# Simulated Social Control for Secure Internet Commerce *

Lars Rasmusson, Sverker Jansson
Swedish Institute of Computer Science
Box 1263, S-164 28 Kista, Sweden
lra@sics.se, sverker@sics.se

## Abstract

*In this paper we suggest that soft security such as* social control *has to be used to create secure open systems. Social control means that it is the participants themselves who are responsible for the security, as opposed to leaving the security to some external or global authority. Social mechanisms don't deny the existence of malicious participants. Instead they are aiming at avoiding interaction with them. This makes them more robust than hard security mechanisms such as passwords, who reveal everything if they are bypassed.*

*We describe our work in progress of constructing a workbench to run simulations of electronic markets. By examining the success of different security mechanisms to avoid maliciously behaving actors we hope to gain insight into how to create electronic markets. The idea of creating* reputations *for the participants is discussed. Finally some legal aspects on using social control and reputation as security mechanisms are discussed.*

## 1 Introduction

The Internet is no longer just a medium for non-commercial informal information exchange between scientists and universities. It has recently become a public network also used to support commercial transactions. The new uses are very different from the former, and it is unclear what will happen when this extremely open network is used in the new context of commerce. It is likely that the introduction of money will be the motivation for criminal activities previously considered uninteresting.

The salient feature of the Internet is it's openness. Anyone is free to add what components (hardware and software) as he/she wishes. This is one of the most important factors for the public acceptance of the net. Even though most Internet shopping malls and advertising services today are closed, company specific systems, the most successful market systems of the future will probably be as open as the Net itself.

*This work was supported by a grant from NUTEK, the Swedish National Board for Industrial and Technical Development.



The components in the open system may be hard- or software originating from different companies. A component have to be able to work with other components conceived long after itself. So, it is impossible to guarantee that the other components in the system will behave in a particular, "nice" way.

As is the case for Internet, no central authority keeps track of who is using it and how. An electronic market with a centralised verifying authority that checks and certifies (human and electronic) participants would be a very non-open solution. Instead, we should look for decentralised and open mechanisms that allow participants in a market to know something about other participants. This should be done without having the participants depend on some central authority, but rather it should be accomplished through the interactions of the participants themselves. We call this family of mechanisms *social control*.

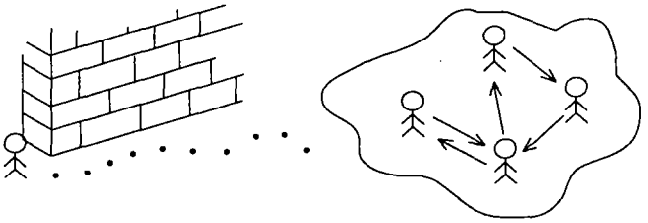

Figure 1: *Hard security, like passwords, leaves the system unprotected if someone finds a way to bypass it. Soft security allows the agents to act as long as they are behaving nicely.*

Social control is a more *soft* security approach than traditional *hard* approaches such as passwords, program verification, access control, capabilities etc. (See figure 1.) Hard approaches demand that you are certain that the components work as intended and if they do so, they will provide a waterproof protection from illegal use. Unfortunately hard mechanisms often fail due to unexpected behaviour of some components. This is what hackers exploit when they use bugs in a mail daemon to gain access to a computer. It doesn't matter that the system used a secure password encryption mechanism since it was possible to sneak in elsewhere. Once the hard security system has been passed, everything lays open to the intruder.

Soft security expect and even accept that there might be unwanted intruders in the system. The idea is to identify them and prevent them from harming the other actors. *There shall never be a key that uncritically opens up all locks on the system.* Social control is therefore a soft security mechanism. An actor is accepted as long as her actions aren't harming anyone else, but if her behaviour changes, she will loose the ability to act, accordingly.

Who will be active in this global open system that we envision? We think that it will be a mix of (human) users and autonomous computer programs. In this paper we use the terminology *actor* when it is unimportant whether it is a human or a computer program who is acting, and we reserve the term *agent* for computer programs. There is a difference between how a system of actors and a system of agents can be allowed to function. In a pure agent system we can allow economic mechanisms to remove agents from the system, mechanisms that could be devastating if they were used on real persons. We also use the term agent when we describe related work in agent oriented computing.

We have built a simple market simulation workbench (available on the net [9]) to simplify discovering and analysing potential threats to an open market system and possible remedies in the form of mechanisms for social control.

In the next few sections, we discuss the notion of social control, describe the tool, illustrate threat scenarios, propose possible remedies, and ask questions about the relation between these and the law.

# 2 Related Work

Underlying our approach is the view of large markets as *ecologic systems* [5], in which the interaction of the participants determines the success of the individual participant. Interesting results can be drawn from game theoretic approaches to the study of ecologic systems and applied to open electronic market systems (or large markets in general) [5]. Non-cooperative games (games with non-binding commitments) can be used to study the cooperation between self-interested participants[7].

The interesting idea of market-oriented programming and how to create design economies in the WALRAS system is described by Wellman in [11]. In their model the market is a tool for *resource allocation*, and it is argued that every computational problem can be transformed into one of resource allocation [12]. The agents' only means of communication is the trade of goods, in an protocol of iterated revealing of preference functions. The model does not include unintended malicious collaboration between subgroups of agents since it doesn't permit side conversations and cooperation. (To assure convergence, the resource usage and utility sets cannot be sub-additive.)

Instead of using market-oriented programming to do resource allocation, we consider a market that is an open environment, possibly supporting actors with malicious intentions and goals. Our interest is to study if forms of social control (of which market economics is one) can make it possible to perform computations even in such environments.

The idea of letting agents themselves answer for the security of the system is inspired by the idea of *mechanism design*. Rosenschein and Zlotkin [10] are using a game theoretic approach for the design of agent communication protocols, assuming that an agent will act *rationally*.

Soft security mechanisms for intrusion detection have been tried by Crosbie and Spafford [2]. They audit program actions and by using a number of sensors that alert if the behaviour of a program seems unusual they are trying to find intruders who are executing programs not normally executed by the users. (Similar methods are being used by credit card companies who try to discover fraud by searching for card owners who drastically changes their buying behaviour.)

One attempt to detect actors who don't behave as expected is being made by A. Rasmusson[8]. A personal security assistant collects information about the actions made by an actor and is trying to find patterns of abnormal behaviour in actors. This could be used to build a very fine grained reputation system that lets other actors know of abnormal actors.

# 3 Social Control

## 3.1 Everyday Examples

An intuitive way to think of social control is by taking examples from our every day world. We have a opinions about numerous things, opinions which help us know how to act in different situations. Often these opinions are formed through social mechanisms.

*Ex:* People who live in small villages sometimes complain about feeling watched by their neighbours. They feel that if they don't behave according to the social norm, they will ridiculed and talked about behind their back. Still, they don't want to move to a big city since there are a lot of crazy people there, and they find it awful that people can live for years door to door to a complete stranger.

Some effects can be good. In a small village, life is governed by the social control of the local community. This is why people can leave their bike unlocked, being sure no-one steals it. In a big city the size of the crowd reduces the social control. If a person buys a bad car, she can't prevent other people from going to the car dealer, and in a big enough city a dishonest dealer can can continue his business regardless of the clients opinion about him. In a small town he would soon loose his clientele if he tried that.

Some effects can also be bad. The social pressure to conform can hinder people who want to do *different* things. An artist or a writer might be considered weird and odd by the locals. A person who once has been excluded from the society (perhaps after having been hospitalised for some mental illness) might never be accepted again. In a big city it is possible to start over, make new friends who base their feelings for the person on his current behaviour.

*Ex:* Another social mechanism is that different groups develop different tastes and preferences. Music is a good example of this. There is no objective scale that can measure the quality of *all* music. Instead, when we want to buy a new record, we use the advice of people who have approximately the same taste as us (friends, music reviewers etc.). Classical music and hop hop can coexist, each of them developing its own quality scale. In fact, often *many* competing scales develop within one genre, which is seen in the plethora of reviewing magazines that exist for (for example) pop music.

In summary, *social control is a group behaviour that indirectly forces the group members to behave in a particular way.*

## 3.2 Social Control in Electronic Commerce

Internet is as an open system more like "a big city" than "a small village", in that it is possible to act in any possible way without anyone being able to stop it. If this is permitted in electronic markets there will be a large amount of fraud and con men doing business there. In a big city you can't know who you are dealing with if you meet for the first time.

How is it considered possible to negotiate, cooperate and perform commerce in an open electronic system if there is

now way of formally knowing the intentions of the other actors? By looking at how the market does commerce today, we find that there exists many informal mechanism and ideas that are absolutely necessary for commerce to work. There are notions such as *reputation*, a *stability over time*, and *estimated risk* of being conned, all of which help one company to choose with whom to do business.

But first, let's introduce a game of two prisoners. This will cast new light on some issues in the security of large electronic markets.

### Prisoner's Dilemma in Global Markets

The *prisoner's dilemma* (name coined by A. Tucker and made famous by R. Axelrod in [1]) is a non-cooperative game with two prisoners (see figure 2). The prisoner's dilemma shows that if chances that the prisoners will ever meet again are low, then the best strategy is to try to cheat on the other prisoner. But, interestingly, if the game is iterated, then the best choice of action is to "cooperate while the other prisoner cooperates." The strategy *tit-for-tat* is found to be an *evolutionary stable strategy*, ESS. Such a strategy is the best strategy for a player provided the other players play by the same strategy.

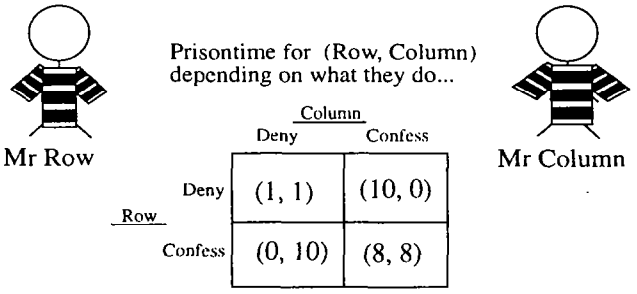


Figure 2: *Should Row and Column confess or deny to minimize their prison time? Since they gain most if they confess, independent of what the other does they will confess, even though they only get one year if both deny.*

The prisoner's dilemma can tell us a lot about automatic electronic markets. A transaction between two parties is most certainly a one shot deal. Probability is low that the two will do business again. This means that the agents find themselves in a game similar to that of the prisoner's dilemma, and that the best game theoretic move is to try to cheat the other agent. What's missing is *a means to bring the system back to an iterated prisoner's dilemma-like game.*

If the actors on an electronic market are dependent on their reputation to act, their earlier actions are important for how other agents will treat them, and this is exactly what makes the iterated prisoner's dilemma support cooperation. A form of social control, or put in other words, a distributed locally implemented mechanism, can help to achieve cooperation in an open system which would otherwise be prone to extremely malicious agents.

## 3.3 Mechanisms for Social Control

Social control is a way for the population of actors to avoid unwanted actors. Let's consider some cooperation mechanisms used by non-cooperative agents. That is, mechanisms that promote cooperation even though they are free not to cooperate if one agent find that to be the best action.

The Clarke Tax (as described in [3]) is a *one-shot* voting-by-bid mechanism for group decision making that eliminates the advantage of trying to manipulate the result. There is no risk that a rational agent chooses not to cooperate (manipulate) since this doesn't maximise the expected utility. Reducing agent communication before the vote, as in Clark Tax, can therefore lead to systems in which cooperation is the most efficient action.

But there must be some way to assure that the agents act rational and that non-rational agents are removed. In computational economies it is common to introduce a fictitious currency to do elimination of bad alternatives. If agents are paid according to their rationality and have to pay a small fee to stay on the market, then non-rational agents will run out of money and will be removed from the market. Without this cost of staying on the market there will be room for all sorts of actors on the markets, some of them non-cooperating, malicious, some of them just wasting system resources.

If the system is open there is no way of making sure that there is no prior communication between the agents before the Clarke Tax voting is made. If some of the agents form a collaborating subgroup (conspiracy) they can undermine the fairness of the voting. Although very important, these mechanisms alone might not be sufficient to assure a desired behaviour of an open system.

Actors on a market must be able to refuse to cooperate with others that are found to be disobeying the rules of the game. There are two reasons why someone might be doing this. The first one is that they are trying to fraud the other party. The second is that they could be acting "non-rationally" because they have incomplete information, because of a mistake by the programmer (a bug) or some fault in the system. It is not safe to trust all security to the game theoretic mechanisms and assume that all actors are completely rational. They must be completed with a way to dismiss non-rational (possibly malicious) agents. And a way to do this without introducing global authority mechanism is through social control among the actors.

All this raises a number of questions. How can it be possible to introduce new, reputationless, agents into the system? Can an agent build up a good reputation and use this to commit crimes? What happens if someone wrongly gets a bad reputation and is hindered from doing business? Is any of this even legal? We will try to address some of this below, but first we'll look at what mechanisms of social control are there that make open agent systems not promoting cheating agents?

Different behaviours of the participants will give rise to different emergent effects, and this can be used to advise implementors of such systems how they can go about to avoid creating systems that are prone to explosions of black markets and other criminal behaviour in the actors

# 4 A Commerce Simulation Workbench

By performing experiments on models of large markets we hope to verify our intuitions and gain further insight to what problems can arise in large open markets with small means of control, and which remedies are effective to avoid and remove those problems.

To do this we are developing a workbench for simulating systems of agents which produce services and sells them on to other agents (see figure 3).

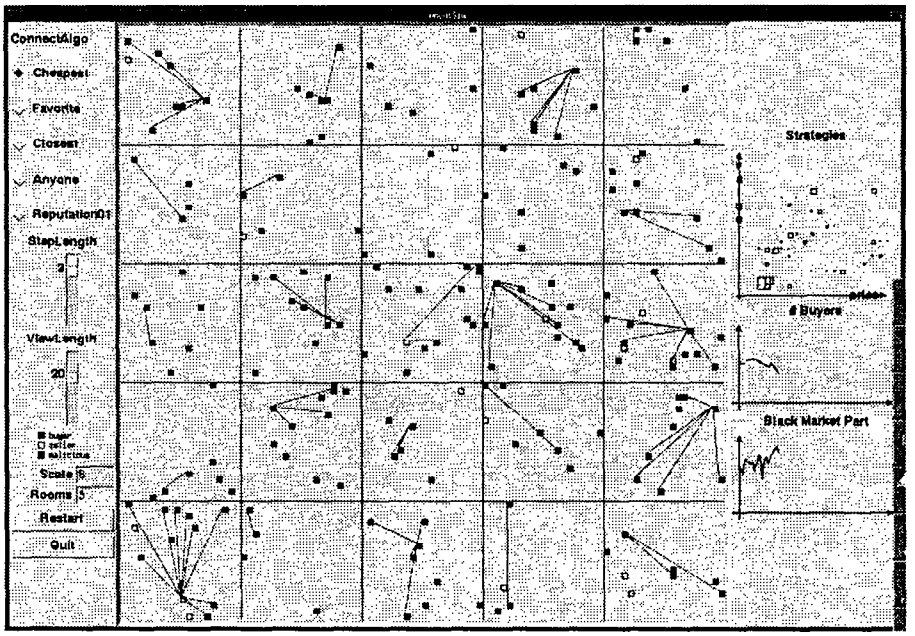


Figure 3: *Screenshot of a commerce simulation workbench*

*Outline of the current model:* The market world is a 2-dimensional surface populated with actors who can be buyers or sellers. The actors range of sight can be constrained either by a maximum length, or by walls subdividing the surface into rooms. The actors move about randomly and for each step they choose who to do business with among the visible actors. If the actors move into another room they can't see anyone in their previous room.

A seller is modeled as an entity offering a service for a fixed *price*. The service provided by the seller has a particular *value*. This value is not known to the buyer when he/she is deciding on doing business or not. An actor who is offering a service with a high price and a low value is either *malicious* or *inefficient*. (The two are the same in this model.) Currently all producer produce the same service to avoid side effects caused by the trade of other services.

To produce a service, the seller must pay a certain amount of money, for example half that of the value. Remember, an agent buying a service has no way of knowing the value of the service, it can only look at the price. After the transaction is done, the *value* of the provided service is added to the capital of the buyer.

Every actor gets an initial capital. For every turn of the simulation the actor has to pay a certain price to participate. When the actor is out of money it is removed from the simulation.

The number of actors is fixed, therefore new actors can only be introduced when another actor is removed. This is done by making a copy of an actor who has a large capital and inserting it at random into the world matrix.

By running simulations on market systems where the actors have different methods for selecting cooperation partners we can see which of these methods that succeed in choosing good sellers, that is, sellers with larger produced value than price plus buyer's participation cost.

If the system stabilises when there is a population of non-malicious sellers in many different price categories and where the price is correlated to the value of the service then the actors have managed to auto-organise them self into a market where a buyer can choose the quality of what he/she is asking for simply by looking at the price. Malicious agents can't make money on their services, since they only provide a very low value.

So far we have studied what happens when an actor pick *anyone*, the *closest* seller, the *cheapest* seller, sticks to a *favourite* seller or to a locally *recommended* seller. We find that by disallowing buyers to choose from all sellers a larger number of seller agents can co-exist. Furthermore, by creating *regions* of commerce (rooms), the system is able to more quickly get rid of malicious agents. This happens as the actors in regions infested by malicious actors go bankrupt, making the malicious actors go bankrupt too. All of this is very dependent on the *migration level* between the different regions. High migration leading to an almost global environment is much more sensitive for bad agents.

# 5 What Can Go Wrong in a World Without Trust

We have already observed some, and hope to observe more of fraudulent behaviours as we expand the model to make the actors more sophisticated. Some of these behaviours are described below.

## 5.1 "Con-mania"

In an electronic market economic transactions take place at a rate much higher than what is usual in the ordinary market. Therefore it is possible to make a lot of money even on very small crimes. By falsely claim to perform a certain small and insignificant task, a person can rapidly con a large number of actors and make a lot of money (see figure 4).

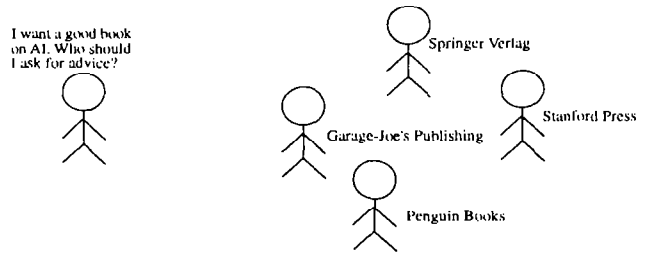


Figure 4: *An actor on an open market is facing the problem of deciding with whom she shall make business.*

It is not probable to think that an actor always is able to decide whether he/she have been ripped off or not. For example, if I want to know the size of the population in Canada anyone can propose to answer and give me a random number between ten and fifty million.[1] By the time the information is taken into use it might not have been possible to verify it.

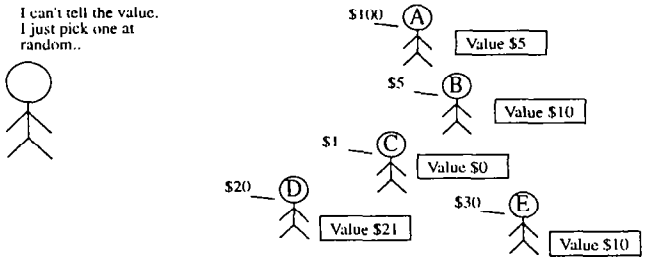


Figure 5: *One way to choose business partner is to pick one at random. This is very vulnerable to con strategies since there is no way to retaliate on bad behaviour.*

If we run a simulation in a one-room world with mostly good sellers where the buyers choose between all the sellers at random (see figure 5), we find that bad sellers who are charging a high price and giving a low value in return will make the most money. Good sellers will loose market shares to bad sellers since more and more buyers go bankrupt and more new bad sellers enter the market. Soon there are so many bad sellers that the buyers disappear completely from the market.

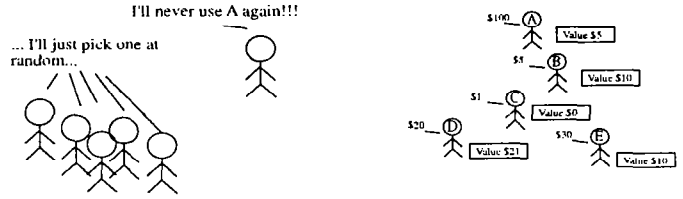


Figure 6: *With many buyers there is always someone new to con.*

On a global market there will be so many customers that even larger, detectable frauds can be profitable (see figure 6). Since there are so many people to con, there will certainly be people who haven't heard of the fraud and who will fall for it.

[1] FYI: The population of Canada was 29,413,100 on January 1st, 1995.

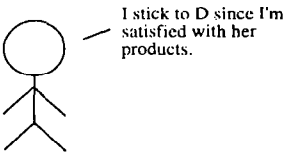


Figure 7: *A buyer who is sticking to the same seller if the seller is nice will manage ok, but he will not find the optimal value.*

A simulation where the buyers use their *favourite* seller, that is, a seller that didn't fool him the last time, illustrates buyers who can distinguish that they have been fooled (see figure 7). We find that the malicious sellers disappear from the market. But, if their view length is limited (35) and they move about a lot (10) they will occasionally choose other sellers since they loose track of their favourite (this illustrates new buyers) . There will be a low but constant number of bad sellers.

If the buyers choose the *cheapest* seller then all the buyers will quickly disappear since all the money will go to cheap malicious sellers. There is no way for the nice sellers to make enough money to stay in business. All the buyers will run out of capital and be removed from the market. But, if there is very little migration on the board then there will sometimes be an some buyers who only find a nice seller within their range of sight. The system will finally stabilise at a population of nice sellers, but at a terrible cost. Almost all of the original buyers will have gone bankrupt. This leads us to conclude that even though several mechanisms can arrive at "nice" populations, some can be more "drastic" than others in that they demand the elimination of a lot of innocent actors.

In an electronic market new electronic actors might be introduced in such a high rate that it is impossible for a single customer to know whether the new actors are actually doing anything, or whether they are only ripping customers off. If it is possible to create a new electronic identity, old actors can reappear in new clothes, repeating a fraud again and again.

If a large part of our day-to-day transactions are made through an open system with largely unknown components, it might be possible for someone to collect information about the persons using the system. This information can be used to create a computer shadow of a person, which can be used for purposes ranging from directed advertising to blackmail. How can we prevent this from happening?

If we are using electronic agents, who are migrating programs, then programs of possibly unknown origin can execute on a person's local computer. It could be a virus or a Trojan horse trying to steal information. Remember that since the system is open, anyone can let anything in into the system. A malicious program doesn't even have to be written in spite. It's not far fetched to imagine an agent turning into a computer virus because of unintended bug.[2]

## 5.2 "Monopoly"

If everybody use the same sources for information about which other actors to use then a monopoly will quickly be established. This leads to extremely sensitive systems both because the buyers are in the hands of the whims of the

[2] Remember the Morris Internet Worm?

seller, and because a change in the demanded services is less likely to be satisfied, since there are no competing, slightly different services being offered.

A simulation with one room where the buyers use the room reputation (a "top ten list" for that room, see figure 8), there will finally be only one seller left since the others have gone bankrupt. If the seller is permitted to mutate to a bad seller the entire population will finally go bankrupt.

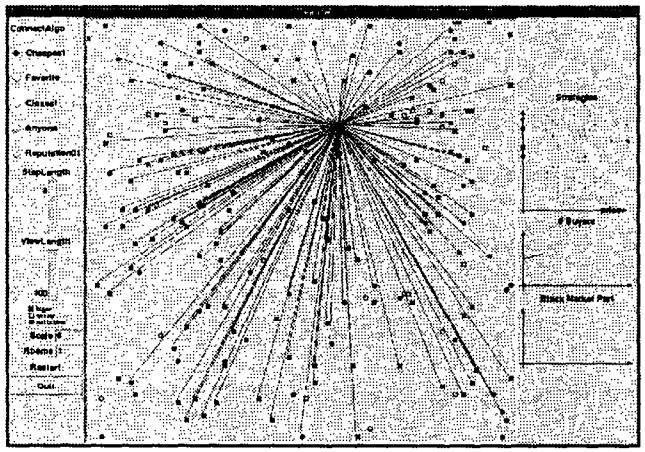

Figure 8: *Monopoly can be established when everybody are choosing using the same algorithm of choice in a global market.*

To avoid monopoly we can either create a system with many different different valuation measures and different price levels. We can also introduce the notion of location in the agent system. If it is costly to move from one place to another then an agent will promote the alternatives close to him. This allows for the proliferation of several different sellers in different regions.

If we rerun the same simulation in a world with, say, 9 rooms, there will be more sellers, about one in each room. They don't have monopoly since the buyers can move from one room to another. If the seller changes to bad behaviour, the buyers who stay will go bankrupt and the ones who move will survive. Finally the bad seller goes bankrupt and a new good seller can be introduced.

By introducing location we have to sacrifice the notion of optimal choice (or redefine it to include proximity in its measure). This can be difficult to accept, but perhaps this is better than the alternative, monopoly.

Still this mechanism yields a very low number of sellers on the market. This makes the buyers very sensitive to when the seller leaves their range of sight. When this happens, they often have to choose at random again, and they are possibly choosing a bad seller. Since the bad seller makes a lot of money on this he can afford only being used very seldom.

# 6 Mechanisms of Social Control

Here we sketch on some possible mechanisms for social control that are inspired by establishing opinions about the other group members. These are either local to each individual, distributed or trusted to a third part.

## Promotion

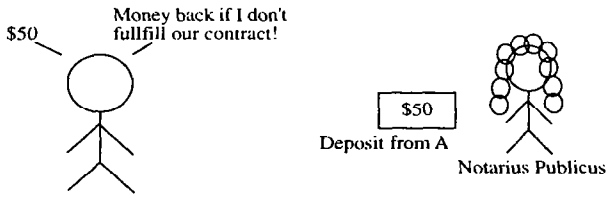


Figure 9: *The seller is offering money back, using a trusted third part*

When a new agent is introduced to the market it has no reputation and therefore it is highly unlikely that anyone will use it. The agent can promote itself by giving a "money back guarantee" and leaving the money to a third, trusted part (see figure 9). Therefore the agent needs some initial capital as an insurance for its client, but gradually, as the reputation builds up, this will be less and less necessary. The value will be in the reputation, and loss of this reputation will be at least to the cost of the interest of lending the money to build up the reputation (not counting the cost of the profitable deals that could be made using that reputation).

## Gossip and Rumours

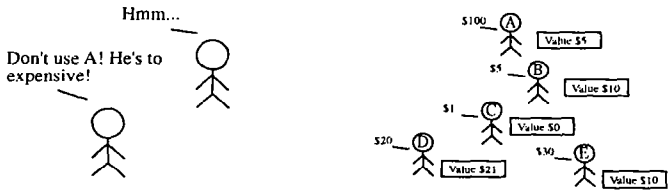


Figure 10: *The actors can themselves spread the rumour about bad agents.*

If an agent discovered a good cooperating partner it could inform the other agents about it, and likewise if it had been cheated. The agents gossip about each other (see figure 10). This quickly ruins the local market for a malicious agent.

However, in ecologic systems, the ones who are to gain the most from this information is by large probability the competitors to the gossiping agent, and therefore an agent might be better of by lying. This is the *paradox of altruistic communication*, which usually ends up in no communication at all.

The question about how the gossip spreads in the global society is also unclear. If the agents have low or none incentive to move around, actors who move faster that its rumour can find new clients.

## Reputation and Reviewing

If we introduce actors who are concerned with the maintenance of reputation of others they can charge others for giving them advice on whom to choose for a particular task. These *reputation agents* can use the money to buy services from the vending agents (see figure 11). This idea is very similar to that of restaurant critics, and for them, as well as for agents, it is very important that they act *incognito*.

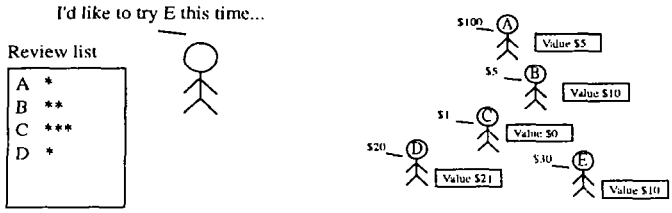


Figure 11: *A reviewer can maintains a list of reputations*

If they weren't incognito they might get a special treatment and they would be unable to give a correct judgment.

It has been suggested [6] that for ideal critics reviews the buyer has a large incentive to be anonymous (otherwise he/she can't trust reviews as a means to find good deals), and the (non-malicious) seller has every reason to prove his identity to the buyers. Otherwise he/she won't get the credit he/she deserves.

Reputation is in many cases a subjective measure, as was exemplified above with the example about music. By permitting different reputation scales to co-exist and compete, it will be possible to use the reputation mechanism as a tool for finding and categorising information. Variations of this idea is today tried for filtering Usenet news.

Possibly reputation can be used in self-improving systems where the reputation corresponds to how well a service is performed [5]. Actors who's services perform badly compared to others' will gradually be replaced by the services of actors with better reputation.

# 7 Security and the Law

The legal aspects of actors in information systems are being investigated by Karnow and others. The question is raised whether electronic agents' digital identities should have their own rights and whether it is an infringement of the rights of the client if an agent acting on his/her behalf is being denied access to a service.

What if there are agents that will only service Swedish users? Is this equal to discrimination of nationality? Or what if a person can't use the global Internet market because he/she once made some crime? Even though the person already has paid his debt, he might find himself forever being refused service on arbitrary grounds. This *denial of service* could be a serious problem for automatic reputation systems.

How should a disagreement between two electronic actors be resolved? Krogh [4] proposes that agent systems should have normative legal systems that should be used when such matters have to be settled. The question is whether or not these rules have to be coded into the system and globally enforced. It would be most satisfactory if this was not the case, but since the Internet surely will have effects on the world outside of it, it might be necessary to make legal systems in different countries be able to work together.

In short, there is need for a way to *prove* and *motivate* denial of service as well as *electronic contracts* that allow unpartial jurors to decide who is being wronged by whom.

An interesting observation is that buyer *anonymity* can help preventing some kinds of discrimination. If an actor is anonymous there is no way anyone can know the ethnic group, sex or nationality of the client, and therefore discrimination on these grounds is made impossible.

# 8 Conclusions and Further Work

## 8.1 Relevance of the model

The current market model is far from optimal. It is unclear how many aspects of real markets it captures. For now there is no notion of advertising, the actors can't choose where to move, so even the "favourite" and "room reputation" algorithms are choosing somewhat randomly. The production of goods is strictly additive, that is, there is no way that sellers can bundle different goods together using co-production advantages to lower the price since there is only one product. Discussions with people more familiar with economic market models will hopefully help to remedy these, and other, shortcomings.

If bankruptcy is to be used as a means for regulating the paricipants on the market, then it migth be difficult for real people to interact on the same market as the electronic agents. In a pure electronic systems it's just a way to create adaptive systems, as in the agoric computing of Miller and Drexler or in the computational ecologies of Huberman, but in systems including real people, these mechanisms might be too strong. By limiting the system to only include electronic agents, there are more parameters to modify, hence more possible social control mechanisms that can be examined. We hope to investigate mixed as well as pure systems.

Social control is not something that is obviously just good. If actors are excluded, rightly or wrongly, from the market, chances are that it is fatal to them. We continue to look further in to what kinds of problems can be caused by these mechanisms.

By using reputation agents, an actor looses control over what information he/she transmits to the other actors. This is all left to the reputation agent. Perhaps the actors can better assure that their information is not lost somewhere if they themselves are the ones responsible for sending out the information. Still, this returns us to the problem of collecting and refining the enormous amounts of information that will be available.

## 8.2 About future markets

Human markets will with extended global connectivity change to be open systems with strongly interacting components through the use of telecommunication, network and agent technologies. If such large systems are ever to be put in practical, everyday use, we must be sure that they cannot run amok.

In open systems it might be impossible to know and to be sure of the detailed behaviour of the components in the system. Furthermore, since components can be added and removed as time goes, the global properties of an open system can change rapidly.

The more the parts in such a system are permitted to interact, the more unpredictable becomes the result of their collective actions. Systems dealing with tasks in our everyday life must be assuredly stable, in the sense that adding components to the system can't open it up to criminal behaviour, permitting new ways to safely fraud people.

In open, Internet based, markets it won't be possible to enforce or assure a certain behaviour of all the actors. By simulating markets with interacting agents with *evolutionary stable strategies* along with some means to promote *rational* behaviour in agents (reward for good work and cost of execution), it seems possible to find strategies giving the

agent society the emergent property to (without any global control) expel malicious agents.

Socially controlled systems are ecologic systems. Ecologic simulations of agents using different rules for choosing collaboration partners will be used to measure different social methods' efficiency in detecting malicious agents. Since hard security fails in open systems, we must understand how to use these mechanism if we want to create secure open systems for Internet commerce.

# References


[1] Robert Axelrod. *The Evolution of Co-operation.* Penguin Books, 1984.

[2] Mark Crosbie and Eugene Spafford. Applying genetic programming to intrusion detecion. Technical report, Purdue University, Department of Computer Sciences, 1995.

[3] E. Ephrati, Gliad Zlotkin, and J.S. Rosenschein. *Meet Your Destiny: A Non-manipulable Meeting Scheduler,* volume Conference on Computer Supported Cooperative Work. North Carolina, October 1994.

[4] Christen Krogh. The rights of agents. volume Intelligent Agents Volume II - Proceedings of the 1995 Workshop on Agent Theories, Architectures and Languages (ATAL-95) of *Lecture Notes in Artificial Intelligence.* Springer-Verlag, 1995.

[5] M.S. Miller and K.E. Drexler. Comparative ecology: A computational perspective. In B.A. Huberman, editor, *The Ecology of Computation.* North-Holland, 1988.

[6] M.S. Miller and K.E. Drexler. Market and computation: Agoric open systems. In B.A. Huberman, editor, *The Ecology of Computation.* North-Holland, 1988.

[7] Eric Rasmusen. *Games and Information.* Blackwell, 1989.

[8] Andreas Rasmusson. Personal security assistant for secure internet commerce. (position paper), unpublished.

[9] Lars Rasmusson. Simwb. URL: http://www.sics.se/~lra/simwb.

[10] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Desiging Conventions for Automated Negotiation among Computers.* MIT Press, 1994.

[11] Michael P. Wellman. A computational market model for distributed configuration design. volume Proceedings of the Twelfth National Conference on Artificial Intelligence. Seattle, WA, August 1994.

[12] Michael P. Wellman. The economic approach to artificial intelligence. ACM Computing Surveys Symposium, September 1995.