Stability conditions for a dynamic market simulation model

–draft for discussion purposes–

Kai Nagel (ETH Zürich) and Martin Shubik (Cowles Foundation)

April 19, 2001

1 Introduction

General equilibrium theory is static. Since there is no dynamics, not enough is known how the system should reach the equilibrium state, or how it behaves once it is away from it.

A consequence of this is that nobody actively sets prices – they are the outcome of the common thought process. Game theory offers a framework which allows to couple the setting of prices to agents, e.g. in the Cournot or Bertrand games.

However, again, game theory is static. Evolutionary game theory (e.g. [1]), then, is a framework which makes game theory dynamic and it is possible to evaluate the dynamical behavior of the system. For example, Nash equilibria are fixpoints of the dynamic evolution of evolutionary game theory. Such fixpoints can be stable or unstable, etc.; and it is also possible to have more complicated limit sets such as cyclic or chaotic ones. If those are stable, 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.

In this paper, we construct a dynamic model of markets. We have agents which set consumption and prices, and we follow the dynamics. We think that the situation is somewhat similar to the situation in Statistical Physics some 100 years ago, where the macroscopic thermodynamic equations were around, but neither was their microscopic foundation nor their extension into non-equilibrium. In consequence, we search for a theory where, at least in some cases, the general equilibrium result will emerge when transients have died out, but where the formulation of the theory also allows for systematic investigations into non-equilibrium situations.

The recent development of "econophysics" is making progress into that direction. What, in our view, is missing is a more direct relation to markets, which includes identifiable goods, prices, consumption, etc. In this paper, we will introduce a family of models which should be capable to serve as a starting point for such investigations.

2 The basic model

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 conversion of work into (dis)utility is convex, meaning that working twice as much is more than twice as unpleasant. In contrast, consumption is concave, meaning that consuming twice as much is less than twice as pleasant.[1]

The model works as follows.

Initialization:

- Each agent i receives an initial amount of money $M_{i,0}$.
- Each agent i sets an initial price $p_{i,0}$.
- Each agent i sets, for each good j, an initial allocation $\hat{x}_{ij,0}$. These allocations set the ratios between the goods, that is, we always use $\sum_j \hat{x}_{ij} = 1$.

One time step:

1. Agents order goods according to their preferences vector $\hat{x}_{ij}$ such that exactly all money is spent. This means a multiplication of the $\hat{x}_{ij}$ by $M_{i,t-1} / \sum_{j \neq i} p_j \, \hat{x}_{ij}$, i.e.
$$x_{ij} = \hat{x}_{ij} \, \frac{M_{i,t-1}}{\sum_{j \neq i} \hat{x}_{ij} \, p_j} \; . \tag{3}$$
$M_{i,t-1}$ is the amount of money left over from the last time step, see below.
2. Goods are produced to order, and sold at the prices previously indicated. That is, each agent produces
$$q_i = \sum_{j \neq i} x_{ji} \tag{4}$$

[1] In this paper, Eq. 1 will be used. However, it can be seen as a specific version of

$$U_i = -A_i \, g(q_i) + \sum_{j \neq i} a_{ij} h(x_{ij}) \; , \tag{2}$$

where $g(X)$ is a convex function and $h(X)$ a concave function. Many of our results should also hold for this more general version of the utility function.

and receives the amount of

$$M_{i,t} = p_i\, q_i \ . \tag{5}$$

Since, according to Eq. 3 each agent has spent all her money with her orders, the result of Eq. 5 is indeed the amount of money with which they go into the next time step.

3. Utilities U_i are calculated.

4. Adaptation of preferences: Every T_x time steps, agents adapt (see Sec. 3) their x_{ij} to maximize utility given prices. Usually, $T_x = 1$.

5. Adaptation of prices: Every T_p time steps, agents adapt (see Sec. 3) 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, $T_p = 10^5$.

3 Adaptation

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. Trial-and-error adaptation for the $\hat{x}_{ij}$ works as follows:

1. After T_x time steps with the "normal" strategy, pick one of your $\hat{x}_{ij}$ randomly. Remember its old value, plus the corresponding utility, $\tilde{U}_i$. Obtain a new "trial" value for $\hat{x}_{ij}$.[2]

2. Operate with the new trial value for T_x time steps and then look at the resulting utility. If this new utility is larger than the original utility, stick with the new $\hat{x}_{ij}$, otherwise return to the old one.

Trial vialues can be obtained via two mechanisms:

- **Mutation**. The old x_{ij} is changed by a small random amount.
- **Copy**. A new value of x_{ij} is taken from another agent.

For this paper, only mutation will be used.

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 price per agent,

[2]According to the above model specification, this means that all other $\hat{x}_{ik}, k \neq j$ will be rescaled such that still $\sum x_{ij} = 1$.

(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 $f_p = 0.1$.

More parameters are necessary to fully describe adaptation; for example, the size of a mutation for $\hat{x}_{ij}$ is $\pm 0.01 \times rnd$, where rnd is a random number between zero and one. 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 preferences can complete even with small mutations.

4 Related work

The model is related to a model by Bak, Norrelyke, and Shubik (BNS) [2]. The main difference is that in the present paper, agents transparently adapt x_{ij} and p_i, which seem to be the plausible economic quantities to work with. In addition, the present paper assumes that everybody is buying from everybody, whereas the BNS model assumes that everybody is buying from their 'left" neighbor only. We will see that this has important consequences for the dynamics and its justification.

5 Analytical approximation

The above model corresponds to a solution of a two-step optimization problem:

1. First, for given prices find an optimal allocation of the $\hat{x}_{ij}$ such that utility is maximized.

2. Second, given that every agent knows every other agent's reaction to price changes, find your optimal price.

Analytical approximations to this are in the appendix. One result of the analytical calculation is that, for large N, the homogeneous solution is

$$p \approx \frac{2^{2/3}\, M}{(N-1)^{1/3}} \;, \; x \approx \frac{2^{-2/3}}{(N-1)^{2/3}} \text{ and } q \approx 2^{-2/3}\,(N-1)^{1/3} \;. \tag{6}$$

That is:

- Prices are proportional to the amount of money M in the system.

- Prices decrease with increasing N.

- Consumption does not depend on prices.

- 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 fact that the slope of the utility function goes to infinity as $x \searrow 0$.

6 Simulation results

6.1 Adaptation of preferences

First, we check if the simulation reacts to price differentials. For that, we set one agent at a different price than all others (Fig. 6.1). One sees that agents slowly shift their consumption towards to 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 in Sec. 6.2.

One also sees that it takes about 40000 time steps until the aggregate consumption has relaxed to the value corresponding to the prices. Including some safety margin, we will use $T_p = 100000$ in the following simulations.

6.2 Adaptation of both preferences and prices

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 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. 6.2. One also sees how the system relaxes towards a price of about 0.87, a production of about 1.2, and a utility of about 5.7.

For comparison, the "large N" solution from Sec. A gives, for the $N = 10$ and $M = 1$ used here,

$$p \approx 0.76 \text{ and } q \approx 1.31 \; . \tag{7}$$

It is clear that the values of the simulation will not completely agree with the values of the analytical calculation since the assumptions are different. For example, the analytical calculation is valid only for large N (compared to $N = 10$ in the simulation), and the analytical calculation does not take into account fluctuations.

7 Simulations with calculated instead of adapted consumption

The analytical solution implies and the simulation results confirm that, instead of slow adaptation, agents could set their consumption directly via

$$x_{ij} = \frac{M_i}{p_j^2 \sum_{j \neq i} \frac{1}{p_j}} \; . \tag{8}$$

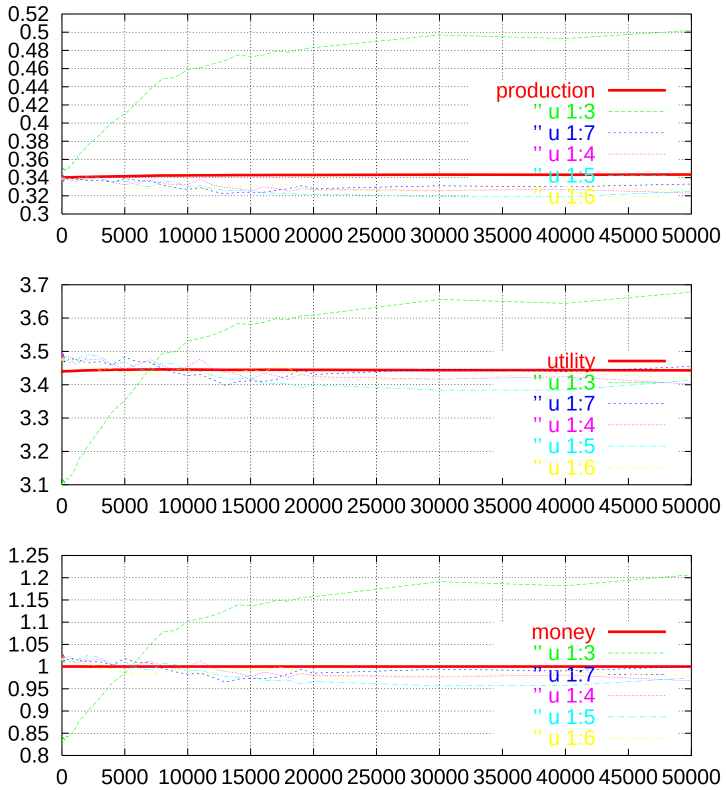

Figure 1: *Simulation where the preferences $\hat{x}_{ij}$ adapt via trial-and-error while prices are fixed. The price of one agent is fixed at 2.4, the prices of all other agents are fixed at 3. TOP: Production. CENTER: Utility. BOTTOM: Money. The fat line describes the average behavior, the single thin line the behavior of the one agent with the lower price, all other thin lines are examples of other agents. One clearly sees that preferences slowly move towards the good with the lower price. 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 in Sec. 6.2.* $N = 10$, $M = 1$.

So a possibility is to use this directly in the simulations. In fact, apart from different fluctuations simulation results look the same as before. This saves enormously in terms of computer time, since we can now do price adaption with $T_p = 10$ instead of $T_p = 10^5$ time steps.

From a conceptual point, this means that we allow agents to locally use mathematics in order to find better (i.e. locally optimal) solutions. Since we have demonstrated that the same solution can be reached via adaption, this can be seen purely as a shortcut in computation. Note however that this equivalence of adaptation and calculation is only valid as long as there is a unique minimum.

A discussion in how far such a quantitative difference may result in a qualitative difference may eventually become necessary but is beyond the scope of this paper.

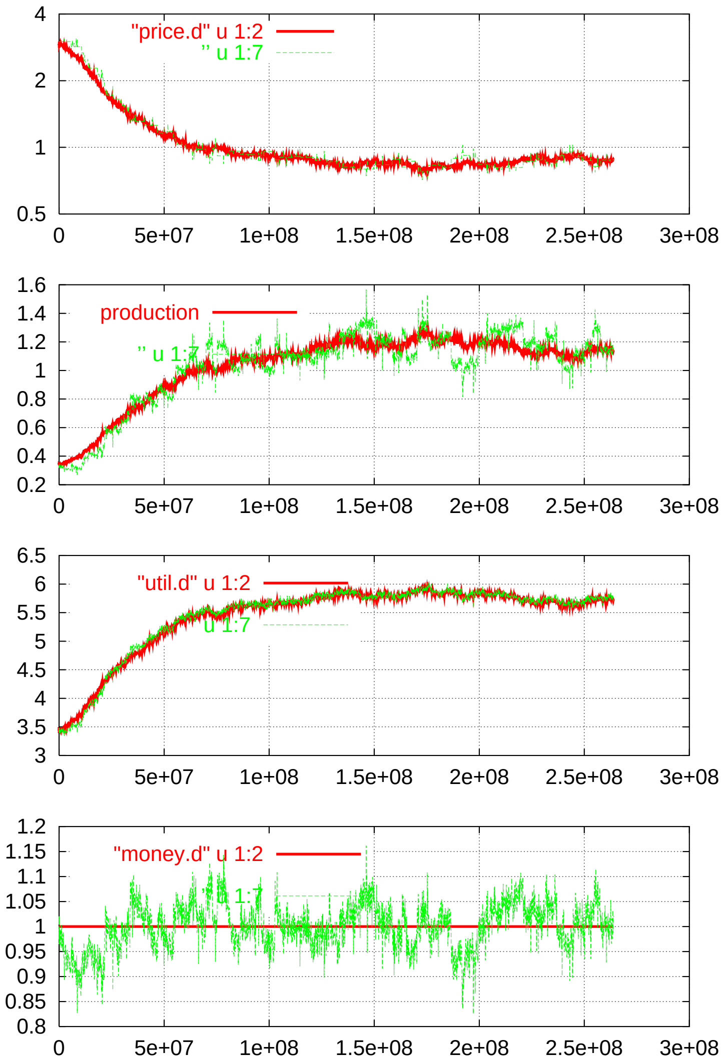

Figure 2: *Simulation where both the preferences $\hat{x}_{ij}$ and prices p_i adapt via trial-and-error. TOP: Prices. SECOND: Production. THIRD: Utility. BOTTOM: Money. The fat line gives the average behavior of the system; the thin line gives the behavior of a single (arbitrary) agent. One sees how the system relaxes towards a price of about* 0.87*, a production of about 1.15, and a utility of about 5.7. The bottom figure shows that the average per capita amount of money in the system remains at one (as it should) but that the individual amount of money fluctuates considerably.*

8 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 preferences adaptation has always completed before the performance of a price change is evaluated.

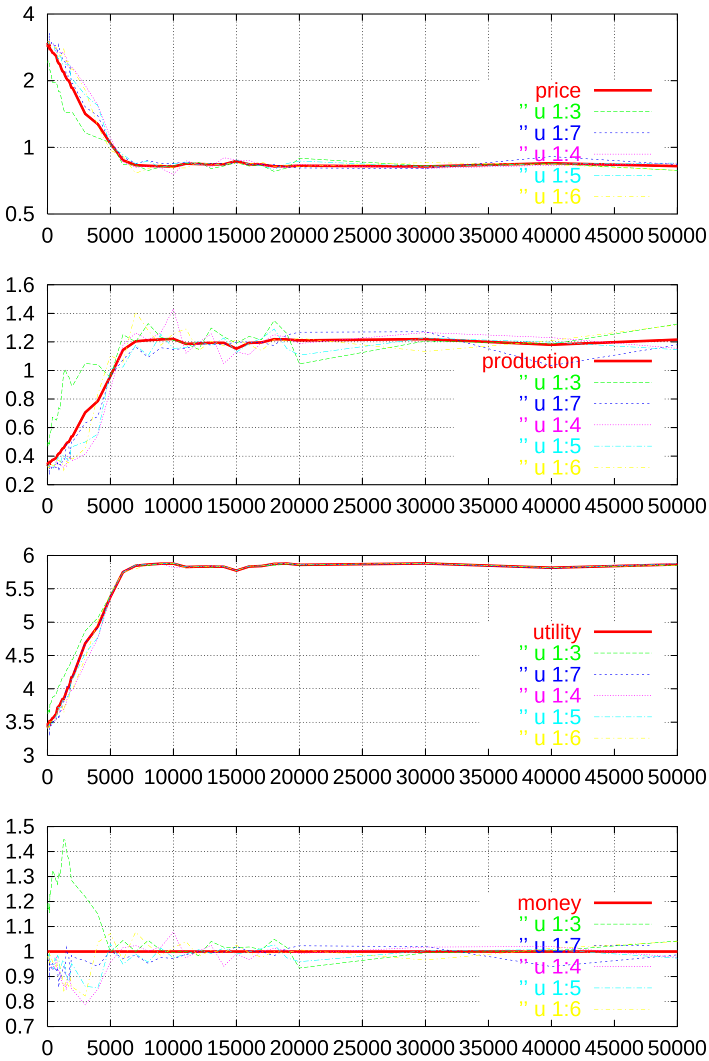

Figure 3: *Simulation where preferences are calculated optimally while prices adapt via trial-and-error. TOP: Prices. SECOND: Production. THIRD: Utility. BOTTOM: Money. The fat line gives the average behavior of the system; the thin lines give the behaviors of single (arbitrary) agents. One sees that the system relaxes towards the same values as in Fig. 6.2. Note the much shorter time scale on which that happens.*

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. The reason for this is simple: A price increase will initially increase my income; the reduction in sales sets in later.

This is also reflected in the Lagrangian solution: One has to solve for the x_{ij} first, when prices are given, and then optimize prices based on the consumption behavior.

9 The BNS model

The results of this paper allows some outlook on the BNS model [2]. In that paper, agents are located on a ring, and every agent buys from the left and sells to the right. This corresponds to the $N = 2$ solution in the appendix, meaning that prices will go to infinity when one uses the dynamics of the present paper (i.e. adapt consumption on a fast time scale and prices on a slow time scale). This can in fact be verified by simulation, and it makes it hard to extend that model towards other scenarios.

The BNS model is however meant differently, and it is important to note that difference. As we have seen in the present paper and elsewhere, for large N, the reaction of the market demand to changing prices should be as

$$q_i \approx \frac{1}{\lambda_i^2 \, p_i^2} \; . \tag{9}$$

BNS now assume that λ is the slowly changing variable, and both consumption and prices are coupled to it via

$$x_{ij} = \frac{1}{\lambda_i^2 \, p_i^2} \qquad \text{and} \qquad p_i = q_i / \lambda_i \; . \tag{10}$$

The question of the BNS paper now is, *given these reactions*, how do they diffuse through the system. This is in fact a different question than the one of the present paper. Nevertheless, one needs the "large N" case in order to justify parts of the dynamics of BNS.

10 Discussion and outlook

10.1 Non-homogeneous situations

In simulations it is easy to change some of the ingredients. This is easiest (but also most consuming in terms of computer time) in the model of Secs. 2 and 6.2 which uses only trial-and-error adaptation. For example, it is possible to give each agent i a different weight a_{ij} for each consumption x_{ij}, leading to a utility function of

$$U_i = -\frac{q^2}{2} + 2 \sum_{j \neq i} a_{ij} \sqrt{x_{ij}} \; . \tag{11}$$

Fig. 4 shows the approach to the relaxed state in such a situation. In this case, the a_{ij} were selected randomly between 0.5 and 1.5 from a flat distribution.

Alternatively, one can use a totally different utility function, e.g.

$$U_i = -\frac{q^2}{2} + \sum x_{ij} \, (1 - x_{ij}) \; . \tag{12}$$

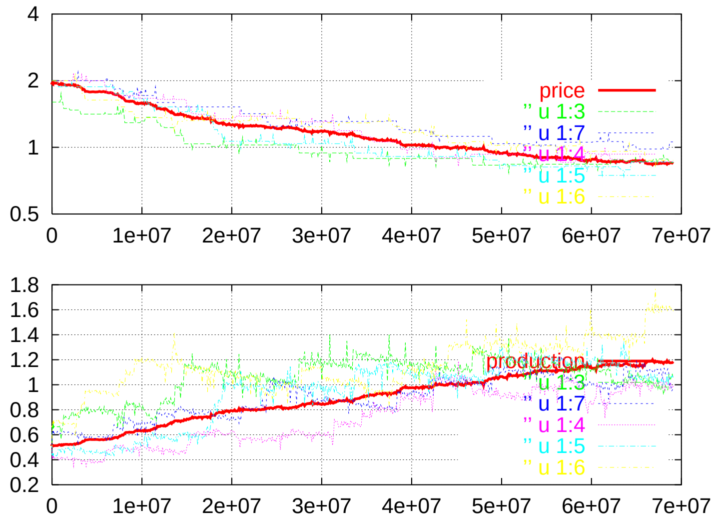

Figure 4: Relaxation for a non-homogeneous situation.

10.2 Approach to equilibrium

This paper considered conditions that make simulations of a simple market model stable. In consequence, the adaptive dynamics was selected in a way that the model converged towards the stable solution.

It is possible to use adaptive mechanisms which show more complicated dynamics. This becomes most obvious when using an underlying spatial structure, for example every agent trading only with her four nearest neighbors on a 2d-grid. In spatial situations, it can happen that information only spreads slowly through the system, leading, for example, to punctuated equilibrium or to avalanches.

10.3 Basin of attraction; path dependence

In the scenario of this paper, agents find the solution via simple hillclimbing. Since there is only one basin of attraction, simulations will go to the same (stochastic) state no matter what the starting conditions and precise rules. The simulations here are not a demonstration of path dependence – however, the approach is a method which can include path dependence if desired.

11 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 bulding simulations models and possibly when regulating the real world.

Acknowledgments

We thank Maya Paczuski and Per Bak for extremely helpful discussions.

A Analytical solution

A.1 Optimal consumption

The optimization proceeds in two steps. In the first step, each agent i maximizes

$$V_i := -q_i^2/2 + 2\sum_{j\neq i}\sqrt{x_{ij}} + \lambda_i\left(M_{i,t-1} - \sum p_j x_{ij}\right) . \tag{13}$$

$M(t-1)$ is the money from the last time step; the Lagrangian multiplier means that all money has to be spent.[3]

Since q_i does not depend on the x_{ij}, the derivatives with respect to x_{ij} and λ_i are

$$\frac{\partial V_i}{\partial x_{ij}} = x_{ij}^{-1/2} - \lambda_i\, p_j \qquad \forall j \neq i \tag{14}$$

and the budget constraint

$$\frac{\partial V_i}{\partial \lambda_i} = M(t-1) - \sum p_j x_{ij} . \tag{15}$$

Setting these derivatives to zero and solving it results in

$$\frac{1}{\lambda_i^2} = \frac{M_{i,t-1}}{\sum_{k\neq i}\frac{1}{p_k}} \qquad \text{and} \qquad x_{ij} = \frac{1}{\lambda_i^2\, p_j^2} . \tag{16}$$

For a homogeneous solution, we can drop all indices and obtain

$$\frac{1}{\lambda^2} = \frac{M\,p}{(N-1)} \qquad \text{and} \qquad x = \frac{M/p}{N-1} . \tag{17}$$

[3]This is already an approximation. Since the model is stochastic, income will not always be the same. However, in the current formulation, the x_{ij} are slowly varying variables. That is, one would have to maximize the expectation value of V_i over the distribution of $M(t-1)$.

A.2 Optimal prices

Now given that one knows the consumer's reaction Eqs. 16, one can calculate the optimal price. For this, one has to replace q_i by $\sum_{k \neq i} x_{ij}$, leading to

$$\tilde{U}_i = -\frac{1}{2p_i^4} \left(\sum_{j \neq i} \frac{M_{j,t-1}}{\sum_{k \neq j} \frac{1}{p_k}} \right)^2 + 2 \sum_{j \neq i} \sqrt{x_{ij}} + \lambda_i \left(\frac{1}{p_i} \sum_{j \neq i} \frac{M_{j,t-1}}{\sum_{k \neq j} \frac{1}{p_k}} - \sum_{j \neq i} p_j \, x_{ij} \right) . \quad (18)$$

Note that λ_i does not depend on p_i.

Solving this in all generality is tedious and beyond the scope of this paper.

Two limiting cases are easy to calculate:

- $N \to \infty$. In this case, the dependence of $\sum_{k \neq j} 1/p_k$ on p_i vanishes and thus all λ_j become independent of p_i. In this case, the derivative of $\tilde{U}_i$ w.r.t. p_i becomes

  $$2p_i^{-5} \left(\sum_{j \neq i} \frac{1}{\lambda_j^2} \right)^2 - \lambda_i \, p_i^{-2} \sum_{j \neq i} \frac{1}{\lambda_j^2} . \quad (19)$$

  Solving for p_i results in

  $$p_i^3 = 2 \, \lambda_i^{-1} \sum_{j \neq i} \frac{1}{\lambda_j^2} . \quad (20)$$

  In the homogeneous case, $\lambda_i^2 = \lambda_j^2 = \lambda^2 = (N-1)/Mp$, and therefore

  $$p = \frac{2^{2/3} \, M}{(N-1)^{1/3}} . \quad (21)$$

  This yields

  $$x = 2^{-2/3} \, (N-1)^{-2/3} \text{ and } q = 2^{-2/3} \, (N-1)^{1/3} . \quad (22)$$

- $N = 2$. In this case, $\sum_{k \neq j} 1/p_k = 1/p_i$. In this case,

  $$\frac{\partial \tilde{U}_i}{\partial p_i} = p_i^{-3} , \quad (23)$$

  and setting this equal to zero means $p_i = \infty$. This may look surprising at first, but makes sense since this is the monopoly situation: Each agent buys only one good, and so the seller can raise prices without bound and still make the same amount of money.

References

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

[2] P. Bak, S.F. Norrelykke, and M. Shubik. Dynamics of money. *Physical Review E*, 60(3):2528–2532, 1999.