets, and for constructing advanced computational agents for participating in such markets. The examples above show that Theorem 4.1 and Theorem 4.2 can be useful. However, they do not apply in all cases. In this section we discuss some pragmatic alternative methods for constructing speculative demand functions in cases where they do not apply. It seems reasonable to assume that in many practical markets, an agent has some—possibly rough—estimate of the aggregate properties of the other agents, and some feeling for how much risk it is willing to take in speculating. In such settings it is not obvious how to construct a reasonable demand function in regions that are far from the expected market price. In this section we introduce a basic demand function that can be tailored to different degrees of speculation and risk. We believe that such a demand could be very useful in practice and that it could serve as a basis for an agent that learns about its competitors during the search for equilibrium (cf. (Wellman & Hu, 1998)). It makes sense that the more information the learner obtains about its environment, the more it can speculate and the larger risk (of loosing if it is really wrong) it can take in order to gain more. This is easily captured by adjusting parameters in the proposed demand function.

$$18. \text{sign}(x) = \begin{cases} -1, & x < 0 \\ 0, & x = 0 \\ 1, & x > 0 \end{cases}$$

$$19. \text{The interpretation of } \frac{\delta f(x)}{\delta x_i} \geq \frac{\delta g(x)}{\delta x_i}, x_j^{\min} \leq x_j \leq x_j^{\max} \text{ is } f(x_1, \dots, x_i + \epsilon, \dots, x_k) - f(x_1, \dots, x_i, \dots, x_k) \geq g(x_1, \dots, x_i + \epsilon, \dots, x_k) - g(x_1, \dots, x_i, \dots, x_k), x_j^{\min} \leq x_j \leq x_j^{\max}, x_j^{\min} \leq x_j + \epsilon \leq x_j^{\max}, \epsilon > 0.$$

$$20. \text{MRS is the marginal rate of substitution defined as } \frac{\partial u_s(\mathbf{z})}{\partial u_s(\mathbf{z})} \text{ (Mas-Colell et al., 1995). (As before, the}$$

partial derivative need not exist. A numerical approximation suffices.) Furthermore, $\text{MRS}(p_1^c(a_1), a_1)$ is the marginal rate of substitution at $\mathbf{z} = (-\underline{z}_1^n(p_1^c, a_1), p_1, \underline{z}_1^n(p_1^c, a_1))$.

Further investigations are, however, required to produce solid arguments for the usefulness of this demand in conjunction with learning.

The proposed demand function is based on the competitive demand—which is a reasonable choice if the agent has very limited information about its competitors. The *risk factor*, β , is used as follows:

$$\gamma_g = 1 - (1 - \gamma_g^0)e^{-\beta(\bar{p}_g^* - p_g)^2}, \quad (40)$$

where γ_g^0 is a *speculation factor*. The definition of the risk factor, β , could also be extended to one risk factor per commodity, β_g , in Equation (40). The revealed demand, $\underline{z}_s(\mathbf{p})$, is given by

$$\underline{z}_{sg}(\mathbf{p}) = \gamma_g z_{sg}(\mathbf{p}), \quad (41)$$

where $\mathbf{z}_s(\mathbf{p})$ again is the demand resulting from competitive behavior of the agent that we are observing. Hence with $\gamma_g^0 = 1$ or $\beta = \infty$, competitive behavior is obtained.

The interpretation of the above is that the speculation is reduced if the price is far from the expected optimal price, \bar{p}^* , and approaches γ_g^0 when the price approaches \bar{p}^* . Reasonably, $\gamma_g^0 \in [\min(\frac{-\underline{z}_g^n(\mathbf{p}^*)}{z_{sg}(\mathbf{p}^*)}, 1), \max(\frac{-\underline{z}_g^n(\mathbf{p}^*)}{z_{sg}(\mathbf{p}^*)}, 1)]$, and $\beta \geq 0$. That is, for every price, the revealed demand should be between the competitive demand and the (estimated) optimal demand, and the distance between the competitive demand and the revealed demand should decrease with increased distance between the actual market clearing price and the estimated market clearing price. We now demonstrate the concept on our example. If $\epsilon_\alpha \neq 0$ and $\epsilon_\rho \neq 0$, it is impossible to determine an optimal demand independently of the probability distribution of the error. First the demand function is plotted for $\gamma_g^0 = \frac{-\underline{z}_g^n(\mathbf{p}^*)}{z_{sg}(\mathbf{p}^*)}$ and a number of different values of β , Figure 10. In other words, if the estimate of \underline{z}^n coincides with the actual value at the market price, the maximal gain from speculation is obtained, and different risks for losses caused by biased beliefs are modeled by β . Note that requirements on the demand function might hinder the speculator from independently choosing γ_g^0 and β . As seen from Figure 10, too high a β causes the demand to increase with price in a certain region, and this might cause some algorithms not to converge to the desired outcome.

The result of using a demand such as the one in Figure 10 is shown in Figures 11 and 12. In both cases $\gamma_g^0 = \frac{-\underline{z}_g^n(\mathbf{p}^*)}{z_{sg}(\mathbf{p}^*)}$, and $n = 2$. In Figure 11 the agent uses $\beta = 0$. This enables the agent to gain from speculation also when there are moderate errors. If errors are larger, however, the speculator loses in utility compared to acting competitively.

In Figure 12 the agent uses $\beta = 25$. In this case it still obtains the maximal gain from speculation if there is zero error, but with increasing error the utility rapidly approaches the utility resulting from competitive behavior. Thus, the area in which a gain is obtained is reduced, but, for this example, there is no risk of a significant loss, independent of the size of the error $\langle \epsilon_\alpha, \epsilon_\rho \rangle$.

For this example it is relatively easy to construct demand functions that result in a gain compared to competitive behavior, even for significant uncertainty. The benefits of this new parameterized demand function compared to traditional oligopoly demands from economics become apparent when comparing Figures 3–6 to Figures 11 and 12. For example, with the Bertrand bid and with $n = 2$ the error for which there is no longer a gain from speculation is approximately $|\epsilon_\alpha| = 0.15\%$. The corresponding error with our new pragmatic demand function is approximately $-50\% \leq \epsilon_\alpha \leq 500\%$, with $\beta = 0$. Thus, compared to




Figure 10: A pragmatic approach to choosing a demand function. $\gamma_g^0 = \frac{-\bar{z}_g^n(p^*)}{z_s(p^*)}$ and $\beta = 0, 25, 100, 400, 1500$ or 5000. A small β means a larger segment where the speculative demand deviates from the competitive demand, i.e. greater risk, but also a greater chance of profit when errors are moderate, cf. Figure 11 and Figure 12. The top curve is the competitive demand and hence the other curves are ordered from bottom to top by increasing β .




Figure 11: The result of the demand of Figure 10 with $\gamma_g^0 = \frac{-\bar{z}_g^n(p^*)}{z_s(p^*)}$, $\beta = 0$, and $n = 2$.

existing methods, this new parameterized demand provides a relatively simple way to obtain significant gains from speculation at relatively low risk.




Figure 12: The result of the demand of Figure 10 with $\gamma_g^0 = \frac{-\bar{z}_g^n(p^*)}{z_s(p^*)}$, $\beta = 25$, and $n = 2$.

5. Strategic behavior by multiple agents

So far in this article we discussed a setting where one agent (or a coalition of agents) is generating a best-response strategy to some fixed strategies of the others. We did not assume that the others act competitively as price-takers (except in the example), but we did assume that the others' strategies are fixed, i.e. they do not depend on what strategy the speculator chooses. However, it might be that the speculator cannot treat others' strategies as fixed because the others would like to tailor their strategies to the specific strategy that the speculator chooses. If the others are rational, they would like to generate a best-response like that.

The analysis of the case where the choice of one agent affects the choices of the other agents cannot be performed as generally as the analysis in the previous sections of this paper. Unless the bids are given in batch mode where each agent submits its entire bid at once, the analysis will unavoidably reflect what type of information is sent to agents at different times, to what extent agents commit to previous bids, etc. Analyzing the construction of demand functions when the demand functions of the other agents cannot be assumed to be fixed is beyond the scope of this article. Instead we will merely describe some of the most important tools from economics and game theory for performing such mechanism-dependent analysis.

The agents' strategies are said to be in *Nash equilibrium* (Nash, 1950; Mas-Colell et al., 1995) if each agent's strategy is his best response to the others' strategies.²¹ This can be viewed as a necessary condition for robustness against manipulation in settings where all agents can act strategically. In mechanisms with multiple steps one can also strengthen the

21. Stated in another way, 1. each agent optimizes his strategy to his beliefs of the others' strategies, and 2. his beliefs about the others' strategies are correct. A relaxation of this, called *self-confirming equilibrium*, maintains 1., but relaxes 2. by requiring that the agent's beliefs about the others' strategies are correct only on the path of play (Fudenberg & Levine, 1993). The question of learning to play according to a Nash equilibrium of a game under imperfect information, when the same game is repeated infinitely many times, is studied further for example in (Kalai & Lehrer, 1993).

Nash equilibrium solution concept in different ways by requiring that the strategies stay in equilibrium at every step of the game. For example, the *perfect Bayesian equilibrium* requires that the agents' strategies and beliefs are in equilibrium at every point of the game given that agents update their beliefs using Bayes rule (Mas-Colell et al., 1995; Kreps, 1990).

Unlike our analysis, the Nash equilibrium outcome is specific to the market mechanism. Important factors impacting the Nash equilibrium—and therefore also the outcome—are the order in which bids are submitted (see e.g. Stackleberg vs. Cournot models (Mas-Colell et al., 1995)), whether the bids are sealed or open (Sandholm, 1996a), whether the mechanism is iterative (the agents can change their excess demand between iterations) or not, whether trades occur during the market algorithm (Sandholm, 1993; Sandholm & Lesser, 1995; Sandholm, 1996b) or only after a general equilibrium has been reached, whether the agents can decommit from their agreements by paying a penalty (Sandholm & Lesser, 1996; Sandholm, Sikka, & Norden, 1999; Sandholm & Zhou, 1999), *etc.*

In some markets, a Nash equilibrium might not exist or it might not be unique. Also, 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. Furthermore, even if both the Nash equilibrium and the general equilibrium exist, they do not coincide in general. This can be shown with a simple example that has only two agents and one commodity (Ellickson, 1993).

A considerable amount of research has gone into constructing mechanisms that support a general equilibrium outcome in Nash equilibrium or in its refinements. In exchange economies with complete information—i.e. where every agent knows all the parameters of the economy—this can be achieved relatively easily (Moore, 1992). However, in the more realistic settings where agents have private information, the mechanism designer has a much harder task even in economies with no producers. The fact that agents have incentives to reveal speculative excess demand functions in the straightforward direct revelation mechanism was observed early on (Hurwicz, 1972). More recently it has been shown that truth-telling in an exchange economy can in general be obtained in strategy-proof equilibrium only by a mechanism that fixes the price ratios at which commodities can be exchanged (Barbera & Jackson, 1995). The mechanism designer has to fix the ratios in advance without knowing the agents' types. It follows that the outcome is not Pareto efficient in general because some efficiency improving trades cannot occur at those price ratios. Furthermore, the inefficiency does not disappear as the number of agents increases.

The mechanism designer's task becomes easier if she (and the agents playing the game) know the agents' priors, and the agents are Bayesians. The allocations that can be implemented in Bayes-Nash equilibrium in such settings have been characterized by Jackson (1991).

Another approach to the mechanism design problem is *virtual implementation*: requiring that the mechanism only guarantees that the desired outcome is achieved in equilibrium with probability $1 - \epsilon$, and with probability ϵ some other outcome will occur. Virtual implementation via iterated elimination of dominated strategies was studied in the incomplete information setting in (Abreu & Matsushima, 1991). Recently it was shown that almost any outcome of the incomplete information game can be virtually implemented in Bayes-Nash equilibrium (Duggan, 1997).

We argue that for each mechanism proposed for implementation of equilibrium markets including self-interested computational agents, a thorough game theoretic analysis should be attempted. However, as discussed above, game theoretic solution concepts sometimes run into non-existence and non-uniqueness problems. In addition, some mechanisms are difficult to analyze game theoretically. For example, in WALRAS, the agents might change their excess 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 learn more about other agents' excess demand/supply functions. If many agents are involved in such probing, time can become an important factor. Some agents might reveal progressively more of their competitive excess 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.²²

Once the other agents' strategies are known (e.g. from a game theoretic equilibrium analysis), the methods of this article can be used to analyze the speculating agent's strategy alternatives. The methods are mechanism 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.

In this article we studied how a speculating agent can maximize its profit given the demand functions of the other agents. When an agent has the possibility to collude with others—i.e. to coordinate the speculative demand function revelations with others—it is essential to be able to evaluate the possible gains from collusion, and to determine joint strategies that drive the market to the desired equilibrium. This can be done with *exactly* the same methods as for a single speculating agent, i.e. the methods presented in this article apply directly. The sole difference is that the object of maximization is the sum of the profits of the colluding agents instead of that of any single agent. The same holds for computing the maximal gain that can be obtained by deviating from a collusive arrangement.

The methods of this article 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 best-response strategies—the theory of speculation under biased beliefs is highly applicable.

6. Conclusions

In computational markets utilizing algorithms that establish a general equilibrium, competitive behavior has usually been 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.

22. Some work has addressed non-competitive behavior in WALRAS (Hu & Wellman, 1996; Wellman & Hu, 1998), 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 (instead of the more deliberative speculations presented in this article). Further analysis is required to determine whether the speculator's optimal strategy can be captured in that simple model. This need not be the case because the optimal strategy may involve some more "aggressive" behavior, e.g. the probing described above.

We presented a method for computing the maximal advantage of speculative strategic behavior in equilibrium markets. It is computed from the other agents' excess 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 not specific to a particular market mechanism, but applies to most existing market mechanisms and ones to come. Specifically, it applies to all market mechanisms where the exchanges are carried out after an equilibrium is reached. More importantly, we also constructed excess demand revelation strategies that guarantee that an agent can drive the market to an equilibrium where the agent's maximal advantage from speculation materializes. In other words, the additional constraint that the market mechanism will converge is satisfied by the speculative demand revelation.

In the presence of imperfect information about the others' aggregate excess demand, constructing an optimal demand function is rather delicate. We demonstrated how classical oligopoly demands (the Cournot quantity bid and the Bertrand price bid) may lead to significant losses even when an agent's beliefs are only slightly biased. Furthermore, we showed how the optimal demand is computed from probability distributions of the behavior of the other agents, and we discussed the difficulties associated with that approach. We also presented conditions when an optimal demand function can be constructed regardless of the probability distribution of the estimation error. It was shown how to compute the optimal demand curve in such a setting and an algorithm for doing that was introduced and exemplified. Some pragmatics of choosing a demand function in the case of imperfect information were given, and we showed that it can be relatively easy to construct demand functions that result in a gain from speculation even when estimation errors are rather large. Finally, we discussed the mechanism dependent game theoretic issues related to multiple agents counterspeculating. As discussed, recent advances in game theory have a lot to say about that setting.

We believe that computational agents representing self-interested real world parties will deviate from competitive behavior in practice if they can benefit from doing so. Almost all previous work on computational equilibrium markets has assumed that agents act as price-takers, although that is not what a rational agent would do. We hope that this article can serve as a blueprint for building agents that act efficiently on behalf of the real-world parties that they represent. The methods of this article enable this by showing how an agent with imperfect information about others should construct demand functions that strike a tradeoff between the potential gains from speculation and the risk that is associated with it, while considering the convergence of the market mechanism.

Source code

For the sake of reproducibility, the `C++` source code of the program used for obtaining the numerical values shown in Table 2, Figure 13, and Figures 3–12 is downloadable through the world-wide web from www.ensearch.se/ygge

Acknowledgments

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

Appendix A. The possible gain from speculation in a specific example with perfect information

In this appendix we derive the maximal gain that is obtainable via speculation under perfect information about the others in the example market of Section 4.2. We get, e.g. from (Wellman, 1994), and the definition of excess demand, that the excess demand of the competitive agents is

$$\underline{z}_1^n(\mathbf{p}) = n \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right). \quad (42)$$

Because prices are only relative, we can set one of the prices arbitrarily, e.g. p_k can be set to 1, i.e. $p_2 = 1$. From the budget constraint ($\mathbf{p} \cdot \underline{z}(\mathbf{p}) = 0$) we then get $\underline{z}_k(\mathbf{p}) = -\sum_{g=1}^{k-1} p_g \underline{z}_g(\mathbf{p})$, i.e. $\underline{z}_2^n(p_1) = -p_1 \underline{z}_1^n(p_1)$. Using this with Equation (37) and Equation (42), we get

$$u_s(x(p_1)) = \ln \left(1 - n \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right) \right) + \ln \left(1 + n \left(\frac{2p_1 + 1}{p_1 + 1} - 2p_1 \right) \right). \quad (43)$$

From Equation (37) we see that x_{sg} 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 Equation (43), and the requirement that $x_{s1} > 0$, we have

$$p_1 > p_1^{min} = -\frac{1}{2(2n + 1)} + \sqrt{\frac{n}{2n + 1} + \left(\frac{1}{2(2n + 1)} \right)^2}. \quad (44)$$

Analogous reasoning for x_{s2} shows that

$$p_1 < p_1^{max} = \frac{1}{4n} + \sqrt{\frac{1}{2} \left(1 - \frac{1}{n} \right) + \left(\frac{1}{4n} \right)^2}. \quad (45)$$

If n 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 , Equation (43), with respect to p_1 is

$$\begin{aligned} \frac{\partial u_s}{\partial p_1} &= \frac{-1}{1 - n \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right)} n \left(2 \frac{1}{p_1(p_1 + 1)} - \frac{2p_1 + 1}{p_1^2(p_1 + 1)} - \frac{2p_1 + 1}{p_1(p_1 + 1)^2} \right) + \\ &\quad \frac{1}{1 - n \left(\frac{2p_1 + 1}{p_1 + 1} - 2p_1 \right)} n \left(\frac{2}{p + 1} - \frac{2p_1 + 1}{(p_1 + 1)^2} - 2 \right). \end{aligned} \quad (46)$$

It turns out that $\lim_{p_1 \rightarrow p_1^{min+}} \frac{\partial u_s}{\partial p_1} > 0$ and $\lim_{p_1 \rightarrow p_1^{max-}} \frac{\partial u_s}{\partial p_1} < 0$, that $\frac{\partial u_s}{\partial p_1}$ is continuous 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 $\underline{z}_{s1} = \frac{p_1 + 1}{2p_1} - 1$. Setting the aggregate excess demand to zero gives $n \left(\frac{2p_1 + 1}{p_1(p_1 + 1)} - 2 \right) +$

$\frac{p_1+1}{2p_1} - 1 = 0$. Solving for the competitive price gives $p_1^c = \sqrt{\frac{2n+1}{4n+1}}$. The results are shown in Table 2.²³

n	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	0.01746	0.01626
2	0.5403	1.000	0.7402	0.7454	0.02202	0.02152
5	0.6303	0.8262	0.7227	0.7237	0.02614	0.02602
10	0.6667	0.7670	0.7154	0.7157	0.02788	0.02784
20	0.6863	0.7372	0.7114	0.7115	0.02884	0.02884
30	0.6931	0.7272	0.7100	0.7160	0.02918	0.02918
100	0.7029	0.7131	0.7080	0.7080	0.02967	0.02967

Table 2: *Acting strategically vs. acting competitively.* n 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 13, the utility is plotted for the situations where the agent acts strategically and where it acts competitively. As expected (see e.g. (Roberts & Postlewaite, 1976)), 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 agents is around five, the gain from strategic behavior is negligible. The two important conclusions from this exercise are: 1) there is always a positive gain, and 2) the gain from speculation often decreases rapidly with the number of agents.

23. Since utility is invariant to positive affine transformations (i.e. multiplying by a positive constant and adding a constant) (Mas-Colell et al., 1995) one should be careful when discussing degrees of improvement.




Figure 13: Comparison of strategic and competitive utility, $u_s(p_1^*)$ and $u_s(p_1^c)$. 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.

Appendix B. Algorithms for finding a market equilibrium

The operational motivation behind market mechanisms is that the agents can find an efficient joint solution—which takes into account tradeoffs between agents and the fact that the values of different commodities to a single agent may be interdependent—while never centralizing all the information or control. There are many algorithms that can be used to search for an equilibrium, with different levels of decentralization. The most common algorithm for this purpose is the *basic price tâtonnement process* (Algorithm 2) which is a steepest descent search method. It is an iterative mechanism, and the trades, production, and consumption occur only after the process has converged. 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.

Clearly, if no equilibrium exists, no algorithm can find it. Furthermore, sometimes the price tâtonnement algorithm fails to find an equilibrium even if equilibria exist. However, there are sufficient conditions that guarantee that an equilibrium is found if it exists. One such sufficient condition is the gross substitutes property which was used in Proposition 2.2. More generally,

Proposition B.1 (Convergence) *The basic price tâtonnement algorithm converges to an equilibrium if $\mathbf{p}^* \sum_i \mathbf{z}_i(\mathbf{p}) > 0$ for all \mathbf{p} not proportional to a market clearing price vector \mathbf{p}^* (Mas-Colell et al., 1995).²⁴*

²⁴. If one price is fixed (which can always be done without loss of generality) the proportionality condition reduces to an equality.

Algorithm for the price adjustor:
 $p_g = 1$ for all $g \in [1..k]$
 Set λ_g to a positive number for all $g \in [1..k - 1]$
 (How to select λ_g is described separately.)
 Repeat
 Broadcast \mathbf{p} to all agents
 Receive a net demand vector \mathbf{z}_i from each agent i
 For $g = 1$ to $k - 1$
 $p_g = p_g + \lambda_g \sum_i z_{ig}$
 Until $|\sum_i z_{ig}| < \epsilon$ for all $g \in [1..k - 1]$
 Inform all agents that an equilibrium has been reached

Algorithm for agent i :
 Repeat
 Receive \mathbf{p} from the adjustor
 Announce to the adjustor a demand vector $\mathbf{z}_i \in \mathfrak{R}_+^{k-1}$
 representing the agents revealed desires
 Until informed that an equilibrium has been reached
 Exchange and consume or produce

Algorithm 2: DISTRIBUTED BASIC PRICE TÂTONNEMENT ALGORITHM.

Strictly speaking, the convergence guarantee only applies to the continuous variant where prices are adjusted according to

$$\frac{dp_g}{dt} = \lambda_g \sum_i z_{ig}(\mathbf{p}),$$

not to the more realistic discrete step version (Algorithm 2). However, these results suggest that even the discrete variant often converges—e.g. under gross substitutes—as long as the λ -multipliers in the algorithm are sufficiently small. There are standard methods to choose λ_g so that the algorithm neither "overshoots" the equilibrium nor is so small that convergence is impractically slow. A simple version of a backtracking algorithm which has proven very useful in practice is now presented. First, let Line 3 in Algorithm 2 read:

Set $\lambda_g = 1$ for all $g \in [1..k - 1]$

Next, replace Line 8 of Algorithm 2 by Algorithm 3. In Algorithm 3 r denotes the current iteration. Line 5 in Algorithm 3 tests if the current step-size "overshoots the equilibrium too far", i.e. the demand changes sign and there is no significant decrease in the absolute value of the demand. If the current step-size is too large, it is divided by two. Again, there are more sophisticated ways to do this, but this is a simple way that has worked well in

```

done = false
Repeat
   $p_g^{r+1} = p_g^r + \lambda_g \sum_i z_{ig}(p_g^r)$ 
  if  $sign(\sum_i z_{ig}(p_g^r)) \neq sign(\sum_i z_{ig}(p_g^{r+1})) \wedge |\sum_i z_{ig}(p_g^{r+1})| > 0.9 |\sum_i z_{ig}(p_g^r)|$ 
     $\lambda_g = \lambda_g/2$ 
     $p_g^{r+1} = p_g^r$ 
  else
    done = true
    if  $\lambda_g \leq 0.5$ 
       $\lambda_g = 2\lambda_g$ 
Until done

```

Algorithm 3: SIMPLE BACKTRACKING ALGORITHM FOR DETERMINATION OF STEP-SIZE.

practice. Loosely we can say that Algorithm 3 resembles a binary search for the proper step-size, whereas more sophisticated methods are more Newton-like, also in the determination of the proper step-size. For a detailed description of such methods, see e.g. Press et al. (Press et al., 1994, p. 385). The complexity of each price update of the basic price tâtonnement is $\mathcal{O}(nk)$ (the backtracking not included) where n is the number of agents, and k is the number of commodities.

Within computational multiagent systems, Wellman has developed a general equilibrium based simulation called WALRAS (Wellman, 1993). The tâtonnement process used in that market environment differs from the classical price tâtonnement. In WALRAS the auctioneer sends a price vector to the agents. Each agent then replies with demand functions for each commodity, treating the prices of every other commodity as fixed. These demand functions can be submitted asynchronously in arbitrary order. When an auctioneer has received one or more new demand functions for a commodity, a new equilibrium for that commodity is computed and a new market clearing price is obtained. This new price is sent to the agents and typically this causes the agents to revise and resubmit their demand functions for *other* commodities. Under certain conditions, this process still converges to an equilibrium (Cheng & Wellman, 1998). As in tâtonnement, trades in WALRAS only occur after the market process has converged (close) to an equilibrium. Because of the asynchronous nature of WALRAS, it is difficult to give relevant complexity measures. What can be said is that the complexity of updating every price once using one demand function for each commodity and for every agent is $\mathcal{O}(nk)$. Similarly, it is difficult to determine the complexity in terms of the accepted error, but empirical results strongly suggest linear convergence, i.e. that the number of iterations is $\mathcal{O}(-\log \epsilon)$, where ϵ is the error (Cheng & Wellman, 1998; Ygge, 1998).

A clear advantage of both basic price tâtonnement and WALRAS is their simplicity. The implementation is straightforward, and following the price update for debugging purposes is very easy. A disadvantage with both algorithms is convergence speed. It has been demon-

strated that in some cases the convergence of WALRAS is too slow to be useful, even in cases where it is guaranteed to converge (Ygge, 1998, pp. 76–77). One approach to significantly speeding up the convergence—under the assumption that the demand functions are reasonably smooth—is to use a Newton-Raphson algorithm. It uses derivative information to adjust the size of the price changes. A standard version thereof updates the prices using the following formula:

$$\mathbf{p}^{i+1} = \mathbf{p}^i - \lambda \cdot \nabla \mathbf{z}^{-1}(\mathbf{p}^i) \cdot \mathbf{z}(\mathbf{p}^i), \quad (47)$$

where $i+1$ and i denote iterations, λ is a step size, and $\nabla \mathbf{z}(\mathbf{p})$ is the gradient matrix defined by $\nabla z_{ij}(\mathbf{p}) = \frac{\partial z_i(\mathbf{p})}{\partial p_j}$. A proper value for λ can be determined at run-time by a backtracking algorithm (Press et al., 1994, pp. 384 – 385).

As the computational task of the auctioneer includes summing the n demand functions and solving the $k - 1$ linear equation system, the time of each update is $\mathcal{O}(nk^{2.496})$ (Pan, 1984)²⁵. Under certain smoothness assumptions and under the assumption that \mathbf{z} is decreasing only if the prices are moving towards the market clearing prices, this *Newtonian price tâtonnement* is guaranteed to converge with quadratic convergence, i.e. the required number of iterations is $\mathcal{O}(\log(-\log \epsilon))$, where ϵ is the error (Press et al., 1994).

Compared to the basic price tâtonnement and the WALRAS algorithms, the Newtonian price tâtonnement requires partial derivatives of the demand functions, and imposes a heavier computational burden on the auctioneer at each iteration. Furthermore, the size of the message from each agent at each round is $k - 1 + (k - 1)^2 (\mathbf{z}(\mathbf{p}^i)$ and $\nabla \mathbf{z}(\mathbf{p}^i)$, the latter possibly inverted) which is relatively large compared to the alternatives. On the other hand, partial derivatives can be easily and inexpensively numerically approximated and the number of iterations is typically reduced so significantly that this algorithm is more efficient both in terms of computation and communication (Ygge, 1998). The presented Newtonian price tâtonnement requires significantly fewer iterations than the alternative methods because the step length is based on the partial derivatives and because all prices are updated in parallel, i.e. it is a *simultaneous* tâtonnement process (Takayama, 1985).

In addition to the above price tâtonnement methods, there is a corresponding *quantity tâtonnement* or Marshallian quantity adjustment method, (e.g. (Takayama, 1985, p. 297)):

$$\frac{\partial q_g}{\partial t} = \lambda_g (D(q_g) - S(q_g)), \quad (48)$$

where $D(q_g)$ is the price that the demand side is willing to pay for quantity q_g of commodity g , and $S(q_g)$ is the price at which the supply side is willing to sell quantity q_g . Essentially, if the price at the supply side exceeds the price at the demand side, the transferred (produced) quantity decreases and vice versa.

Traditionally, the tâtonnement methods have been used for analyzing the stability of equilibrium and as an argument for why certain markets strive toward an equilibrium. More recently, these basic mechanisms have been used as the algorithms for finding market clearing prices in computational markets. Indeed, for example the WALRAS system is more or less a direct implementation of basic price tâtonnement, with some extensions in terms of, e.g., asynchrony. However, finding efficient algorithms based on quantity tâtonnement

25. One should bare in mind though that the type of algorithms with better asymptotic behavior than standard methods (typically with complexity $\mathcal{O}(k^3)$) are normally uninterestingly slow for $k < 300$.

is not as straightforward. In a real setting with more than two agents, it is not obvious how to apply Equation (48). There may not be an obvious distinction between the supply and demand sides: some agents may choose to buy or sell depending on the market price. Furthermore, even if such a distinction is possible, it is nontrivial to distribute q_g for the respective sides so that each side announces only a single price. That is, when determining S in Equation (48) (analogously for D) it has to be determined how q_g is provided by the different suppliers so that all production is performed at the same (reported) marginal price. Determining this division of supply is significantly harder than determining q_g itself. If one accepts a very broad generalization of quantity tâtonnement so that it includes any algorithm in which the allocations of the agents are updated as a function of previous allocations, it is possible to use quantity tâtonnement for finding an equilibrium. Such generalizations are rather far from the original idea of quantity tâtonnement as captured by Equation (48), and some times the terms *resource-based* or *resource-oriented* mechanisms are used instead to denote algorithms that search for equilibrium with the resource as the free search parameter, see e.g. (Kurose & Simha, 1989; Ygge, 1998). As we show below, useful quantity tâtonnement algorithms for realistic settings have little in common with Equation (48) above.

The principles of quantity tâtonnement differ fundamentally from the principles of price tâtonnement. Instead of asking an agent how much it is willing to buy or sell at a specific price, one might ask it how much it is willing to pay for an infinitesimal additional amount of each commodity (a set of prices) at the current allocation. That is, each agent, i , can be viewed as holding a price function $\underline{p}_i(\mathbf{z}_i)$, rather than a demand function $\underline{z}_i(\mathbf{p})$. If there is a bijective mapping between $\underline{p}_i(\mathbf{z}_i)$ and $\underline{z}_i(\mathbf{p})$ (i.e. $\underline{p}_i^1 \neq \underline{p}_i^2 \Leftrightarrow \underline{z}_i^1 \neq \underline{z}_i^2$), then the equilibrium condition, $\sum_i \underline{z}_{ig} = 0$, corresponds to

$$\begin{cases} \underline{p}_{ig}(\mathbf{z}_i) \leq \underline{p}_g, & z_{ig} = \underline{z}_{ig}^l \\ \underline{p}_{ig}(\mathbf{z}_i) = \underline{p}_g, & \underline{z}_{ig}^l < z_{ig} < \underline{z}_{ig}^u \\ \underline{p}_{ig}(\mathbf{z}_i) \geq \underline{p}_g, & z_{ig} = \underline{z}_{ig}^u, \end{cases} \quad (49)$$

where \underline{z}_i^l and \underline{z}_i^u are (revealed) lower and upper bounds of the net demand of agent i . The basic idea in quantity tâtonnement is to try different reallocations until Equation (49) is fulfilled, instead of evaluating different prices until $\sum_i \underline{z}_{ig} = 0$ as in price tâtonnement.

One advantage of quantity tâtonnement is that the net total reallocated resource is always kept at zero, i.e. every allocation in the search for the equilibrium is feasible. Therefore, the algorithm is an anytime algorithm: it can be terminated at any time with a feasible solution in hand.²⁶ Another advantage is that the price functions used in quantity tâtonnement are more closely related to utility functions than the demand functions of price tâtonnement which results in higher computational efficiency. Most treatments of market-based search only discuss the complexity of finding an equilibrium once the agents' supply and demand functions are known. However, it may be computationally complex for each agent to generate its optimal supply/demand decision given the current prices. Solving for the demand function from a given utility function and endowments can be a nontrivial optimization problem. The price function, on the other hand, is merely the quotient of two

26. There are recent ways to construct anytime algorithms also based on price tâtonnement under some restrictions (Ygge & Akkermans, 1997).

partial derivatives and can be inexpensively numerically estimated in one iteration (Ygge & Akkermans, 1998). Depending on the application, the advantage of being able to utilize price functions instead of demand functions varies. In some cases the hard problem is to come up with the utility function or production possibilities set. For example, if the agent is a manufacturer, it may need to solve several planning and scheduling problems just to construct its production possibilities set from which it has to choose the profit maximizing production plan. Furthermore, each agent has to go through this local deliberation at every iteration of the market protocol because prices change, and that affects what the optimal plan for each agent is.

For the two commodity case a *Netwonian quantity tâtonnement* is given by (Ygge & Akkermans, 1998, 2000):

$$z_{i1}^{r+1} = z_{i1}^r - \lambda \frac{p_i(z_{i1}^r) - \frac{\sum_{j=0}^n \frac{p_j(z_{j1}^r)}{p'_j(z_{j1}^r)}}{\sum_{j=0}^n \frac{1}{p'_j(z_{j1}^r)}}}{p'(z_{i1}^r)}. \quad (50)$$

The quotient including the sums is a weighted average of all prices (weighted by the derivatives of the prices), and the change in allocation is the difference between the price of the agent under observation and this weighted average, divided by the derivative of the price times a step size, λ .

In the multi-commodity case the update algorithm is somewhat more complicated (Ygge & Akkermans, 1998, 2000):

$$z_i^{r+1} = z_i^r - \lambda \cdot \diamond p_i^r (p_i^r - \langle p \rangle^r), \quad (51)$$

where p_i^r and $\diamond p_i^r$, are abbreviations for $p_i(z_i^r)$, and $(\nabla p_i(z_i^r))^{-1}$, respectively. The term $\langle p \rangle^r$ is defined by

$$\langle p \rangle^r = p_n^r + \nabla p_n^r \left(\sum_{j=1}^n \diamond p_j^r \right)^{-1} \cdot \diamond p_n^r \sum_{j=1}^n \diamond p_j^r (p_j^r - p_n^r) \quad (52)$$

and can be interpreted as the *expected price*. (∇p_i^r is an abbreviation for $\nabla p_i(z_i^r)$.) This would have been the market clearing price if the current value of ∇p had been independent of the allocation (i.e. if $p_{ig}(z_i)$ had been linear functions).

From the above we see that even though the termination condition is the same, the intermediate allocations, i.e. the path to the solution, will depend on the ordering of agents. The final solution (as captured by Equation (49)) is the same though.

The update of z includes matrix inversion, and the computation is done separately for each of the n agents, so each iteration takes $\mathcal{O}(nk^{2.496})$ time (Pan, 1984). Since the algorithm (under certain smoothness conditions) is a quadratic scheme, the required number of iterations is $\mathcal{O}(\log(-\log \epsilon))$, where ϵ is the error. The gradient matrices are typically computed locally (analytically or numerically) by the participating agents and communicated to the auctioneer.

A delicate issue in quantity tâtonnement is proper management of boundaries.²⁷ At each step of the algorithm some z may end up outside its boundaries. For example, if a

27. In price tâtonnement this is managed locally by each agent. When posting a demand at a specific price, the agent ensures that the demand does not violate its boundaries. Therefore, this is not a concern for the auctioneer.

consumer has a lower bound on z_1 of 0 (e.g., it has no endowment and cannot produce), but have a very low valuation compared to the other agents, it might be assigned a negative value by the algorithm. Therefore, some mechanism is needed to manage the boundaries to maintain feasible allocations. In this example, resource need to be reallocated from other agents in order to obtain $z_1 = 0$ for the agent in question. The proper management of the boundaries is rather complicated, and a thorough presentation is found elsewhere (Ygge & Akkermans, 1998, 2000).

As seen from Equations (51) and (52), and because of the required management of boundaries, quantity tâtonnement is relatively complicated conceptually. Because this type of algorithms may speed up the computation in various applications, we recommend them mainly for settings where simpler methods are too inefficient (Ygge, 1998, pp. 82–86).

In distributed environments with a large number of agents, naive implementations of the algorithms above are not very efficient. In a recent approach, multiple auctioneers were used in a hierarchy to better distribute the computational burden while keeping the communication low (Andersson & Ygge, 1998).

The general multi-commodity case in highly distributed environments is very complicated and presumably application specific algorithms are required for very high efficiency. For the two commodity case on the other hand, a recent algorithm, COTREE, provides highly efficient distributed computation, both for price and quantity tâtonnement (Andersson & Ygge, 1998). The basic idea is distributed pairwise combinatorial aggregation of demand/price functions in a binary tree, in which all producers and consumers are represented as leaves. Once all functions have been aggregated to one function (at the root node), the equilibrium can be determined and the resource is recursively allocated downwards in the tree.

References

- Abreu, D., & Matsushima, H. (1991). Virtual implementation in iteratively undominated strategies II: incomplete information. Tech. rep.. mimeo.
- Andersson, A., & Ygge, F. (1998). Managing large scale computational markets. In El-Rewini, H. (Ed.), *Proceedings of the Software Technology track of the 31th Hawaiian International Conference on System Sciences (HICSS31)*, Vol. VII, pp. 4 – 14. IEEE Computer Society, Los Alamos. ISBN 0-8186-8251-5, ISSN 1060-3425, IEEE Catalog Number 98TB100216.
- Barbera, S., & Jackson, M. O. (1995). Strategy-proof exchange. *Econometrica*, 63(1), 51–87.
- Borenstein, S., & Bushnell, J. (1998). An empirical analysis of the potential for market power in California's electricity industry. (Available at <http://nberws.nber.org/papers/W6463>).
- Cheng, J. Q., & Wellman, M. P. (1998). The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, 12, 1–24.
- Duggan, J. (1997). Virtual Bayesian implementation. *Econometrica*, 65(5), 1175–1199.
- Ellickson, B. (1993). *Competitive equilibrium: Theory and Applications*. Cambridge University Press.
- Fudenberg, D., & Levine, D. K. (1993). Self-confirming equilibrium. *Econometrica*, 61(3), 523–545.
- Hu, J., & Wellman, M. P. (1996). Self-fulfilling bias in multiagent learning. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pp. 118–125 Keihanna Plaza, Kyoto, Japan.
- Hurwicz, L. (1972). On informationally decentralized systems. In McGuire, C., & Radner, R. (Eds.), *Decision and Organization*. Amsterdam: North Holland.
- Ilic, M. (1999). Generation strategies for gaming transmission constraints - will the deregulated electric power market be an oligopoly?. In *Proceedings of the 32nd Hawaiian International Conference on System Sciences (HICSS32)*. IEEE Computer Society.
- Jackson, M. O. (1991). Bayesian implementation. *Econometrica*, 59, 461–478.
- Kalai, E., & Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, 61(5), 1019–1045.
- Kehoe, T. J. (1991). Computation and multiplicity of equilibria. In Hildenbrand, W., & Sonnenschein, H. (Eds.), *Handbook of mathematical economics*, Vol. 4, chap. 38, pp. 2049–2144. Elsevier Science Publishers.
- Klemperer, P. D., & Meyer, M. A. (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6), 1243–1277.

- Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton University Press.
- Kurose, J. F., & Simha, R. (1989). A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5), 705–717.
- Malinvaud, E. (1985). *Lectures on Microeconomic Theory*. North-Holland.
- Mas-Colell, A., Whinston, M., & Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Moore, J. (1992). Implementation in environments with complete information. In Laffont, J.-J. (Ed.), *Advances in Economic Theory: The Proceedings of the Congress of the Econometric Society*. Cambridge University Press.
- Mullen, T., & Wellman, M. P. (1995). A simple computational market for network information services. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pp. 283–289 San Francisco, CA.
- Nash, J. (1950). Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36, 48–49.
- NordPool (1998). The elspot market – the spot market. (Available from www.nordpool.no).
- Pan, V. (1984). How can we speed up matrix multiplication?. *SIAM Review*, 26, 393–416.
- Press, W., Teukolsky, S., Vetterling, W., & Flannery, B. (1994). *Numerical Recipes in C* (2nd edition). Cambridge University Press.
- Roberts, D. J., & Postlewaite, A. (1976). The incentives for price-taking behavior in large exchange economies. *Econometrica*, 44(1), 115–127.
- Sandholm, T. W. (1993). An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pp. 256–262 Washington, D.C.
- Sandholm, T. W. (1996a). Limitations of the Vickrey auction in computational multi-agent systems. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pp. 299–306 Keihanna Plaza, Kyoto, Japan.
- Sandholm, T. W. (1996b). *Negotiation among Self-Interested Computationally Limited Agents*. Ph.D. thesis, University of Massachusetts, Amherst. Available at <http://www.cs.wustl.edu/~sandholm/dissertation.ps>.
- Sandholm, T. W., & Lesser, V. R. (1995). Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pp. 328–335 San Francisco, CA. Reprinted in *Readings in Agents*, Huhns and Singh, eds., pp. 66–73, 1997.
- Sandholm, T. W., & Lesser, V. R. (1996). Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pp. 126–133 Portland, OR. Extended version: University of Massachusetts at Amherst, Computer Science Department technical report 95-72.

- Sandholm, T. W., Sikka, S., & Norden, S. (1999). Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 535–540 Stockholm, Sweden. Extended version: Washington University, Department of Computer Science technical report WUCS-99-04.
- Sandholm, T. W., & Ygge, F. (1997). On the gains and losses of speculation in equilibrium markets. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 632–638 Nagoya, Japan.
- Sandholm, T. W., & Zhou, Y. (1999). Revenue equivalence of leveled commitment contracts. In *AAAI-99 Workshop on Negotiation: Settling conflicts and identifying opportunities* Orlando, FL, July. Extended version: Washington University, Department of Computer Science technical report WUCS-99-03.
- Scarf, H. E. (1967). The approximation of fixed points of a continuous mapping. *SIAM Journal of Applied Mathematics*, 15, 1328–1343.
- Simmons, G. F. (1972). *Differential Equations with Applications and Historical Notes*. McGraw-Hill.
- Takayama, A. (1985). *Mathematical Economics*. Cambridge University Press.
- Varian, H. R. (1992). *Microeconomic analysis*. New York: W. W. Norton.
- Wellman, M. (1993). A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1, 1–23.
- Wellman, M. (1994). A computational market model for distributed configuration design. In *Proc. 12th National Conference on Artificial Intelligence (AAAI-94)*, pp. 401–407 Seattle, WA.
- Wellman, M. P., & Hu, J. (1998). Conjectural equilibrium in multiagent learning. *Machine Learning*, 33, 179–200.
- Wolak, F. A. (1998). Market design and the behavior of prices in restructured electricity markets: An international comparison. (Available from <http://www.stanford.edu/~wolak/>).
- Yamaki, H., Wellman, M. P., & Ishida, T. (1996). A market-based approach to allocating QoS for multimedia applications. In Tokoro, M. (Ed.), *Proceedings of the Second International Conference on Multi-Agent Systems ICMAS'96*, pp. 385–392. AAAI Press, Menlo Park, CA.
- Ygge, F. (1998). *Market-Oriented Programming and its Application To Power Load Management*. Ph.D. thesis, Department of Computer Science, Lund University. ISBN 91-628-3055-4, CODEN LUNFD6/(NFCS-1012)/1-224/(1998) (Available from <http://www.enersearch.se/ygge>).

- Ygge, F. (1999). Energy resellers – an endangered species?. In Moukas, Sierra, & Ygge (Eds.), *Proceedings of the Agent-Mediated Electronic Workshop in conjunction with the International Joint Conference on Artificial Intelligence 1999*.
- Ygge, F., & Akkermans, J. M. (1996). Power load management as a computational market. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pp. 393–400 Keihanna Plaza, Kyoto, Japan.
- Ygge, F., & Akkermans, J. M. (1997). Duality in multi-commodity equilibrium computations and any-time algorithms. In Zhang, C., & Lukose, D. (Eds.), *Proceedings of the Third Australian Workshop on DAI*, pp. pp. 65 – 78 Perth, Australia.
- Ygge, F., & Akkermans, J. M. (1998). On resource-oriented multi-commodity market computations. In Demazeau, Y. (Ed.), *Proceedings of the Third International Conference on Multi-Agent Systems ICMAS'98*, pp. 365 – 371. IEEE Computer Society.
- Ygge, F., & Akkermans, J. M. (2000). On resource-oriented multi-commodity market computations. *Autonomous Agents and Multi-Agent Systems*, 3(1). In press.