

# A Framework for Analyzing the Economics of a Market for Grid Services\*

Robin Mason<sup>1</sup>, Costas Courcoubetis<sup>2</sup>, and Natalia Miliou<sup>2</sup>

<sup>1</sup> University of Southampton, Highfield, Southampton, United Kingdom

<sup>2</sup> Athens University of Economics and Business, Athens, Greece

**Abstract.** This paper provides a single broad model for the analysis of a range of issues underlying a market for Grid services. The demand and the supply sides of such a market are being treated separately and the relation between the two sides is being studied. We provide numerical results in order to derive conclusions about the viability of a market for Grid services. Underlying our model are parameters such as the cost technologies, the random processes driving demand and supply and the size of the market. We study the effect of the model's parameters, such as risk aversion or the durability of resources, on the system's behavior, eg. on the clearing price or volume of trade.

**Keywords:** Grid economics, Grid services, Grid market, spot market, market failure, competitive equilibrium.

## 1 Introduction

The remarkable evolution of computational power, as well as the growth of internet, have changed the approach of research on computational resources. Namely, computational resources have become less expensive and storage capabilities have increased dramatically. Clustering geographically remote computational resources and using them as one unified resource has become not only feasible but also economically beneficial. In this work we treat the question of how computational resources should be traded. We demonstrate the benefits (to both sellers and buyers of Grid services) from trade. And we examine how these benefits arise in a particular trading mechanism.

We analyse a competitive market model, with no market power on either the buyer or provider side. We opt for this approach, for two main reasons; it employs a mechanism that is simple yet effective with desirable economic properties, and it is an extension of the standard stock market mechanism for storable commodities. We view this as the limit of a double auction in which the number of traders becomes large.<sup>1</sup> The properties of double auctions in static settings have been

---

\* Work partially supported by the EU-funded research project GRIDECON, FP6 - 2005 - IST5 - 033634.

<sup>1</sup> This double auction model has been adopted in the market mechanism developed in the project GRIDECON.

investigated both theoretically and empirically. The theory of double auctions (see e.g. [2]) shows that every non-trivial equilibrium of a double auction converges to the competitive market outcome as the number of auction participants grows. Moreover, the rate of convergence is fast, namely of the order of  $1/N^{2-\alpha}$  for any  $\alpha > 0$ , where  $N$  is the number of auction participants. Satterthwaite and Williams, [6], have shown that the worst-case speed of convergence of the double auction is as high as for any other type of auction. This work strengthens earlier results of e.g. [5].

The experimental evidence backs up these theoretical results. Even from early work, such as [1] and [10], experiments have found that double auctions yield outcomes that are close to fully efficient, even with a small number of human participants. In [3] it was shown that the same is true with "zero intelligence" traders. As it is explained in [4], the rules of the double auction are such that only a small number of participants is needed to get (very close) to the competitive outcome. Work by Cliff, [7], has challenged this conclusion. For example, he shows that, under some circumstances, agents who are less than fully rational can fail to get close to equilibrium in a double auction. But note that this does not necessarily imply that the efficiency of a double auction is low (since a high degree of efficiency just requires that high value buyers are generally matched with low value sellers). Moreover, recent work (by e.g., Preist and van Tol in [8]) extending Cliff's analysis suggests that the rate at which actual play approaches the competitive outcome is quite high. See Parsons et al. in [9] for a very good summary of the literature.

Some complementary work is in [11]. The authors analyze there the use of real options traded at an additional contract market, to efficiently manage economical issues arising from the realization of a flexible resource reservation scheme. Necessary conditions so that even risk neutral agents have incentives to participate in such a market, as it increases their expected utility, are being derived.

In this work, we develop a framework for analysing the economic properties of the competitive equilibrium of a market for Grid services. The objectives of the framework are mainly to provide a single broad model, instead of a set of separate models, within which a range of issues can be assessed and to look separately at the demand and supply sides of the market, while recognising the relation between the two sides.

This framework also permits validation of the model, by using numerical and simulation methods to assess the broad features of the model and derives a set of conclusions about the viability of a market for Grid services (both spot and forward), based on the underlying parameters of the analysis: for example, the cost technologies, the random processes driving demand and supply, the size of the market (number of participants).

In section 2 a discussion on the competitive market model and our basic assumptions is being presented. The definition of our model, as well as the parameters studied in this paper is being described in section 3. A general solution for the competitive equilibrium with a Grid spot market is being presented in section 4. The simulation environment in which our experiments are being

conducted is described in section 5. Finally, the analysis of the situation when the market fails to satisfy the demand in computational resources is being analyzed in section 6.

## 2 Our Approach

The starting point of economic analysis is usually a competitive market model, with no market power on either the buyer or provider side. In this framework, the price and quantity of Grid services are determined by competitive equilibrium in which demand and supply are equal. In order to analyse this equilibrium outcome, we must first model explicitly the demand and supply sides of the market for Grid services. We shall do so by considering the dynamics of resource allocation for both users and providers of Grid services.

On the demand side, each user is subject to random arrival of jobs that need to be completed. These jobs might also be of different sizes or durations. Each user has to decide whether to use the Grid or buy its own resources. Economies of scale, through e.g., multiplexing, mean that the Grid may have lower costs; but uncertainty over total demand and supply means that the user cannot be certain of getting adequate resources at a given price from the Grid. The user's decision optimally trades off these factors. The aggregate of different users' decisions determines at any time the total demand for Grid resources.

The supply side is modelled similarly. Each provider of Grid resources has its own capacity for computational tasks; and faces random usage patterns for that capacity. It releases capacity to the Grid whenever its own demand for resources is less than its own capacity. The Grid then offers a way in which a provider of computational resources can recover the costs of capacity investment, in the event that its own demand falls short of capacity. The aggregate of different providers' decisions determines at any time the total supply of Grid resources.

Hence, the demand and supply sides of the model are related in a two-stage procedure. In the first stage, the economic agents choose the amount of computational resources to buy for themselves. In the second stage, each agent observes its own idiosyncratic demand for computational resources. Those with too little capacity buy resources from the Grid, forming the demand side of the model. Those with too much capacity offer resources to the Grid: the supply side. Differences between the agents—primarily heterogeneity in the size and duration of the jobs—mean that some agents buy very little capacity and are systematically on the demand side of the market.

Under this approach, it is assumed that there is a single homogeneous good that is to be sold through the market. Hence, no differentiation among resources is feasible. Such a commodity market assumes a single market price for a unit of the good sold. There are clear indications that current virtualization technologies justifies such a commoditization assumption, for instance Amazon EC2 cloud sells virtual machines of similar specifications. That is, it is assumed that all market actors, both consumers and producers, have a small fraction of the total market share. This way, they cannot affect the market price by means of

a strategy, thus the market is perfectly competitive and all agents are price-takers. It is clear that this assumption does not hold in the existing oligopolistic commercial Grid environment where few large providers dominate the market.

The first benchmark to analyse is the sequence of competitive spot markets, where demand and supply are equalised in each period. The spot market is analysed under a number of different scenarios concerning the stochastic processes determining the demand and supply sides of the market. In order to analyse the use of forward and options markets, we shall measure the extent to which Grid market participants are willing to pay to avoid the risk attached to an uncertain Grid market price. This is one key source of uncertainty facing Grid participants. (The other is their own demand for computational services; we do not consider insurance against this sort of risk. In a sense, the Grid market performs this role, since it allows users to buy and sell resources, in the event that their demand exceeds or falls short of their own supply of computational resources.) Risk averse users, faced with uncertainty about demand and supply conditions, will want to buy insurance against that uncertainty. With a very large number of users (in the perfectly competitive limit, a continuum), there is no aggregate uncertainty. But for each user and provider, of course, there would be uncertainty in each time period. The most direct analysis is to calculate the risk premium that users are willing to pay to insure themselves entirely against market fluctuations. The aggregate risk premium gives a measure of the potential for forward and options markets for Grid services to develop. This risk premium can be calculated numerically under a number of different scenarios, using simulations of the competitive market outcome over time.

### 3 The Model

There are  $N$  agents, indexed by  $i$ . There are an infinite number of discrete time periods, indexed by  $t$ . Each time period  $t$  is divided into (up to) three stages. During the first stage the amount of resources to be bought is being determined: agent  $i$  chooses to buy an amount  $x_i(t) \in \mathbb{R}_+$  of computational resource. The cost of buying an amount  $x$  of this resource is  $c(x)$ . We make the following general assumptions about the cost function: buying nothing costs nothing  $c(0) = 0$ ; variable cost of acquiring the resource is increasing and convex in the resource amount:  $c'(\cdot) > 0$ ,  $c''(\cdot) > 0$  for  $x > 0$ . Taken together, these features mean that (a) agents can rationally choose not to buy any resource; (b) demand for resources is bounded (due to the convexity of variable costs).  $c(\cdot)$  is a key (functional) parameter of the model.

There are a number of modelling options for how resources evolve over time. One extreme is that a resource is non-durable, so that each agent starts each period with a zero stock of the resource. The opposite extreme is that the resource is infinitely durable; each agent therefore accumulates the resource over time, with no loss or degradation of the resource. In this case, the resource available to agent  $i$  in period  $t$  is  $X_i(t) = X_i(t-1) + x_i(t)$ . In between these two extremes, the resource may be durable but suffer from depreciation

(e.g., technical obsolescence). We model this by using the *depreciation* coefficient  $1 - \beta$ , where  $0 \leq \beta \leq 1$ . Hence, the resource available to agent  $i$  in period  $t$  becomes  $X_i(t) = \beta X_i(t-1) + x_i(t)$ .

The second stage is initiated when the demand is being realized. Agent  $i$ 's demand parameter  $y_i(t) \geq 0$  is a random variable at the time that it chooses its resource level. The value of the parameter is realized after the resource level is chosen, but before the Grid market operates.

A key modelling question is the process that  $y_i$  follows. One representation is  $y_i(t) = \alpha_i y_i(t-1) + \epsilon_{it}$  where  $y_i(0)$  is a parameter; and  $\epsilon_{it}$  is drawn from a distribution with zero mean (e.g., a normal), subject to  $y_i \geq 0$ . In this first-order autoregressive process,  $\alpha_i$  is the degree of persistence in the process for  $\{y_i(t)\}$ . If  $\alpha_i = 0$ , then  $\{y_i(t)\}$  is a random walk.

The third stage is the stage of the actual operation of the Grid market. Agents participate in the Grid, according to their resource levels and demand realizations. Potentially, an agent's maximization problem is very complex: choices in a given period can depend on choices made in the previous period; the current period's environment (e.g., realisations of demands); and the evolution of the system in the future. It would be extremely difficult to solve this system fully. Instead, we make the simplifying assumption that agents are 'myopic', solving in each period a static maximization problem subject only to the constraints generated by previous choices. In particular, the effect of current choices on future maximization problems is ignored.

In this case, agent  $i$ 's maximization problem in period  $t$  is

$$\max_{z \geq 0} \{y_i(t)u(z) - q(t)(z - (\beta X_i(t-1) + x_i(t)))\}. \quad (1)$$

$u(\cdot)$  is a utility function that is common to all agents and  $z$  models the aggregate need in computational resources; all heterogeneity is therefore contained within the realizations of  $\{y_i\}$ .  $u(\cdot)$  is assumed to be (weakly) increasing: for all agents, a greater use of resources cannot decrease utility.  $q(t)$  is the unit resource price that prevails in the Grid market at time  $t$ .

An important modelling question is the functional form of the utility function. Since we shall be interested in considering forward markets, which effectively provide insurance for agents, risk aversion will be a key component of the model. For this, we need concavity of the utility function.

## 4 Solving for Competitive Equilibrium with a Grid Spot Market

In a competitive Grid spot market, a market-clearing condition must hold:

$$\sum_i z_i^* = \sum_i X_i(t), \quad (2)$$

where  $z_i^*$  is agent  $i$ 's optimal choice of  $z$ . Note that this market-clearing condition will indeed hold with equality. There cannot be more (instantaneous) demand

than there is supply: that is, it is not possible for  $\sum_i z_i^* > \sum_i X_i(t)$ . If there were excess supply—if  $\sum_i z_i^* < \sum_i X_i(t)$ —then this would imply that the price of Grid resources,  $q(t)$ , would be zero. Then there would be no cost to increasing use of the resource (see equation (1)); consequently, each agent would raise its  $z_i$  until the market-clearing condition (2) holds with equality.

Each agent  $i$  perceives that the Grid price  $q(t)$  is unaffected by its choice of resource  $x_i(t)$  and resource consumption  $z_i$ . (This follows from the competitive market environment that we consider.) In this case, the first-order condition characterizing the optimal consumption of resources is

$$z_i^*(q(t); y_i(t)) = \begin{cases} u'^{-1}\left(\frac{q(t)}{y_i(t)}\right) & y_i(t) > 0, \\ 0 & y_i(t) = 0. \end{cases} \quad (3)$$

Note, therefore, that agent  $i$ 's optimal consumption of the resource is independent of its own resource level  $x_i$ .  $q(t)$  satisfies

$$\sum_i z_i^*(q(t); y_i(t)) = \sum_i X_i(t), \quad (4)$$

and hence is a function of the vector of demand realizations  $\mathbf{y}(t)$  and the total resource level  $X(t) \equiv \sum_i X_i(t)$ . Denote the market-clearing Grid price as  $q(\mathbf{y}(t), X(t))$  to emphasize this dependence.

Supposedly agent  $i$  solves the following optimization problem, when choosing its resource level

$$\max_x \mathbb{E}_y [y_i(t)u(z_i^*(q(\mathbf{y}(t), X(t)) - q(\mathbf{y}(t), X(t)))(z_i^*(q(\mathbf{y}(t), X(t)) - (\beta X_i(t-1) + x) - c(x))]. \quad (5)$$

Consistent with the competitive market setting, agent  $i$  does not consider the effect of its choice of  $x_i$  on  $q$  since  $X(t)$  is much larger than  $x$ . Therefore, the optimization problem that is being solved in (5) is

$$\max_x \mathbb{E}_y [q(\mathbf{y}(t), X(t))]x - c(x). \quad (6)$$

Essentially,  $x_i$  maximizes the net benefit of the user from the sale to the Grid at price  $q$ . Hence the first-order condition for its resource choice is

$$x_i^* = c'^{-1}(\mathbb{E}_y [q(\mathbf{y}(t), X(t))]) \quad (7)$$

if  $x_i^* > 0$  is optimal.

However, note that this is not the exact optimization problem that an agent faces when choosing its resource level. Namely, the agent has to take into consideration the fact that resources being bought at time slot  $t$  in general last for multiple time slots and their net present value by selling them to the Grid should be used in the optimization problem 5 (instead of the revenue at time  $t$  only).

Namely, if  $r$  is the rate of return that could be earned on an investment with similar risk, the net present value (NPV) of  $x$  (depreciated by  $1 - \beta$  in each

period) is  $\sum_{s=t}^{\infty} xq(s) \frac{\beta^s}{(1+r)^s}$ . This problem is in general difficult to solve since the future price values are not known. Instead, we assume for the purpose of an approximation of the NPV that the price remains constant and equal to the current price. In this case the NPV simplifies to  $\frac{r+1}{r+1-\beta} xq(t)$  and the subproblem 6 is modified to

$$\max_x \frac{r+1}{r+1-\beta} \mathbb{E}_y [q(\mathbf{y}(t), X(t))] x - c(x). \quad (8)$$

As a consequence, problem (7) is modified to

$$x_i^* = c'^{-1} \left( \frac{r+1}{r+1-\beta} \mathbb{E}_y [q(\mathbf{y}(t), X(t))] \right). \quad (9)$$

Equation (9) with equations (3) and (4) define the market-clearing price  $q(t)$ ;  $\mathbf{y}$ ,  $\mathbf{X}(t-1)$ .

Hence the sequence of events to solve the model in any given period  $t$  is

- Compute the optimal resource demands  $\mathbf{x}^*(t)$  (which depends on expectations of the realizations of the agents' types  $\mathbf{y}(t)$  and the existing stock of resources  $\mathbf{x}^*(t-1)$ ).
- Take the instantaneous realization of the agents' types  $\mathbf{y}(t)$ .
- For that realization and the agents' resources, compute the market-clearing price  $q(t)$ .
- Given the market-clearing price  $q(t)$  and the realization of agents' types  $\mathbf{y}(t)$ , compute the optimal demand  $\mathbf{z}^*(t)$ .
- Compute metrics of the performance of the Grid market:  
the aggregate indirect utility of participants:

$$U(t) \equiv \sum_i \mathbb{E}_y [(y_i(t)u(z_i^*(t)) - q(t)(z_i^*(t) - (\beta X_i(t-1) + x_i^*(t))) - c(x_i^*(t)))] \quad (10)$$

and the realized volume of trade:  $V(t) \equiv \sum_i |z_i^*(t) - (\beta X_i(t-1) + x_i^*(t))|$ .

#### 4.1 The Autarkic Benchmark

The outcome with a Grid spot market can be contrasted to the situation in which there is no trading of resources subsequent to demand realizations. Each agent's maximization problem for the demand stage is then

$$\max_z \{y_i u(z)\} \quad \text{subject to } z \leq x_i. \quad (11)$$

The solution to this is simply  $z_i^\dagger = x_i$ : each agent consumes its own entire resource. For the optimal choice of resource under autarky, the agent faces

$$\max_x \{\mathbb{E}_y [y_i] u(x) - c(x)\} \quad (12)$$

which has a solution  $x_i^\dagger$  given by

$$\mathbb{E}_y [y_i] u'(x_i^\dagger) = c'(x_i^\dagger). \quad (13)$$

## 5 A Simulation Environment

In this section, we illustrate the main properties of a competitive Grid spot market by means of presenting a simple example and then establishing a general simulation environment.

### 5.1 An Example

To illustrate the model, this subsection looks at a very simple example. For the functions in the model, take the following forms

$$c(x) = \gamma x^2, \quad \gamma > 0; \tag{14}$$

$$u(z) = \ln(z). \tag{15}$$

Suppose that the resource is entirely non-durable, so that the agents start each period with a zero resource stock. Suppose that each agent's type  $y_i$  is drawn from the uniform distribution on  $[0, 1]$  independently each period; the types of any two agents are independent draws from the same distribution.

Given the realization of its type, agent  $i$ 's optimal demand is given by

$$z_i^* = \frac{y_i}{q}. \tag{16}$$

Hence total demand is

$$\sum_i z_i^* = \frac{\sum_i y_i}{q} \equiv \frac{Y}{q}. \tag{17}$$

Market clearing requires that  $\sum_i z_i^* = \sum_i x_i^*$ , so that

$$\frac{Y}{q} = \sum_i x_i^* \equiv X \tag{18}$$

or  $q = Y/X$ .

Agent  $i$ 's optimal resource level is given by

$$x_i^* = \frac{1}{2\gamma} \mathbb{E}_y[q] = \frac{1}{2\gamma X} \mathbb{E}_y[Y] = \frac{N}{4\gamma X}. \tag{19}$$

Since  $x_i^* = X/N$ , this means that

$$\frac{X}{N} = \frac{N}{4\gamma X} \tag{20}$$

so that

$$X = \frac{N}{2\sqrt{\gamma}}; \quad x = \frac{1}{2\sqrt{\gamma}}. \tag{21}$$



Elementary calculations show that  $Y$  is a random variable with a mean of  $N/2$  and density function of

$$\frac{1}{2(N-1)!} \sum_{k=0}^N (-1)^k \frac{N!}{k!(N-k)!} (y-k)^{N-1} \text{sgn}(x-k). \quad (22)$$

These facts can be used to compute the (expected) volume of trade

$$V = \mathbb{E} \left[ \sum_i \left| \frac{Ny_i}{2\sqrt{\gamma}Y} - \frac{1}{2\sqrt{\gamma}} \right| \right] \quad (23)$$

and the expected utility of each agent

$$U^* = \mathbb{E}_y \left[ y_i \ln(y_i) - y_i \ln(q) - y_i + \frac{q}{2\sqrt{\gamma}} - \frac{1}{4} \right]. \quad (24)$$

Note that there are two sources of uncertainty for the agent: its own demand, determined by  $y_i$ ; and the aggregate market conditions, summarised by the market-clearing price  $q$ . Because of risk aversion, both sources of risk decrease the agent's utility; that is

$$\mathbb{E}_y [y_i \ln(y_i)] < \mathbb{E}_y [y_i] \ln(\mathbb{E}_y [y_i]), \quad \text{and} \quad \mathbb{E}_y [y_i \ln(q)] < \mathbb{E}_y [y_i] \ln(\mathbb{E}_y [q]). \quad (25)$$

Since we are interested in the role of forward contracts, we calculate the size of the 'risk premium' attached to the market price, which is

$$RP \equiv \mathbb{E}_y [y_i] (\ln(\mathbb{E}_y [q]) - \mathbb{E}_y [\ln(q)]). \quad (26)$$

The risk premium is the cash amount that an agent would be willing to pay in each period in order to avoid the uncertainty attached to the market-clearing price. It is, therefore, the fee that a broker could charge to write a forward contract that insures the agent against fluctuations in the Grid price.

These variables can be contrasted with the outcome from autarky:

$$x_i^\dagger = \frac{1}{\sqrt{4\gamma}}, \quad z_i^\dagger = x_i^\dagger; \quad (27)$$

the expected utility from autarky is<sup>2</sup>

$$U^\dagger = -\frac{1}{4} \ln(4\gamma) - \frac{1}{4}. \quad (28)$$

---

<sup>2</sup> Note that the expected utility under autarky may well be negative. This is not a problem: it is simply a question of normalization. The utility functions used in this model are, technically, von Neumann-Morgenstern utility functions, and hence are unique up to affine transformations. Put more simply, any constant can be added to the agents' utilities without changing the outcome of the model. Hence negative utilities can be made positive by adding a large enough positive number. We take the simpler course of using  $U^\dagger$  as the benchmark utility level, and comparing it to  $U^*$ , the utility level with a Grid market.

## 5.2 A More General Simulation Environment

In this subsection a more general simulation environment is being presented. Before describing the environment we provide three important remarks. First, we introduce demand parameters  $\{y_i\}$  of general distributions. In the simulation experiment, we suppose that an agent's demand parameter in period  $t$  is drawn from a Gamma distribution, with a scale parameter of 1 and a shape parameter  $k_i$ . An agent's shape parameter is determined initially (before the market opens) as a random draw from a uniform distribution on  $[0,1]$ . Hence agents are now *ex ante* asymmetric: agent  $i$  has a mean demand  $\mathbb{E}[y_i] = k_i$ , so that a high  $k_i$  corresponds to high average demand. Draws between periods, and across agents, are conditionally independent (i.e., conditional on the realizations of  $\{k_i\}$ ). But note that the realization of  $k_i$  induces intertemporal dependence in the realizations of agent  $i$ 's demand levels  $\{y_i(t)\}_{t \geq 0}$ . Agents do not know in any period the initial realizations of  $\{k_i\}$ , nor the per-period realizations  $\{y_i(t)\}$ . It is common knowledge how the  $\{k_i\}$  and  $\{y_i(t)\}$  are generated. This is an extremely flexible way of specifying demand. The motivation behind choosing Gamma distribution consists in properties such as the divisibility and the relationship of Gamma distribution to the Gaussian and exponential distributions.

Secondly, we allow agents to accumulate the resource over time either completely or partially, as has been already mentioned.

Thirdly, we suppose that the utility function  $u(\cdot)$  is given by

$$u(z) = \begin{cases} \frac{z^{1-\rho}}{1-\rho} & \rho > 0, \rho \neq 1 \\ \ln(z) & \rho = 1. \end{cases} \quad (29)$$

In general,

$$u'(z) = z^{-\rho} > 0; \quad u''(z) = -\rho z^{-(1+\rho)} < 0; \quad \frac{-zu''(z)}{u'(z)} = \rho. \quad (30)$$

These three properties mean that:

1. Utility is increasing in the amount consumed.
2. Utility is strictly concave, so that agents are risk averse.
3. The Arrow-Pratt coefficient of relative risk aversion, defined as  $-zu''/u'$ , is constant. Hence agents' attitudes towards risk are parameterized completely by the constant  $\rho$ .

We assume that the cost of acquiring a resource level  $x$  is  $\gamma x^2$ , where  $\gamma > 0$ .

With this set-up, agent  $i$ 's optimal demand in period  $t$  is

$$z_i^*(t) = \left( \frac{y_i(t)}{q(t)} \right)^{\frac{1}{\rho}}. \quad (31)$$

Market-clearing requires that

$$\sum_i \left( \frac{y_i(t)}{q(t)} \right)^{\frac{1}{\rho}} = X(t) = \sum_i X_i(t), \quad (32)$$

which means that

$$q(t) = \frac{\left(\sum_i y_i(t)^{\frac{1}{\rho}}\right)^\rho}{X(t)^\rho}. \quad (33)$$

Hence

$$\mathbb{E}[q(t)] = \frac{\mathbb{E}\left[\left(\sum_i y_i(t)^{\frac{1}{\rho}}\right)^\rho\right]}{X(t)^\rho}. \quad (34)$$

Agent  $i$ 's optimal choice of resource in period  $t$  is

$$x_i^*(t) = \frac{1}{2\gamma} \frac{r+1}{r+1-\beta} \mathbb{E}[q(t)], \quad (35)$$

so that

$$x^*(t) \equiv \frac{N}{2\gamma} \frac{r+1}{r+1-\beta} \mathbb{E}[q(t)]. \quad (36)$$

Since  $X(t) = \beta X(t-1) + x^*(t)$ , this gives a non-linear equation in  $X(t)$ :

$$(X(t) - \beta X(t-1))X(t)^\rho = \frac{N}{2\gamma} \frac{r+1}{r+1-\beta} \mathbb{E}\left[\left(\sum_i y_i(t)^{\frac{1}{\rho}}\right)^\rho\right] \quad (37)$$


which can be solved numerically. With the resulting expression for  $X(t)$ , the rest of the system can be computed.

The outcome with trading in a competitive Grid spot market can be contrasted with autarky:

$$x_i^\dagger = \left(\frac{k_i}{2\gamma}\right)^{\frac{1}{1+\rho}}, \quad z_i^\dagger = x_i^\dagger. \quad (38)$$

The simulation experiment proceeds as follows:

1. We fix the number of agents.
2. Each agent starts with zero resources:  $X_i(0) = 0$  for all  $i$ .
3. A vector of demand parameters  $\mathbf{k}$  is drawn, where each  $k_i$  is an independent draw from the uniform distribution on  $[0,1]$ .
4. Given its  $k_i$  parameter, and the consequent expectations about its demand  $y_i(0)$  and the market-clearing price  $q(0)$ , agent  $i$  chooses its resource level  $x_i^*(0)$  (as described above).
5. Agent  $i$ 's demand level  $y_i(0)$  is realised, as a draw from the Gamma distribution  $\Gamma(k_i, 1)$ . Given this demand, and the prevailing market-clearing price, agent  $i$  chooses its optimal demand  $z_i^*(0)$ .
6. The market-clearing price ensures that aggregate demand is equal to aggregate supply.




**Fig. 1.** The market clearing price as a function of time. We observe that as resources have a higher depreciation rate, equilibrium prices increase, since less resources are available in the Grid market for trading.

7. The experiment proceeds to the next period  $t = 1$ , and the steps are repeated.
8. We use multiple (independent) random draws of  $\{y_i(t)\}$  in each period  $t$  to calculate numerical averages and expectations (i.e., we use a Monte Carlo method).


The output of the experiment is then the average expected utility of the agents (where the average is taken over the demand parameters  $\{k_i\}$ ) under Grid trading and autarky, in each period; the average volume of trade and the risk premium, in each period; and the total level of resources  $X(t)$  over time. Such results are shown in the figures 1–4, in which there are 20 agents; the risk aversion parameter either  $\rho = 2$  or  $\rho = 3$ ; and the cost parameter  $\gamma = 1$ .

Experiments showed that accounting for the net present value of the resource  $x$  instead of myopically considering the revenue at time  $t$  only, does not result to a substantially different behavior of the system. Therefore we conduct our experiments, for computational convenience, with a fixed rate of return 0.2.

The figures show some common patterns. The activity levels of the Grid market increase over time, in terms of the average expected utility of the market participants and the average volume of trade. Of course, in any particular period, the volume of trade (figure 2) may decrease, as a result of low realizations of demand. But on average, over a number of periods of trading, trade volumes increase. As agents accumulate resources, the aggregate supply in the market increases: see figure 3. As a result, the market-clearing price is generally decreasing: figure 1. (Again, in any given period, the market-clearing price may increase, as a result of high realizations of demand.) As the resource stock increases over time, it acts as a buffer to smooth out fluctuations that arise due to random demand realizations. Hence the variability of the market-clearing price declines over time; consequently, the average risk premium also falls: see figure 4. Changing the



**Fig. 2.** The average volume of trade as a function of time. We observe that as resources have a higher depreciation rate, the volume of trade decreases. This may be due to the higher equilibrium prices, see Fig. 1.



**Fig. 3.** The average resource level per agent as a function of time

depreciation coefficient  $1 - \beta$  also introduces changes in the market-clearing price and the average volume of trade, as it can be seen in figures 1 and 2 respectively. More specifically, when the resources are less durable, ie. the parameter  $\beta$  decreases, the market-clearing price increases, since less resources are available in the Grid market, see figure 1 . The increase of the equilibrium prices results to a lower value for the volume of trade, as can be seen in figure 2.




Fig. 4. The average risk premium as a function of time

## 6 Market Failure

The analysis so far has demonstrated that the operation of a Grid spot market can improve the expected utility of market participants, relative to autarky (no Grid trading). Further, since there is a positive risk premium that risk-averse agents attach to variability in the market-clearing price, there is scope in the market for the use of forward contracts. The risk premium decreases with the size of the market, and as the aggregate supply in the market increases. Nevertheless, it is always positive. Furthermore, we have assumed that all agents have the same degree of risk aversion. In practice, agents will differ in their risk preferences. Forward contracts will be particularly attractive to more risk-averse agents; and will be written by less risk-averse agents.

The fact that Grid trading increases agents' utilities is not surprising, given the general results from economics about the efficiency of competitive markets. Note, however, that the solution that we have analysed does not achieve full efficiency, since we have assumed that agents behave myopically in each period, ignoring the consequences in future periods of current choices.

The main set of assumptions underlying the analysis are those that make up a competitive market:

1. Each agent ignores the effects of its own choices on the market-clearing price (even when there is a very small number of agents in the market).
2. There is no asymmetric information in the market: for example, all agents are equally well-informed about the characteristics of Grid resources.
3. There are no external effects: an agent's choices of resource level and demand have no direct effects on other agents (other than those that operate through the market-clearing price).

Of these three assumptions, assumption 2 is likely to be the most critical for the existence and operation of a Grid market. Market power—the ability of agents to influence the market price—will affect the volume of trade and utility levels associated with the market. But it is very unlikely to cause a complete collapse of the market. Similarly, externalities may limit the scale of the market, but are unlikely to cause collapse. But if agents are poorly informed about the quality of resources traded in the market, then it may be that the market fails to operate at all.

In order to assess the scope for this market failure, we study the simplified case, when  $\beta = 0$  and  $\rho = 1$ , where

$$\begin{aligned} c(x) &= \gamma x^2, & \gamma > 0, \\ u(z) &= \ln(z); \end{aligned}$$

the resource is entirely non-durable; and each agent's type  $y_i$  is drawn from the uniform distribution on  $[0, 1]$  independently each period. Now suppose that an agent knows for certain that its own resource is reliable and can be used with probability 1 to service its demand. If an agent has to buy resources in the Grid, then it perceives that there is some probability  $\phi$  that the Grid resources fail to satisfy its demand. In the event that this happens, the agent has to pay the cost of buying the resource in the market, and bears a utility cost of  $\lambda$ . The agent's expected utility is then

$$u_i = \begin{cases} y_i \ln(z) - q(z - x_i) & z \leq x_i, \\ y_i((1 - \phi) \ln(z) + \phi\lambda) - q(z - x_i) & z > x_i. \end{cases} \quad (39)$$

Agent  $i$ 's optimal demand in this case is

$$z_i^* = \begin{cases} \frac{y_i}{q} & y_i \leq qx_i, \\ x_i & qx_i < y_i \leq \frac{qx_i}{1-\phi}, \\ (1 - \phi) \frac{y_i}{q} & \frac{qx_i}{1-\phi} < y_i. \end{cases} \quad (40)$$

The expected market-clearing price is


$$\mathbb{E}[q] = \sqrt{2\gamma \frac{(1 - \phi)(1 - \sqrt{1 - \phi})}{\phi}}. \quad (41)$$

The dependence of the expected market-clearing price on the failure parameter  $\phi$  is shown in figure 5. As the figure shows, the market-clearing price is higher when the failure probability is lower. In the limit, when any Grid resource fails with probability 1, the market-clearing price is 0, and the market collapses for sure.


Consequently, each agent buys

$$x_i^* = \sqrt{\frac{(1 - \phi)(1 - \sqrt{1 - \phi})}{2\gamma\phi}} \quad (42)$$

of the resource.



**Fig. 5.** The market-clearing price against the failure parameter  $\phi$



**Fig. 6.** Expected utilities against the failure parameter  $\phi$

These facts can then be used, along with the realizations of the demand parameters  $\{y_i\}$ , to determine numerically the equilibrium outcome, as a function of the failure parameter  $\phi$ . It can be seen in figure 6, that if  $\phi$  is sufficiently low agents will trade in the Grid market. Once Grid resources become sufficiently unreliable, autarky is preferred.



Obviously, a key parameter determining the ‘crossing point’—the critical value of  $\phi$  at which the Grid market collapse—is the utility cost  $\lambda$ . We have modelled this in a very simple way here: as a constant that is the same for all agents. If this is an important feature of an emerging Grid market, then this aspect can be modelled further. For example, the utility cost could be allowed to be different for different agents. It could also be related to agents’ realized demand types  $\mathbf{y}$ .

## 7 Conclusions

In this paper we analyzed the economics of a competitive market for Grid services. Namely, we provided an analytical model, as well as simulation results regarding the critical parameters of such a model. Accounting for different levels of risk aversion, we studied the behavior of the market clearing price, the volume of trade, the resource level per agent and the average risk premium. In all those cases we assumed that the resources do not necessarily vanish after the end of each period, but last longer and each agent accumulates a stock of such resources. We observed that in general the market clearing price decreases and the aggregate supply in the market increases as agents accumulate resources. We also analyzed the situation where Grid resources are not reliable and the market fails to satisfy the demand and observed that, unless Grid resources are highly unreliable, the agents still prefer to trade in the Grid market than to choose autarky.

## References

1. Chamberlin, E.H.: An Experimental Imperfect Market. *Journal of Political Economy* 56(2), 95–108 (1948)
2. Cripps, M.W., Swinkels, J.M.: Efficiency of Large Double Auctions. *Econometrica* 74(1), 47–92 (2006)
3. Gode, D.K., Sunder, S.: Allocative Efficiency of Markets with Zero Intelligence (ZI) Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy* CI, 119–137 (1993)
4. Gode, D.K., Sunder, S.: What Makes Markets Allocationally Efficient? *Quarterly Journal of Economics* 112(2), 603–630 (1997)
5. Rustichini, A., Satterthwaite, M.A., Williams, S.R.: Convergence to Efficiency in a Simple Market with Incomplete Information. *Econometrica* 62(1), 1041–1063 (1994)
6. Satterthwaite, M.A., Williams, S.R.: The optimality of a simple market mechanism. *Econometrica* 70(5), 1841–1863 (2002)
7. Cliff, D.: Minimal-intelligence agents for bargaining behaviours in market-based environments, Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England (1997)
8. Preist, C., van Tol, M.: Adaptive agents in a persistent shout double auction. In: *Proceedings of the 1st International Conference on the Internet, Computing and Economics*, pp. 11–18. ACM Press, New York (1998)

9. Parsons, S., Marcinkiewicz, M., Niu, J., Phelps, S.: Everything you wanted to know about double auctions but were afraid to (bid or) ask, Technical Report, Department of Computer & Information Science, Brooklyn College, City University of New York (2005)
10. Smith, V.L.: An Experimental Study of Competitive Market Behavior. *Journal of Political Economy* LXX, 111–137 (1962)
11. Meinel, T., Neumann, D.: A Real Options Model for Risk Hedging in Grid Computing Scenarios. In: *Proceedings of the 42nd Hawaii International Conference on System Sciences*, pp. 1–10 (2009)