Published in: Physica A, 2004, Vol. 340(4), 668-677

The importance of timescales: Simple models for economic markets

Kai Nagel, [a,b,1] Martin Shubik, [c] and Martin Strauss [a]

[a] *Inst. f. Computational Science, ETH Zürich, Zürich, Switzerland*

[b] *Transport Systems Planning and Transport Telematics, Technical University of Berlin, Germany*

[c] *Cowles Foundation for Research in Economics, Yale University, New Haven, Connecticut*

Abstract

This paper considers a simple model of an economy. The economy consists of agents. Each agent produces exactly one good. The good is sold on the market and the agent uses the resulting money to buy many other goods. All agents have the goal to maximize their own utility, which consists of a positive contribution from consumption, and a negative contribution from work. The problem for the agent thus is to balance work and consumption. In contrast to many other economic models, this model prescribes the process in all completeness. The paper looks both at analytical solutions and at simulation results. A particularly important results is that a well-defined market only emerges when prices adapt on a much slower time scale than consumption. This makes clear that a functioning market does not just emerge by itself.

Key words: Multi-agent simulation, economic markets, Nash Equilibrium
PACS: 89., 89.65.-s, 89.65.Gh, 07.05.Tp

1 Introduction

Economic General Equilibrium Theory is static. Since there is no dynamics, it is not specified how the system reaches the equilibrium state, or how it behaves once it is away from it. A consequence is that nobody actively sets prices – in that theory, they are the outcome of the common thought process, which is not realistic with respect to the real world.

[1] Corresponding author; nagel@kainagel.org

Game theory offers a framework which allows to couple the setting of prices to agents, e.g. in the Cournot or Bertrand games (e.g. [1], Chap. 6). However, again, game theory is static: It only searches for a situation in which nobody wants to change, but does not consider how the system reaches that state. Evolutionary game theory (e.g. [2]), then, is a framework which makes game theory dynamic and it is possible to evaluate the dynamic behavior of the system. For example, Nash equilibria (NE) are fixpoints of the dynamic evolution of evolutionary game theory. Such fixpoints can be stable (attractive) or unstable (repulsive); if such a fixpoint is unstable, it means that the corresponding NE is essentially irrelevant under the given adaptive rule, and possibly irrelevant under *any* adaptive rule. It is also possible to have more complicated limit sets such as cyclic or chaotic ones. If they are attractive, they have a basin of attraction, meaning that if the system ever gets into the basin of attraction, then it will go to the limit set and remain there. This is, however, only true for deterministic systems; in stochastic systems, the noise can drive the system out of the basin of attraction into another one.

One can see from this that convergence to a NE is only a subset of the behavior that an evolutionary game can display. The question thus becomes if one can get some insight into what dynamic behavior is indeed plausible. For this, one needs to construct dynamic models. In this paper, we construct a dynamic model of markets. Our thinking, and thus our model, is based on the assumption that there are *individual actors* in the economic world, and that we want to be able to identify these individual actors in our simulation model. Our model is constructed around a simple production-consumption economy, i.e. each actor produces a single good, and trades that for many other goods from many other actors in order to be able to consume a diversity of goods. Our approach is to define the rules of the system, and then to analyze the resulting dynamics. Our longterm working hypothesis is, consistent with our analysis above, that there will be some cases in which the dynamics converges to results which are similar to results of General Equilibrium Theory, while there will be other cases in which the dynamics displays a completely different behavior. In this paper, we will concentrate on the first aspect, i.e. the analysis in how far a simple model is consistent with conventional static equilibrium theory.

2 The basic model

Our model derives from work by Bak, Norrelykke, and Shubik (BNS) [3]. Yet, the relation turns out to be somewhat complicated. Its discussion is therefore postponed to Sec. 6. The basic model consists of N agents, each of them producing exactly one good. Each agent is characterized by a utility function

$$U_i = -\frac{q_i^2}{2} + 2 \sum_{j \neq i} \sqrt{x_{ij}} \; . \tag{1}$$

q_i is the amount of work of agent i; x_{ij} is the amount of good j (bought from agent j) which agent i consumes. The marginal disutility of work, $-\partial U_i/\partial q_i$, is increasing, meaning that working twice as much is more than twice as unpleasant. In contrast, the marginal utility of consumption, $\partial U_i/\partial x_i$, is decreasing, meaning that consuming twice as much is less than twice as pleasant.The model is initialized by each agent i receiving an initial amount of money $M_{i,0}$, setting an initial price $p_{i,0}$, and setting, for each good j, an initial preference $\hat{x}_{ij,0}$. Each time step then consists of the following sub-steps:

(1) Agents order goods according to their preferences vector $\hat{x}_{ij,t}$ such that exactly all money is spent. This means a re-scaling of the $\hat{x}_{ij,t}$, i.e. $x_{ij,t} = \hat{x}_{ij,t}\, M_{i,t-1}/(\sum_{j \neq i} \hat{x}_{ij,t}\, p_{j,t})$. $M_{i,t-1}$ is the amount of money left over from the last time step, see below. Note that with this equation $\sum_{j \neq i} x_{ij,t}\, p_{j,t} = M_{i,t-1}$, which is the *budget constraint*. Agents to not buy from themselves, that is, $\forall t : x_{ii,t} = 0$. This simplifies some of the mathematics.
(2) Goods are produced to order, and sold at the prices previously indicated. That is, each agent produces $q_{i,t} = \sum_{j \neq i} x_{ji,t}$ and receives the amount of money of $M_{i,t} = p_{i,t}\, q_{i,t}$. Since, according to Item (1), each agent has spent all her money with her orders, the result is indeed the amount of money with which they go into the next time step.
(3) Utilities $U_{i,t}$ are calculated.
(4) Preferences are adapted: Every T_x time steps, agents adapt (see below) their x_{ij} to maximize utility given prices. Usually, $T_x = 1$.
(5) Prices are adapted: Every T_p time steps, agents adapt (see below) their p_i to maximize utility. Adaptation of prices happens on a much slower time scale than adaptation of preferences, i.e. $T_p \gg T_x$. Usually in this paper, $T_p = 10^5$.

The above model description does not specify how adaptation takes place. In fact, many different adaptation schemes seem to work. Throughout this paper, we use "trial-and-error" adaptation. This means that, from time to time, agents try different strategies. If the performance (= utility) of a changed strategy turns out to be better than the performance of the previous strategy, then they will stick with the new strategy. For each agent i, a strategy consists of the $(x_{ij})_j$ and of p_i.[2] Trial-and-error adaptation for the $\hat{x}_{ij}$ of agent i works as follows:

(1) After T_x time steps with the "normal" strategy, the agent picks one of its $(\hat{x}_{ij})_j$ randomly. It remembers its old value, plus the corresponding utility, $\tilde{U}_i$, and then obtains a new "trial" value for $\hat{x}_{ij}$.
(2) The agent operates with the new trial value for T_x time steps and then looks at the resulting utility. If this new utility is larger than the original utility, it sticks with the new $\hat{x}_{ij}$, otherwise it return to the old one.
Trial values are obtained via mutation: The old x_{ij} is changed by a small random amount.

[2] Double indices of type $(X_{ij})_j$ refer to the vector $(X_{i1}, ..., X_{iN})$.

The above explanation was in terms of preferences $\hat{x}_{ij}$. Adaptation for p_i happens accordingly, the only three differences being that (1) there is only one value per agent, (2) the trial time T_x is replaced by T_p, and (3) agents enter into trial mode only with probability f_p where we use $f_p = 0.1$. The latter is done so that only a small fraction of the system is in trial mode since otherwise the agents that are in trial mode are not faced with an approximation of the "normal" system.

More parameters are necessary to fully describe adaptation; for example, a mutation for $\hat{x}_{ij}$ is, with probability $1/2$, either $\hat{x}_{ij,new} = \hat{x}_{ij,old} \times f_{x,rand}$ or $\hat{x}_{ij,new} = \hat{x}_{ij,old}/f_{x,rand}$ where $f_{x,rand}$ is a random number uniformly distributed between 1 and 1.01. For price adaption we mutate accordingly with a factor $f_{p,rand}$ which is a random number between 1 and 1.1. Our simulations indicate that our results are robust as long as the mutations remain small, and as long as T_p remains long enough so that the adaptation of the $\hat{x}_{ij}$ can reach a steady state even with small mutations (see Fig. 1 plus related text in Sec. 4, and Sec. 5).

3 Analytical approximation

Finding relations between dynamical models and "static" game theory is often not straightforward. In particular, it is critical that one is clear about what are the strategies, and what is known to each agent. In our model, each agent i's strategy consists of the vector of the $(\hat{x}_{ij})_j$ and of the price p_i. Since prices are kept constant for long periods of time, one can assume that agents find $(\hat{x}_{ij})_j$ such that utility is maximized when prices are given. On a second level, agents optimize prices *against that reaction of their fellow consumers*. That is, prices are optimized against the fact that consumers display optimal reaction to prices. In consequence, the above model corresponds to a solution of a two-step optimization problem: (1) First, for given prices $(p_j)_j$ and money $M_{i,t-1}$ find an optimal allocation of the $\hat{x}_{ij}$ such that utility is maximized. This is done in Sec. 3.1. (2) Second, given that every agent knows every other agent's reaction to price changes, find an optimal price. See Sec. 3.2.

3.1 Optimal consumption For given prices $(p_j)_j$ and amount of money $M_i = M_{i,t-1}$, the optimization problem that each agent i faces is to maximize Eq. (1) under the constraint that the amount of money $M_{i,t-1}$ is exactly spent. This is as usual achieved by first adding the constraint via a Lagrangian multiplier,

$$U_i := -q_i^2/2 + 2\sum_{j\neq i}\sqrt{x_{ij}} + \lambda_i\left(M_{i,t-1} - \sum p_j x_{ij}\right), \qquad (2)$$

and then setting first derivatives with respect to the decision variables equal to zero.

Now consider a fixed agent i. Since $q_i = \sum_{j \neq i} x_{ji}$ and $M_{i,t-1}$ do not depend on x_{ij}, the derivative with respect to x_{ij} is simple. Together with the budget constraint this results in

$$x_{ij} = \frac{1}{\lambda_i^2 \, p_j^2} \; , \; \frac{1}{\lambda_i^2} = \frac{M_{i,t-1}}{\sum_{k \neq i} \frac{1}{p_k}} \; , \text{ and } x_{ij} = \frac{M_{i,t-1}}{p_j^2 \sum_{k \neq i} \frac{1}{p_k}} \; . \tag{3}$$

3.2 Optimal price Optimal prices are now computed by taking Eq. (1) and taking the first derivative with respect to p_i. Note that it is no longer necessary to include the budget constraint, since this is automatically fulfilled as long as all agents follow Eqs. (3).

However, one is faced with a new difficulty. Changing a price also changes the distribution of money. That is, in order to evaluate the effect of a price change by agent i, one needs not only compute the reactions of the buyers, but also the resulting re-distribution of money, the resulting 2nd-order reaction of the buyers, etc. Thus, one needs to find the *steady state* solution of the system in reaction to a price change and only then take derivatives.

We were in fact able to find a closed form formulation of that steady state solution [4], but it is beyond the scope of this paper. Important insight is already gained from the case $N \to \infty$. In this case, the change of the steady state value of M_j ($j \neq i$) goes to zero so that it can be neglected. Similarly, the influence of p_i on $\sum_{k \neq j} \frac{1}{p_k}$ decreases as $1/N$. Together, this also decouples λ_j from p_i. With this, one can express everything using either p_i or some terms that explicitly do not depend on p_i. After entering these conversions into Eq. (1), taking the first derivative w.r.t. p_i and setting it to zero, one makes the assumption of a homogeneous system (Nash Equilibrium) by setting $p_i = p$, $\sum_{k \neq i} \frac{1}{p_k} \to \frac{N-1}{p}$, etc. The final results are

$$p \approx \frac{2^{2/3} \, M}{(N-1)^{1/3}} \; , \; x \approx \frac{2^{-2/3}}{(N-1)^{2/3}} \; , q \approx 2^{-2/3} \, (N-1)^{1/3} \; , \tag{4}$$

and $U \approx (2^{2/3} - 2^{-7/3})(N-1)^{2/3}$ (see Tab. 1). That is:

- Prices are proportional to the amount of money M in the system, while the “physical” variables x, q, and U do not depend on money or price levels at all.
- Prices decrease with increasing N. This is plausible because larger N means that it is easier to substitute one good by another, increasing competition.
- The consumption of each individual good decreases with N, as it should. However, the sum of all consumption, $\sum_j x_{ij}$, *in*creases with N. This is a consequence of the fact that the slope of the utility function goes to infinity as $x \searrow 0$, meaning that agents always want “a little bit of everything”, no matter how large N is. This also implies that $q \to \infty$ when $N \to \infty$, which is physically implausible. Other utility functions do not have this effect but cause other “problems” [4,5].

N	p	x	q	U
10	0.76	0.15	1.31	6.01
50	0.43	0.05	2.31	18.60
1000	0.16	0.01	6.30	138.80

Table 1
Some numbers for the "large N" approximation (with $M = 1$).

Note that we are looking for the *Nash* equilibrium and not the system optimum (SO): All agents together could agree upon a price p_{SO} which would maximize the utility of all of them at the same time. This price can be calculated by inserting the "homogeneous budget constraint" $x = \frac{M/p}{N-1}$ into the utility function and maximize the resulting function w.r.t. p. This results in $p_{SO} = M/(N-1)^{1/3}$, which is lower than the NE price. Reinserting p_{SO} into U give them the SO utility $U_{SO} = \frac{3}{2}(N-1)^{2/3}$, which is larger than the NE utility. The situation is the same as in the "prisoners dilemma" where both players would be better off if they played together, but because every player tries to maximize his own profit, both are worse off in the end.

The case $N = 2$ can also be solved exactly but is somewhat peculiar (see Sec. 6). For all other N, we were unable to find a simple solution.

4 Simulation results

First, we check if the simulation reacts to price differentials. For that, we set one agent at a different price than all others (Fig. 1). One sees that agents slowly shift their consumption toward the good with the lower price, visible in the graphs as a larger production for that agent. One also sees that this results in a higher utility for that agent, meaning that in this case average prices are too high and an individual has an advantage when charging a lower price. This effect is what will drive the adaptation of prices further below.

The simulation is confirmed by the analytical result. The "large N" Nash Equilibrium price, Eq. (4), is 0.53, since $M = 1$ and $N = 10$. This indicates that a base price of 3, as in Fig. 1, is too large. One also sees that it takes about 30 000 timesteps until the aggregate consumption has relaxed to the value corresponding to the prices. For $N = 50$ agents, relaxation takes about 70 000 time steps (not shown). Including some "safety margin", we will use $T_p = 100\,000$ in the following simulations.

We now allow both the preferences $\hat{x}_{ij}$ and the prices p_i to adapt. As explained above, the adaptation of the $\hat{x}_{ij}$ happens on a fast time scale, while adaptation of the p_i happens on a slow time scale. In practice, this means that in every second

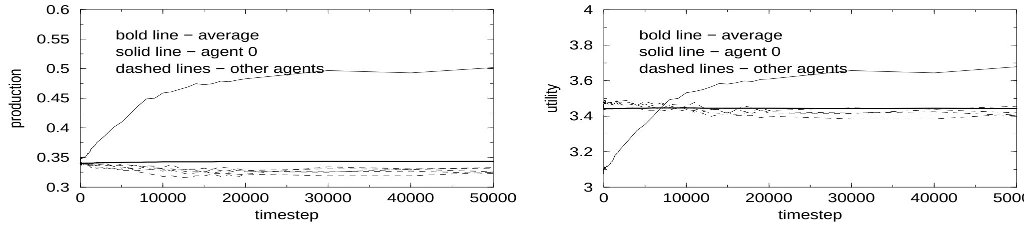

Fig. 1. Simulation where the preferences $\hat{x}_{ij}$ adapt via trial-and-error to price differentials. Initially (at negative times, not shown), the simulation is run such that all agents find optimal $\hat{x}_{ij}$ against a uniform price of 2.4. Then, at $t = 0$, the price of agent 0 is set to 2.4, and the reaction of the system is displayed. LEFT: Production. RIGHT: Utility. The fat line describes the average behavior, the single thin line the behavior of the one agent with the lower price, all other thin lines are examples of other agents. $N = 10$, $M = 1$.

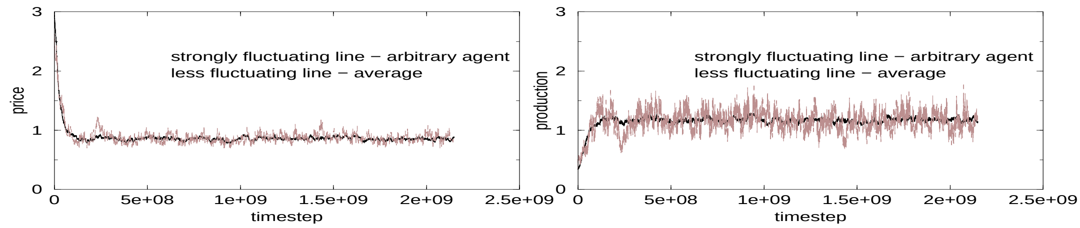

Fig. 2. Simulation where both the preferences $\hat{x}_{ij}$ and prices p_i adapt via trial-and-error, system of 10 agents. Initial prices are at 3. LEFT: Prices. RIGHT: Production.

time step every agent tries a new $\hat{x}_{ij}$. In contrast, new prices are tried out only every $T_p = 10^5$ time steps, and they are left in place for the same duration before they are evaluated.

As a result, this simulation takes much longer, note the time scale in Fig. 2. One also sees how the system relaxes toward a price of about 0.87, a production q of about 1.2, and a utility U of about 5.7. These values are similar to the "large N" values given in Tab. 1 ($N = 10$). This similarity gets better with larger N. Larger sizes than $N \approx 50$, however, cannot be tested because they take several days of computer time. In addition, the analytical calculation does not take into account fluctuations, which may also explain part of the difference between simulation and analytical result.

The analytical solution implies and the simulation results confirm that, instead of slow adaptation, agents could set their consumption directly via $x_{ij,t} = M_{i,t-1}/p_{j,t}^2 \sum_{j \neq i} \frac{1}{p_{j,t}}$. So a possibility is to use this directly in the simulations. The motivation behind this is that one could argue that agents should capable of doing "local" optimizations, in the sense that they should be able to compute a best reply against a fixed environment, here given by known prices. This would correspond to an agent who knows everything about him-/herself and is able to compute a corresponding mathematically optimal reaction to a fixed environment. Such an agent would, however, *not* be capable of predicting other agents' reactions to that best reply. Apart from different fluctuations simulation results look the same as before. This saves enormously in

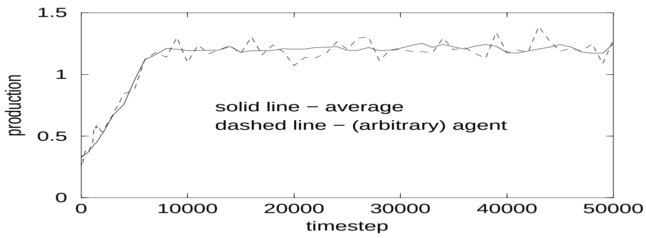

Fig. 3. Production in a simulation where preferences are calculated optimally while prices adapt via trial-and-error. Initial prices are 3. One sees that the system relaxes toward the same value as in Fig. 2. Note the much shorter time scale on which that happens.

terms of computer time, since we can now do price adaption with $T_p = 10$ instead of $T_p = 10^5$ time steps. Fig. 3 shows the result for the same parameters as those used for Fig. 1. But much larger systems can now be simulated (not shown).

5 Separation of time scale and stability

The stability of the above model hinges critically on the fact that adaptation of prices is much slower than adaptation of preferences. In effect, price adaptation is made so slow that the adaptation of preferences and the redistribution of money has always completed before the performance of a price change is evaluated. As shown earlier by direct simulation, one needs to wait about 50000 timesteps until preferences and money have adapted to a new price situation. This corresponds to a complete separation of time scales.

If preferences are *calculated*, as in Sec. 3.2, this adaption time is greatly reduced to about 5 timesteps [3] and it is therefore enough to wait only 10 timesteps between two consecutive price changes. It is very important to note that changes in price must still be made on a slower time scale than adaption of preferences as even if the agents always buy the optimal amount still some time is needed until a steady state distribution of the money is established. However, the main aspect of the "separation of time scales" is now achieved by the agent itself, which reaches the best reply instantaneously.

In terms of stability, it is *not* possible to violate this rule: Making price adaptation faster relative to consumption adaptation leads to continuously increasing prices, which then just go to infinity. The reason for this is simple: A price increase will initially increase the corresponding income; the reduction in sales sets in only later. If thus the effect of the price increase is evaluated very soon, it will seem like a utility-increasing move.

[3] This figure is correct only if N is not smaller than ≈ 10. For smaller numbers of agents, the redistribution of money takes longer.

This effect is also reflected in the analytical solution: One has to solve for the x_{ij} first, under the assumption that prices are given. Only then does one optimize prices based on the consumption behavior. Any attempt to solve this in a different sequence leads to $p_i \to \infty$ or does not work at all.

This result, albeit plausible once one understands the model and its dynamics, came as a surprise to some of the authors, because the implication is that the market (in the sense of competitive prices) does not play under all circumstances, but only if the time scales of adaptation are separated in the correct way. One could argue that consumers would develop strategies to cope with quickly changing prices. Yet, certainly for simple models and possibly for the real world, some regulation may be necessary to achieve competition. The issue is related to "sequential games", "subgame perfection", and "Stackelberg games" in game theory, and also shows up in transport simulations [6].

6 The BNS model

Our results allow some outlook on the BNS model [3]. In that paper, agents are located on a ring, and every agent buys from the left and sells to the right.

The solution with respect to consumption remains the same as in Eq. (3). However, since $\sum_{k \neq i} \frac{1}{p_k} = p_j$, one also obtains $x_{ij} = \frac{M_{i,t-1}}{p_j}$. We will write $x_{ij} \equiv x_i$, since there is only one good that each agent consumes.

With respect to optimal price, first note that in the steady state, $M_i = M_j = M$, since otherwise one agent accumulates money. With this, $U_i = -\frac{1}{2}\, q_i^2 + 2\,\sqrt{x_i} = -\frac{1}{2}\, M^2/p_i^2 + 2\, M/p_j$. Taking the first derivative w.r.t. p_i and setting it to zero leads to the unexpected result $p_i^{-3} = 0$ or $p_i = \infty$. This gets explained by noticing that the amount of money that the agent i spends is always $x_i\, p_j = \frac{M}{p_j}\, p_j = M$ and is thus *totally independent* from the price that j sets. On another level, this is clear since our model prescribes that agents have to spend all their money in every time step. So if there is only one person to buy from, that person always gets all the money, and for that reason that other person will raise the price as high as possible, since that reduces the amount of work (as long as they do it unilaterally, as the NE solution prescribes).

This means that a straightforward interpretation of the BNS model in terms of the present paper is not possible. One way to explain the differences is that in the present paper, the independent strategic variables are given via the vector $(p_i, x_{i1}, ..., x_{iN})$. In contrast, in the BNS paper the only independent strategic variable is λ_i, while p_i and x_{ij} are given via $x_{ij} = 1/(\lambda_i^2\, p_j^2)$ and $p_i = 2^{1/3}\, \lambda_i^{-1/3}\, p_i^{2/3}\, q_{i,t-1}^{1/3}$, where $q_{i,t-1}$ is the demand from the *previous* time step. The interest of the BNS paper is in the dynamics of the λ_i, which are re-calculated in every time step so that the budget is approximately balanced, while locally optimal strategies are maintained. In consequence, in spite of a lot of apparent similarity between BNS and the present paper, the actual approach and interpretation are rather different.

7 Summary

We have presented a simple dynamic model of a market. Certain versions of the model can be solved analytically. Simulation offers the possibility to go beyond the analytically solvable cases. In both cases, for stable solutions it is crucial to select the dynamics correctly. In the model of this paper, price adaptation has to happen on a much slower time scale than consumption adaptation, otherwise prices go to infinity. This is intuitively plausible; nevertheless, it needs to be taken into account both when building simulations models and possibly when regulating the real world.

Economic dynamics uses the institutions of our society. This dynamics is far from ideal process-free economic equilibrium theorizing. The size of the time lags in the adjustment processes matters and there appears to be a trade-off between the speed of adjustment and stability. This raises both empirical questions as well as questions in theory in characterizing both price and quantity adjustment processes.

Acknowledgments Part of the work on this paper was done during a visit of KN at Imperial College, London, in spring 2001. KN also acknowledges the hospitality of the Cowles Foundation during summers 2001 and 2002. KN and MS thank Maya Paczuski and Per Bak for extremely helpful discussions.

References

[1] H. R. Varian, Microeconomic Analysis, W.W. Norton & Company, 1992.

[2] J. Hofbauer, K. Sigmund, Evolutionary games and replicator dynamics, Cambridge University Press, 1998.

[3] P. Bak, S. Norrelykke, M. Shubik, Dynamics of money, Phys. Rev. E 60 (3) (1999) 2528–2532.

[4] M. Strauss, Dynamic market simulations, Diploma thesis, ETH Zurich, Switzerland, see e-collection.ethbib.ethz.ch (2001).

[5] K. Nagel, M. Strauss, M. Shubik (in preparation).

[6] H. v. Zuylen, H. Taale, Urban networks with ring roads: A two-level, three player game, Paper 04-1659, Transportation Research Board Annual Meeting, Washington, D.C. (2004).