JASSS

Oliver Will and Rainer Hegselmann (2008)

A Replication That Failed: On the Computational Model in 'Michael W. Macy and Yoshimichi Sato: Trust, Cooperation and Market Formation in the U.S. and Japan. *Proceedings of the National Academy of Sciences*, May 2002'

Journal of Artificial Societies and Social Simulation vol. 11, no. 3 3
<http://jasss.soc.surrey.ac.uk/11/3/3.html>

For information about citing this article, click here

Received: 30-Aug-2007 Accepted: 29-Apr-2008 Published: 30-Jun-2008

Abstract

The article describes how and why we failed to replicate main effects of a computational model that Michael Macy and Yoshimichi Sato published in the *Proceedings of the National Academy of Sciences* (May 2002). The model is meant to answer a fundamental question about social life: Why, when and how is it possible to build trust with distant people? Based on their model, Macy and Sato warn the US society about an imminent danger: the possible break down of trust caused by too much social mobility. But the computational evidence for exactly that result turned out not to be replicable.

Keywords:
Replication, Social Dilemmas, Simulation Methodology, Cooperation, Trust, Agent-Based Modelling

Introduction

1.1

In May 2002 the *Proceedings of the National Academy of Sciences* published an article written by Michael Macy and Yoshimichi Sato. The abstract immediately arouses interest:

> Compared with the U.S., Japan is believed to have a collectivist culture that nurtures high trust. Results from laboratory and survey research, however, show that Americans are more likely to trust strangers than are Japanese. Why would trust be lower in a collectivist culture? We use an agent-based computational model to explore the evolutionary origin of this puzzling empirical finding. Computer simulations suggest that higher social mobility[1] in the U.S. may be the

> explanation. With low mobility, agents rarely encounter strangers and thus remain highly parochial, trusting only their neighbors and avoiding open-market transactions with outsiders. With moderate mobility, agents learn to read telltale signs of character so that they can take advantage of better opportunities outside the neighborhood. However, if mobility is too great, there is too little trustworthiness to make the effort to discriminate worthwhile. This finding suggests that higher mobility in the U.S. may explain why Americans are more trusting than Japanese, but if mobility becomes too high, the self-reinforcing high-trust equilibrium could collapse. (Macy and Sato 2002, p.7214).

1.2

The last three sentences at the end of the article send at mild shudder through the reader. There we read:

> In particular, the nonmonotonic effects of mobility raise a warning flag for the U.S. Although American society enjoys relatively high trust and participation in global markets, there is no guarantee that this will continue indefinitely. Rapid advances in telecommunications could undermine the embeddedness of social relationships needed to make trust and trustworthiness self-reinforcing. (Macy and Sato 2002, p.7220)

A lot is at stake: namely, modern western social life. And it is an agent-based computational model that bears the burden of proof for the threatening cultural diagnosis.

1.3

Macy and Sato's research was originally motivated by a seeming anomaly in findings on the cultural differences between the U.S. and Japan:

- On the one hand there are — explicitly mentioned by Macy and Sato — social psychologists like Casson (1991), Ouchi (1981) and Dore (1987) who state that the Japanese live under more group obligations, less deviance, stronger relational commitment and lower social mobility. Macy and Sato quote Hagen and Choe (1998) who summarise that — among social scientists — there is a "widely held view that Japan is a collectivist culture in the sense that people's self-identification tends to be deeply rooted in group membership." (Macy and Sato 2002, p.7214[2]). Macy and Sato believe it is commonly thought that these findings lead to the expectation that the Japanese are more likely to trust and cooperate than U.S. citizens who live in an individualist culture (Macy and Sato 2002, p.7214[2]).
- On the other hand, serious research does not support this view. Macy and Sato report on the work of Yamagishi (1988), whose laboratory research found that in the absence of sanctioning systems the Japanese have weaker trust and cooperate less than subjects from the U.S. (Macy and Sato 2002, p.7214[3]). He further approved this result in a survey on Japanese businesses, which turned out to shun better deals on the open market in favour of suppliers with whom they had established relations, (Yamagishi and Yamagishi 1994).

Why? "Lower trust in a collectivist society poses an interesting puzzle" (Macy and Sato 2002, p.7214[4]) — and with this theoretical diagnosis Macy and Sato are quite right.

1.4

Yamagishi attempted an explanation that resorts to two different 'operating systems' (our words; O.W and R.H.):

The Japanese

> … tend to believe that trustworthy behavior occurs only when it is prudent, as when transactions are embedded in tight social networks where malfeasance is discouraged by the need to maintain a good reputation in a set of dense social interactions (Macy and Sato 2002, p.7214[4]).

The Americans

> … are more likely to believe in character that produces trustworthy behavior even when someone could cheat with impunity. Trustworthy individuals learn to send and receive reliable signals that allow them to avoid miscreants. Thus, Americans focus not on social or physical proximity but on a Calvinist concern for telltale signs of character. Strangers can be trusted if they display appropriate emotional, cultural, and social cues that cannot be easily faked. This strategy allows them to participate in unembedded exchanges in the open market in relative safety (Macy and Sato 2002, p.7214[4]).

1.5

Instead of making them two different operating systems — as Yamagishi does it — the central idea behind the Macy-Sato-model is to combine the embeddedness-approach with a read-telltale-signs-approach. The former is necessary to learn the latter.

> We hypothesize that embedded relations are not just a provincial alternative to participation in the open market but are the classroom where the skills are developed for navigating outside the neighborhood. (Macy and Sato 2002, p.7214[5])

Obviously the Macy-Sato-model is meant to answer a fundamental question about societal life: Why, when and how is it possible to build trust with distant people? By contrast, it is not that difficult to understand why cooperation normally works in small family or tribal groups. But trusting of strangers? A riddle!

1.6

The question has a long scientific tradition that goes back to the ancient Greeks. In one of Plato's dialogs, the sophist Protagoras gives a very modern answer which — after some deciphering of the myth in which it is couched (a myth about Prometheus and Epimetheus) — amounts to saying: A high blood toll was paid to learn the lessons. But finally mankind invented both moral virtues and enforcement agencies (Plato 1961). The two together made it possible to live together in comparatively wealthy large-scale societies where a high proportion of interactions are no longer based on family ties or good personal acquaintanceship. About two thousand years later, David Hume (1975a, 1975b) gave a very similar answer — though no longer presented by telling a myth. Additionally, Hume stresses the importance of division of labour and mechanisms of reputation.[2]

1.7

Thus, in the tradition of Protagoras, Hume and other modern scholars (some of them mentioned in the first chapters of their article), Macy and Sato address an old and fundamental problem of understanding social life. But they do it using new methods: modelling and simulation. That is a good idea: Plato delivered a myth and Hume only some draft thoughts. With the new methods at hand we can hope for much more, for instance, to make old myths and draft ideas more precise. We now should be able to analyse in detail the critical components and the loops in their interplay. It should be possible to identify inter- and intra-agent processes that are important. Based on that there should be a fairly good chance to achieve a much better understanding of why, when, and how it is possible to build trust with distant people.

1.8

New chances rarely come without new problems. That holds for modelling and simulation as weill. One of them is replication and replicability: a computational model is normally a complicated 'machinery' — even if it does not involve thousands of lines of code. Modelling errors and programming errors of all sorts may have occurred, among them those that make nonsense out of the main results.

1.9

Replication by others is an obvious answer to that problem. Nevertheless, there are many more models than replication. The reason is clear: if one succeeds in replicating a model, one has nothing new to tell — and even worse, nothing new to publish. Therefore, investment in one's own models seems more promising than squandering time and effort on other's models! There is evidently an incentive problem.

1.10

However, we did a replication of the Macy-Sato-model. Additionally, together with others we are working on the EU-funded international EMIL-project "Emergence in the loop: simulating the two way dynamics of norm innovation" (6. Framework, Information Society Technologies). Under the project's perspective the Macy-Sato-model looked interesting: Firstly, the model contains intra-agent processes by which agents can become more or less morally reliable agents. Secondly, the model ties together intra-agent processes with inter-agent exchange processes. That is exactly what the EMIL-project aims at. Thirdly, the EMIL-project focuses on norm innovation, as it is ongoing among humans probably since prehistoric times. Lots of norms had and have a short-lived existence. But there is a very special norm innovation process ongoing since the Palaeolithic or even earlier ages: It is the process by which at least in parts of the world and for parts of the population the circle of those who should be treated trustworthy, fair, just etc. stepwise expanded from family members, tribal fellows over fellow citizens to, finally, human beings of whatever race, colour, gender, world view or other orientation. At first glance the Macy-Sato-model looked as if it could be a useful first step to model and thereby understand what one might call "expanding-circle-effect" (Singer 1981) — the probably most fundamental moral innovation process in human history.

1.11

There was a further reason for the replication, simply curiosity, induced by Macy and Sato's warning that there might be a threatening 'too much' of mobility that, in the end, could lead to a total collapse of trust into others. That we wanted to understand in detail.

1.12

As the title of this paper suggests, *the replication failed with regard to important aspects of the model*. In particular, *we could not reproduce the effect that above a certain level of mobility trust breaks down*. We tried our best, but the descriptions of the used model were often not detailed enough. Some parts of the descriptions were inconsistent. In several e-mails we asked the authors to deliver more information and the source code but they did not. Following a principle of charity we tried to complete the missing information ourselves.

1.13

The following reports on our difficulties and results. It begins with a description of Macy and Sato's model on a level that does not go much below the surface. It contains information about the aspects of the model that seemed to be clear to us and thus gives the reader a good overview. Afterwards we report on our trial in replication and deliver detailed information on the models assumptions and the problems that arose in the process of replication. We will elaborate why and that one important aspect of the original model could not be re-implemented: the concept of transaction costs. A comparison of results from the original and the replicated model supports this thesis. We end with some concluding remarks.

1.14

Before we start, a comment on citation: As the topic of this work is the replication of Macy and Sato's model, there are many references to their paper. To make it easy to find text passages, the references to Macy and Sato's publication contain page numbers and — in brackets — the number of the paragraph we are referring to. The first words on a new page are always referred to as paragraph 1 even if that paragraph is a continuation of the last one of the previous page.

Macy and Sato's model

2.1

Influenced by Yamagishi's empirical research, Macy and Sato hypothesize that there are two strategies to deal with strangers. The first — referred to as the parochial strategy — compels an agent to always trust its known neighbours but never trust strangers. This strategy is reported to "minimize transactions costs" (Macy and Sato 2002, p.7214[3]) that might arise from being cheated by strangers. On the other hand it also implies opportunity costs because one may miss good deals with agents from outside the neighbourhood. The second trust strategy is to base the decision out on signals of the potential partner. In contrast to the parochialism, this signal reading strategy allows agents to deal with strangers.

2.2

Macy and Sato hypothesize that social mobility plays an important role in the evolution of trust in strangers because it gives local actors an incentive to learn how to interact effectively with new partners. Actors learn to send and read signals in local interaction and afterwards signalling systems "diffuse to other regions through movement of neighbours to socially distant regions" (Macy and Sato 2002, p.7214[6]). Their aim was to find out if agents could "evolve strategies for coping with opportunity costs (the possibility of a better deal outside the local network) as well as transaction costs (notably, the chance of being cheated by one's partner)" (Macy and Sato 2002, p.7215[2]).

The concept of 'trust'

2.3

The first thing to notice is Macy and Sato's concept of trust because it is different from what is usually known as the trust game. Macy and Sato's concept of trust makes use of the prisoner's dilemma, but does not define 'trust' as cooperation in a PD but as the decision to play the PD and not to exit. Furthermore, trust is assumed to be "symmetrical" (Macy and Sato 2002, p.7216[4]), i.e. mutual trust is needed to get to the PD: If only one of the partners refuses to play, both gain the payoff of the exit option. Figure 1 displays the payoff structure chosen by Macy and Sato. They also experimented with an exit option of zero but the main results stem from simulations with an exit option of −0.2.

Figure 1: Payoff matrix of the Macy and Sato's trust game

				player 2			
		cooperates		defects		exits	
	cooperates	0.7 – O	0.7 – O	−0.5	1.0 – O		
player 1	defects	1.0 – O	−0.5	−0.2	−0.2		
	exits					−0.2	−0.2

Opportunity and transaction costs

2.4

When agents interact, the probability that they find a partner that fits them (an agent *A* wants to trade products that agent *B* needs for products that *B* has) depends on the number of potential partners. This must not be true for every real world interaction, but one can easily think of situations where this is plausible. For example, if you are an organic farmer on a mountain pasture, it is surely harder to find an associate partner for your milk business in your neighbourhood, where everyone has milk of his own, than in the big city some miles away. To look for partners in a pool that is smaller than the number of potentially available partners is costly. The *O* in the payoff matrix of the trust game represents these costs. Macy and Sato work around complicated pairing mechanisms by introducing opportunity costs as a simple function of the size of the pool of possible partners *n*, the number of all agents in the population *N* and the heterogeneity of the agents *h* ($0 \leq h \leq 1$):

$$O_n = 1 - \left[\frac{n-1}{N-1}\right]^h \qquad (1)$$

There is not much to say about the implementation of transaction costs right now. What seems to be unambiguous is that transaction costs are not implemented as easy as opportunity costs. They "depend on the level of trustworthiness, which is endogenous to the model" (Macy and Sato 2002, p.7215[9]). This topic is discussed in detail in section 3.2.4.

Social structure of the model world

2.5

The model world consists of a fixed number of 1000 agents that are distributed randomly among a number of neighbourhoods that is altered from 10 to 100 as an exogenous parameter.[3] Agents do not have the ability to move willingly from one neighbourhood to another. Rather, for each agent, in each time step, there is a certain chance of being moved to a randomly chosen neighbourhood. This probability of being moved is the implementation of social mobility that Macy and Sato hypothesize to be important in the evolution of trust. It is equal for all agents and varied as a controlled exogenous variable between zero and unity. Agents that changed neighbourhoods are marked as newcomers for one time step. Parochial agents would not be willing to trust newcomers.

2.6

Social mobility is the first possibility to mingle with unknown agents. The second one is the global market. In each time step agents decide whether they want to interact in their neighbourhood or on the global market. When all agents have decided on the location of their interaction, pairs are formed to play the 'trust-game' described above in which they have to decide to trust (based either on a parochial decision rule or on signal reading) and cooperate.

The agents' decisions

2.7

As seen above, agents face three decisions: whether to cooperate, whether to trust, and whether to enter the open market. Each agent's tendency to act in a certain combination is stored in its "vector of three stochastic decisions" that we refer to as the 'action vector' in this text. Agents learn from experience and modify their behaviour according to the payoff. They learn, i.e. adapt the values in their 'action vector' according to past behaviour, either from reinforcement or social learning (Macy and Sato 2002, p.7216[8,9]).

Replicating the model

Two replications

3.1

As a replicated model is a model itself, we thought that it would be a good idea to verify our replication to reduce the chance of programming errors. Therefore we replicated Macy and Sato's model twice. One implementation was done in NetLogo and a second implementation was written in Fortran 95. While it might not be surprising that we chose NetLogo, our use of Fortran might be intriguing, but it had the major advantage that Fortran is very different from NetLogo. Thus it was less likely that we would make the same programming errors in both languages. To ensure that both models work in exactly the same way we did a test for identity of results. The source code of both implementations can be downloaded from http://pe.uni-bayreuth.de/?coid=21&q=detail&mid=125.

3.2

To do the test, we had to store a large number of random numbers in a file. All calls for random numbers in the implementations were replaced by a function that gives back the next random number from the file. Hence, both implementations were getting the same random numbers in the same order. In a second step the implementations were ensured to work isomorphicly. Here, we speak of an isomorphism between both implementations, if they are not only calling the same random numbers in the same order, but are also calling the same random number for exactly the same task. After this was successfully completed, both implementations produced exactly the same results. For that reason we could be confident that they did what we intended them to do, as it is rather unlikely — though not impossible — that we made the same programming error twice in two different programming languages.

Stumbling blocks of ambiguity and inconsistency

The action vector

3.3

Maybe the most confusing construct of the Macy and Sato's publication is the action vector. They state that "Action a is a vector of three stochastic decisions: trust (play or exit), trustworthiness (cooperate or defect) and location (neighborhood or market)" (Macy and Sato 2002, p.7216[9]).

3.4

The component 'location' is sufficiently explained in the quotation. It is simply an agent's probability to interact in the neighbourhood or on the global market. It seems clear that — analogously — 'trustworthiness' is simply an agent's probability to cooperate given himself and his partner are trusting each other in the sense described above. Though this understanding of trustworthiness is not in itself problematic, it has serious impacts on the understanding of transaction costs that are discussed below.

3.5

We did not find a better understanding of these two components of the action vector, but it should be noted that this causes major problems with Macy and Sato's text. There is a passage that says that "the combination of propensities for cooperation and entering the market gives each agent the ability to develop a conditional strategy for cooperation, depending on whether the relationship is likely to be ongoing (in the neighborhood) or one shot (in the market)" (Macy and Sato 2002, p.7216[3]). According to our interpretation of the propensities for cooperation and location, this is simply not true. The agent vector has only one component for a propensity to cooperate. Thus an agent cannot cooperate at iteration *t* with a probability of P_C if it is on the global market and with a probability different from P_C if it interacts with a known neighbour. This would be true only if the trustworthiness depended on whether the partner is a known neighbour or a stranger but this is not the case. The footnote on page 7216 states that the authors "could have assumed that an agent has two conditional probabilities to cooperate, depending on whether the partner is a neighbor or stranger. In this study, however, we focus on differences in trust that emerge across conditions even when there is no difference in trustworthiness (or propensity to cooperate)." We will return to this in the section on transaction costs.

3.6

What proved to be really tangling is the trust component. Given Macy and Sato's description that "Action *a* is a vector of three stochastic decisions: trust (play or exit), trustworthiness (cooperate or defect) and location (neighborhood or market)", it could be understood on the lines of the two components already discussed. If this was the case, then the trust component is nothing but the probability of trusting one's partner. But if so, how could there be agents that are parochial and others that rate their partner based on signal reading? Macy and Sato explicitly write that agents "decide whether to trust the partner (based either on the partners location or signal)" (Macy and Sato 2002, p.7217[1]). Furthermore, they state that they measured trust as "the manifest behavior (the rate at which agents actually trust and cooperate in the market and neighborhoods)" *and* "the underlying rule (the average propensity to base trust on the partner's social cues rather than social proximity)" (Macy and Sato 2002, p.7217[5], emphasis by the authors). This conflicts with an interpretation of the trust component as a mere probability to trust.

3.7

The simplest interpretation we could find is that the trust component is the probability of choosing one of two decision strategies on how to decide on trusting. The strategies at stake are parochialism and signal reading. To decide on trusting is thus a sequence that starts with the choice between signal reading and the parochial decision rule. A random number between zero and unity is drawn and compared with the agent's trust component. If this propensity is larger, then the agent reads his partner's signal, otherwise she acts parochially. We discuss below the implementation of signalling in detail. Parochial agents that go for the open market do not trust anyone. If they decide to interact in their neighbourhood they do not trust newcomers. Those parochial agents that interact in their neighbourhood and are newcomers themselves do not trust anyone.

3.8

It is obvious that this interpretation of the action vector's 'trust' component goes way beyond the description in the text that merely says, "trust (play or exit)" (Macy and Sato 2002, p.7216[9]). Nevertheless we did not find an alternative that allows agents to apply different decision strategies. Figure 2 illustrates our understanding of the agents' decision on trusting in Macy and Sato's model. There is no information on whether or not a parochial newcomer trusts its new neighbours or not. Our experiments indicated that we come a bit closer to Macy and Sato's results if we assumed that they are distrusting. This is plausible to some extent, in the same way as newcomers are strangers to the neighbours, the newcomer's new neighbours are strangers to the newcomers.

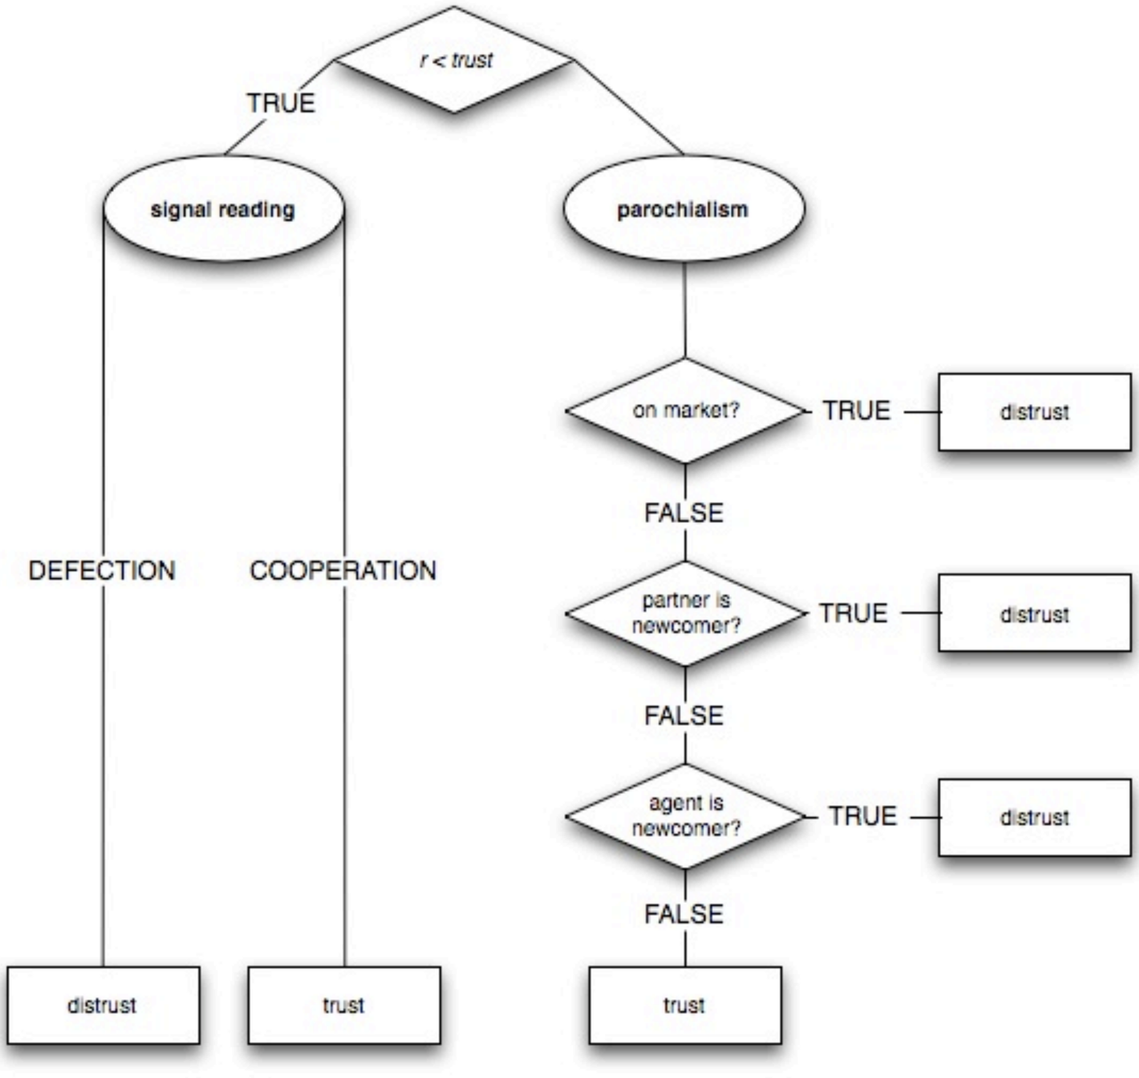

Figure 2. How agents are assumed to decide on trusting

Social learning

3.9

Social learning is well described in the Macy and Sato's text. It is explicitly stated that "Agents identify as a role model the neighbor with the highest cumulative payoff" (Macy and Sato 2002, p.7126[8]). The one thing that is noteworthy is that there is no mention of a discounting of payoffs. It seems natural to expect agents to rate payoffs in an inversely proportional relation to the amount of time that has gone by since they were gained. This is especially true because otherwise present payoffs would have less impact with every iteration ticking by. However, Macy and Sato do not seem to have implemented this and if they did so, they made replication impossible because there is no hint to the amount of discounting in the past.[4] We assume therefore cumulative payoffs to be calculated by absolute values in their model.

Reinforcement learning

3.10

Reinforcement learning is implemented by a Bush-Mosteller learning algorithm. On first view the associated formulae given in the text look perfectly plausible (Equation 2). They imply that if in a binary decision between action A and action B, the agent chose A at time step t, the probability of repeating A at time step $t + 1$ depends on whether or not and to what extend the payoff that was gained is above a threshold of zero. If, for example, in t the agent's propensity to cooperate was 0.5 and he gained a payoff of 0.5 then the probability of cooperating in $t + 1$ increases to 0.75 and therefore the probability of defecting decreases to 0.25. If he did not cooperate and gained 0.5, then the probability of defecting rises to 0.75 and the propensity to cooperate in the action vector falls to 0.25.

$$P_{a,t+1} = \begin{cases} P_{a,t} + \left(1 - P_{a,t}\right) \cdot \pi_{a,t}, & \pi_{a,t} \geq 0 \\ P_{a,t} + P_{a,t} \cdot \pi_{a,t}, & \pi_{a,t} < 0 \end{cases} \quad (2)$$

3.11

As clear as this understanding of reinforcement learning is, problems come up if we think about the implications that this learning mechanism has, given the payoff structure of Macy and Sato. Consider a population of agents in which the mean propensity to cooperate is on a very low level. Now, two agents that decided to trust each other and play the PD are likely to end up in mutual defection. Both earn -0.2, which is less than zero, and therefore the agents' probability of cooperating in the next time step increases. Given the combination of reinforcement learning and the payoff structure chosen by Macy and Sato, the mean propensity to cooperate must stay on a substantial level. To give an idea of what "substantial" means: an agent with $P_C = 0.01$ that ends up in mutual defection cooperates in the next time step with a probability of 0.208. What was said about P_C is true likewise for the 'location' component of the action vector. This is true, because the costs of exiting are -0.2 in all the simulations that Macy and Sato's main results stem from. Therefore, it is an a priori truth that the levels of cooperation and market size, understood as means over the population, cannot fall to near zero. The probability to read signals does not have any direct influence on the payoff. Therefore the mechanism described for the other components does not work for the signal reading component. It is, however, influenced indirectly due to the fact that there are always agents on the open market. Those agents only have the chance to earn more than the outcome of the exit option if they read signals. Thus, signal reading is also ensured to be substantially larger than zero. All in all, a nice world. But if one wants to show that a population can move to a state of high trust, cooperation and market size, it does not seem to be a very good idea to make use of assumptions that exclude situations without substantial amounts of these.

3.12

The combination of the chosen reinforcement algorithm and the payoff matrix described above has other noteworthy implications. Again, imagine an agent in a population of very low cooperation. Its probability to cooperate is 0.01 in t. It is likely, that it ends up in mutual defection, gains a payoff of -0.2 and therefore its propensity to cooperate rises to 0.208. The interpretation is that zero works as a threshold that determines whether or not the agents rate a payoff as a positive or negative experience. Actions that lead to positive payoffs become more likely to be repeated and those that lead to negative outcomes are repeated with a lower probability in the future. So, in the situation in which defection is common, a cooperating agent gains the even worse outcome of -0.5. Thus, the agent in mutual defection refines its propensities in a way that is likely to make its state even worse. This implication does not sound plausible anymore. One could try to justify this behaviour by saying that agents do not know about the actions and payoffs of others. However, this argument is not consistent with social learning, which assumes agents have perfect knowledge about each other's propensities and payoffs.

Signalling

3.13

Macy and Sato suppose that "reliance on signaling, rather than parochial relations, allows the system to move to a Pareto superior equilibrium" (Macy and Sato 2002, p.7214[last]). Thus signalling plays a crucial role in their work. Nevertheless the authors did not describe in detail how the agents send and interpret signals. The most informative part is the following:

> Trustworthy agents have an interest in signaling in a way that is difficult for cheaters to mimic, but we assume that these are imperfect. The reliability of the signal depends on the agents level of trustworthiness. The stronger the agents commitment to honest (or dishonest) behavior, the less likely their signal will be misleading. Signals from ambivalent agents are highly unreliable (no better than a coin toss) (Macy and Sato 2002, p.7216[3]).

As trustworthy agents are described to be "imperfect" in "signaling in a way that is difficult for cheaters to mimic", one could conclude that agents that intend to defect are in some way able to mimic cooperative agents. But is this true? If we take a look at the passage that describes the actions committed by all agents in each iteration, we learn that they "decide whether to cooperate" and "give the appropriate signal" (Macy and Sato 2002, p.7217[1]). It is not clear, what "appropriate" means in this context. A signal of defection may be 'appropriate' for an agent that intends to defect in the sense that it is the correct signal. On the other hand, from the agent's point of view a signal of cooperation is perfectly 'appropriate' for it to maximize her payoff if she intends to defect. Though the passage on signalling is ambiguous, we are fairly sure that agents are not able to mimic signals. The hint is in the conclusion, where Macy and Sato discuss opportunities of further research and mention the possibility to "allow agents to learn how to fake signals of trustworthiness" (Macy and Sato 2002, p.7220[3]). This sentence would not make sense if agents in this study were able to mimic certain signals. We therefore assume that agents always send a correct signal.

3.14

What is 'imperfect' about signalling in the model is merely the reception of the signal. For each agent, there is a certain probability that its partner misunderstands the signal sent. This probability depends on the trustworthiness of the signal-sending agent. 'Trustworthiness' is the agent's probability to cooperate and therefore a number between zero and unity. But, it has to be clarified in which way the trustworthiness of the sender and the correct understanding of the signal-receiving agent are related. The easiest possibility is a proportional relation which in the simplest case would mean that both probabilities are simply on par. Unfortunately, this would imply that an agent with a very low level of trustworthiness, i.e. likely to defect and to signal defection, would also be likely to be misunderstood and therefore likely to be interpreted as intending to cooperate. If this were the case, it would be a really big advantage for defecting strategies and it would be hard to imagine how cooperation might emerge in a population.

3.15

The most plausible interpretation of what Macy and Sato say about imperfect signalling can be mapped by a u-shaped function that describes the relation between probability of being misunderstood and trustworthiness. There is at least one point on the curve that we can be very sure about. This point can be derived from Macy and Sato's statement that "signals from ambivalent agents are highly unreliable (no better than a coin toss)" (Macy and Sato 2002, p.7216[3]). It is reasonable to conclude that this means that an agent, whose trustworthiness is 0.5, has a chance of 0.5 of being misunderstood. The u-shape fits well to the passage in which it is stated: "The stronger the agents commitment to honest (or dishonest) behavior, the less likely their signal will be misleading." Obviously, we know that the function's minimum is at 0.5 and the maxima are to be found at 0 and 1 (agent always defects, agent never defects). We also know that the probability of being understood correctly is 0.5 at the minimum, but we do not know what values it takes when trustworthiness is minimal or maximal. It is likewise unclear whether the relation is linear, concave, convex or of whatever other type. Without being sure that Macy and Sato made the same assumption, we assumed that they assumed signalling to be perfect (i.e. a probability of one to be understood correctly) when trustworthiness is either zero or one and a linear relation, as this is the most salient solution in our opinion. Figure 3 displays the function we used (black) to describe the relationship between an agent's trustworthiness and its probability of being understood correctly. Some other possible functions are also plotted (dashed grey).

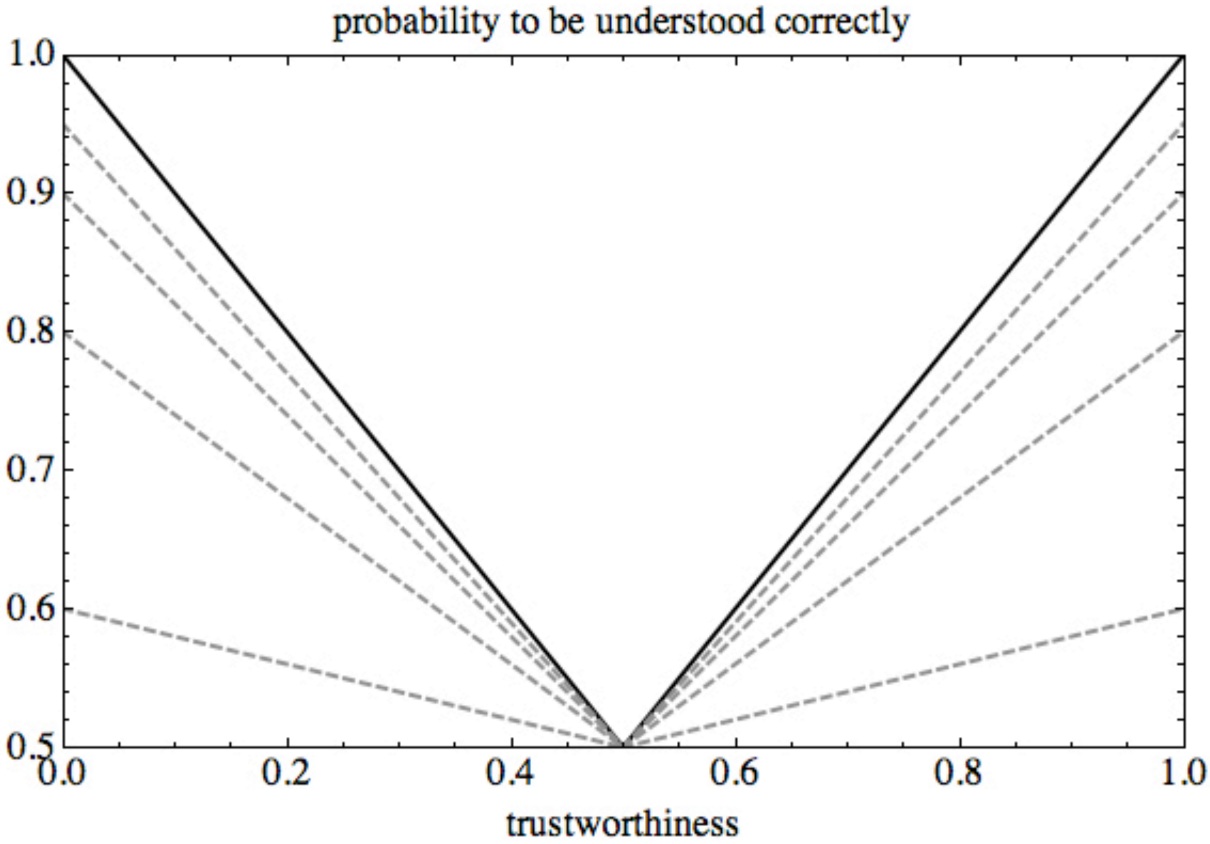

Figure 3. Imperfect signalling

Transaction costs

3.16

Macy and Sato would surely agree that the area of conflict between parochial and non-parochial strategies is very important for their model. On the one hand parochialism "impose[s] opportunity costs" but on the other hand parochial agents "minimize transaction costs associated with the risk of being cheated by strangers" (Macy and Sato 2002, p.7214[3]). It seems natural to expect Macy and Sato to deliver a comprehensive description of their concept of opportunity and transaction costs. The fact is, they do not.

3.17

The clearest statement they make is the one where they say, "the transaction cost of exchange depends on the level of trustworthiness, which is endogenous to the model" (Macy and Sato 2002, p.7216[5]). As the authors do frequently speak of the "risk" that is correlated with interaction between strangers, (Macy and Sato 2002, p.7214[3], p.7219[4], p.7219[8]) this implies less trustworthiness in interactions between strangers than in those between neighbours.

3.18

One could easily operationalize this by using a case specific trustworthiness component in the action vector that is higher in neighbourhood interaction than in interaction with strangers. As agents are able to distinguish strangers from neighbours when they decide on trusting, it would be perfectly plausible that their decision on cooperation can depend on whether or not they know their partner. Unfortunately Macy and Sato exclude this possibility by stating, they "could have assumed that an agent has two conditional probabilities to cooperate, depending on whether the partner is a neighbour or stranger" and that they analyze the emergence of trust "even when there is no difference in trustworthiness (or propensity to cooperate)" (Macy and Sato 2002, p.7216[footnote]).

3.19

As the text provides good evidence that trustworthiness does not depend on whether or not the agent knows the partner, how could trustworthiness in interactions with strangers be lower? Similar to the implementation of opportunity costs, the payoffs of agents could be subtracted by a fixed amount of transaction cost. Although Macy and Sato's statement that transaction costs are endogenous excludes this option, we tested it. We tried to implement transaction costs of 0.2 in this way but the results were not interesting. Independent of mobility rate and number of neighbourhoods, each run generated very similar results.

3.20

There are several ways in which transaction costs might occur endogenously. If transaction costs arise endogenously and the agents' propensity to cooperate does not discriminate between known agents and strangers, trustworthiness could be higher in neighbourhood interaction because interactions are repeated. There are hints in the text that it is the "ongoing" character of local interaction that causes the minimization of transaction costs (Macy and Sato 2002, p.7216[3], p.7219[5]). But how can a relation in the model be ongoing if in "a new iteration" agents "randomly choose a partner from the available pool (either the market or the neighborhood)" (Macy and Sato 2002, p.7216[last])?

3.21

Passages, describing relations in the neighbourhood as "embedded", point in the same direction (Macy and Sato 2002, p.7214[3,4,5,footnote1], p.7215[2,4,8], p.7216[5], p.7217[6], p.7219[2,3,7], p.7220[1]). This embeddedness is described as being characterized by "stability and transitivity" (Macy and Sato 2002, p.7214[footnote1]) but there are no hints on how this could have been implemented.

3.22

A difference in transaction costs on the market and in the neighbourhoods could also be due to a correlation between the agents' propensity of entering the market and their propensity to defect. If agents that are likely to defect were more likely to enter the market, then the risk of being cheated would be lower in neighbourhood interaction.

3.23

Finally, we have to state that we could not find a satisfying way to implement transaction costs. As described in detail in the next section, the comparisons of results from the replicated and the original model indicate that the replicated model lacks the implementation of transaction costs. Thus it is unlikely that we simply failed to understand how they are caused by the other elements of the model.

Comparison of results

Experimental design

4.1

Before reporting on the comparison of results from our replicated and Macy and Sato's original model, some information on Macy and Sato's experimental design are useful. All results discussed here stem from simulations with 1000 agents, costs to exit of 0.2 and heterogeneity of 1. Mobility rate is varied from 0 to 1 in steps of 0.1 and neighbourhood size from 10 to 100 in steps of 10. All combinations of these settings are repeated 20 times. The parameter settings are clearly described in the text but there are some questions left on how the neighbourhood size is exactly altered.

4.2

The authors state that they "varied the size of neighborhoods from 10 to 100 (in a population of 1,000), in 10 steps of size 10" and that agents are "randomly allocated to positions in a social network" (Macy and Sato 2002, p.7217[3], p.7215[11]). Unfortunately they do not tell us what numbers of neighbourhoods were actually used to reach the desired neighbourhood size. Below you see a table in which the numbers of neighbourhoods needed to obtain the desired and expected sizes of neighbourhoods are calculated. In some cases, where the size cannot be represented by an integer, we ran a share of simulations with the next smaller and a share with the next larger whole number of neighbourhoods.

Measurement

4.3

Macy and Sato tell us that they "allowed agents to play for 1,000 iterations before recording results for an additional 1,000 iterations" and that "this was more than sufficient time for the population to move to either a parochial equilibrium [...] or to a non-parochial equilibrium" (Macy and Sato 2002, p.7217[2]). The term 'equilibrium' indicates that in their simulations the simulations end in a situation in which the agents' action vectors no longer change. Unfortunately the combination of the chosen Bush-Mosteller learning algorithm and payoff matrix only allows for three stable situations:

1. All components of the action vectors of all agents are one and all agents gain a payoff above zero.
2. All components of the action vectors of all agents are zero and all agents gain a payoff above zero.
3. All agents gain a payoff of zero.

We can ignore the third option because the chosen payoff structure does not allow for outcomes of 0 alongside the situation in which an agent fails to find a partner. It is not possible that all agents end up without a partner because there are only up to 101 (100 neighbourhoods and the market) pools of partners, which would all have to consist of no more than one agent, which is not compatible with a fixed population size of 1000. In addition, Macy and Sato report results that are more diverse than the first two possibilities. Thus, when speaking of an 'equilibrium' they must have something different in mind.

4.4

A different concept of equilibrium does not demand stable strategies but is satisfied if there is a sequence of strategy combinations that reappears again and again. It seems unlikely that Macy and Sato used those 'cyclic equilibria' without mentioning it. Therefore we think that they used a much simpler concept and measured the relevant data during the last 1000 of 2000 time steps and calculated mean values. We followed this procedure in our simulation. It remains unclear how the authors can be sure that "this was more than enough time" (Macy and Sato 2002, p.7217[2]).

Effects of mobility rate

4.5

Figure 4 shows plots from Macy and Sato's and our replicated model that illustrate the distribution of trust in strangers in dependence on the rate of mobility. Figures from both implementations show a unimodal distribution if mobility is very low and a bimodal distribution if mobility is increased. Compared to the original model, the modi of the distributions from the replicated model are shifted to the right. This means that the 'equilibria' in the replicated model indicate more trust than those in the original model. The replicated model qualitatively reproduces the positive effect of moderate rates of mobility on trust in strangers.

4.6

There is, however, one major difference. While the plot from the original data indicates a break down of trust in strangers if mobility is set to 1, this is not the case in the replicated model. This suggests that our above suspicion that we could not replicate Macy and Sato's implementation of transaction costs is correct.

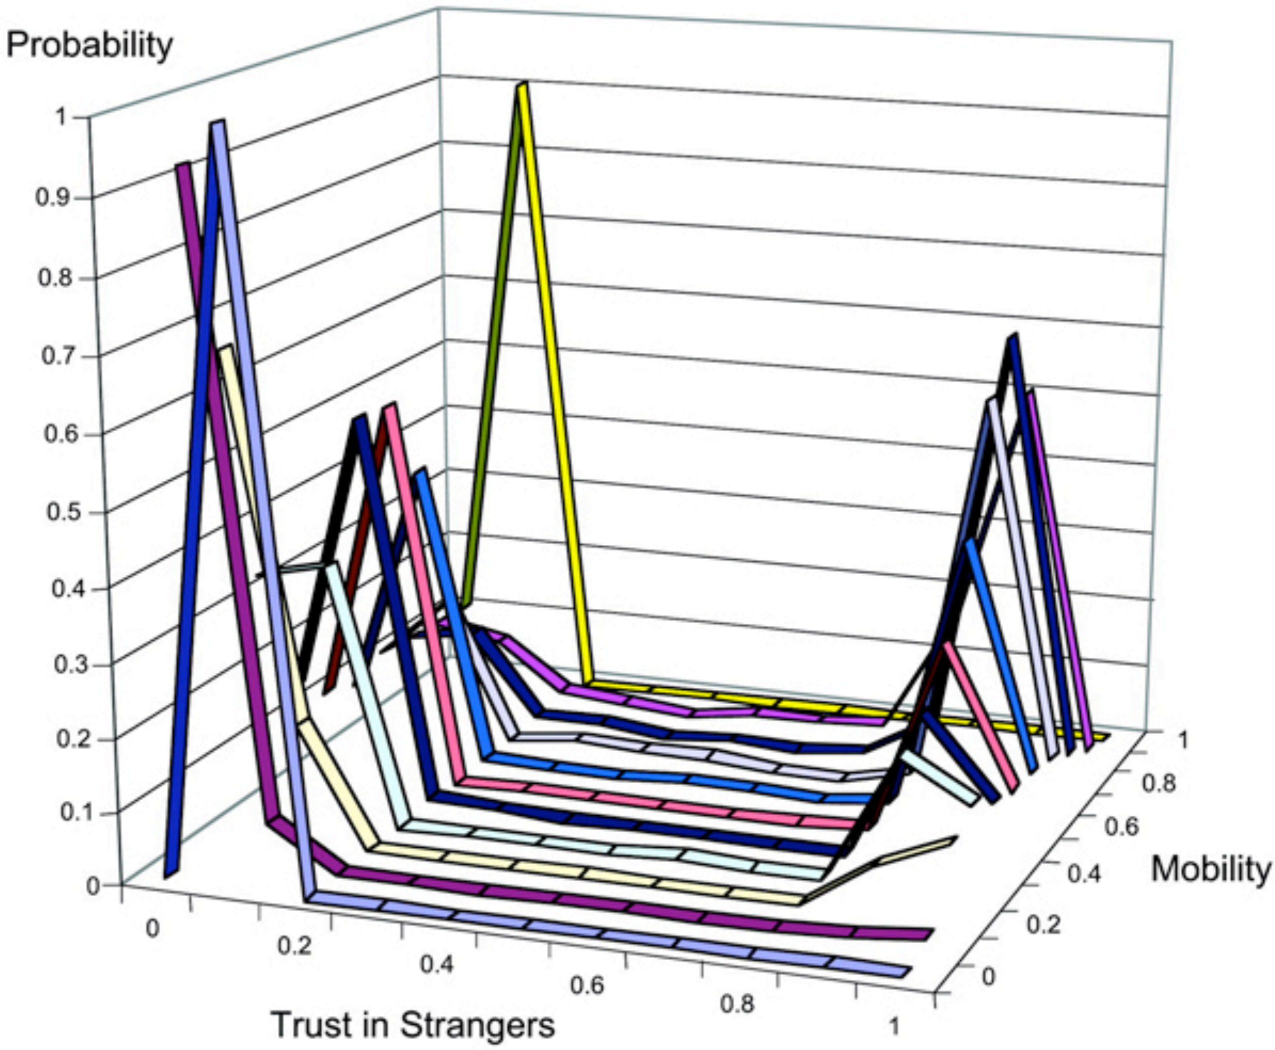

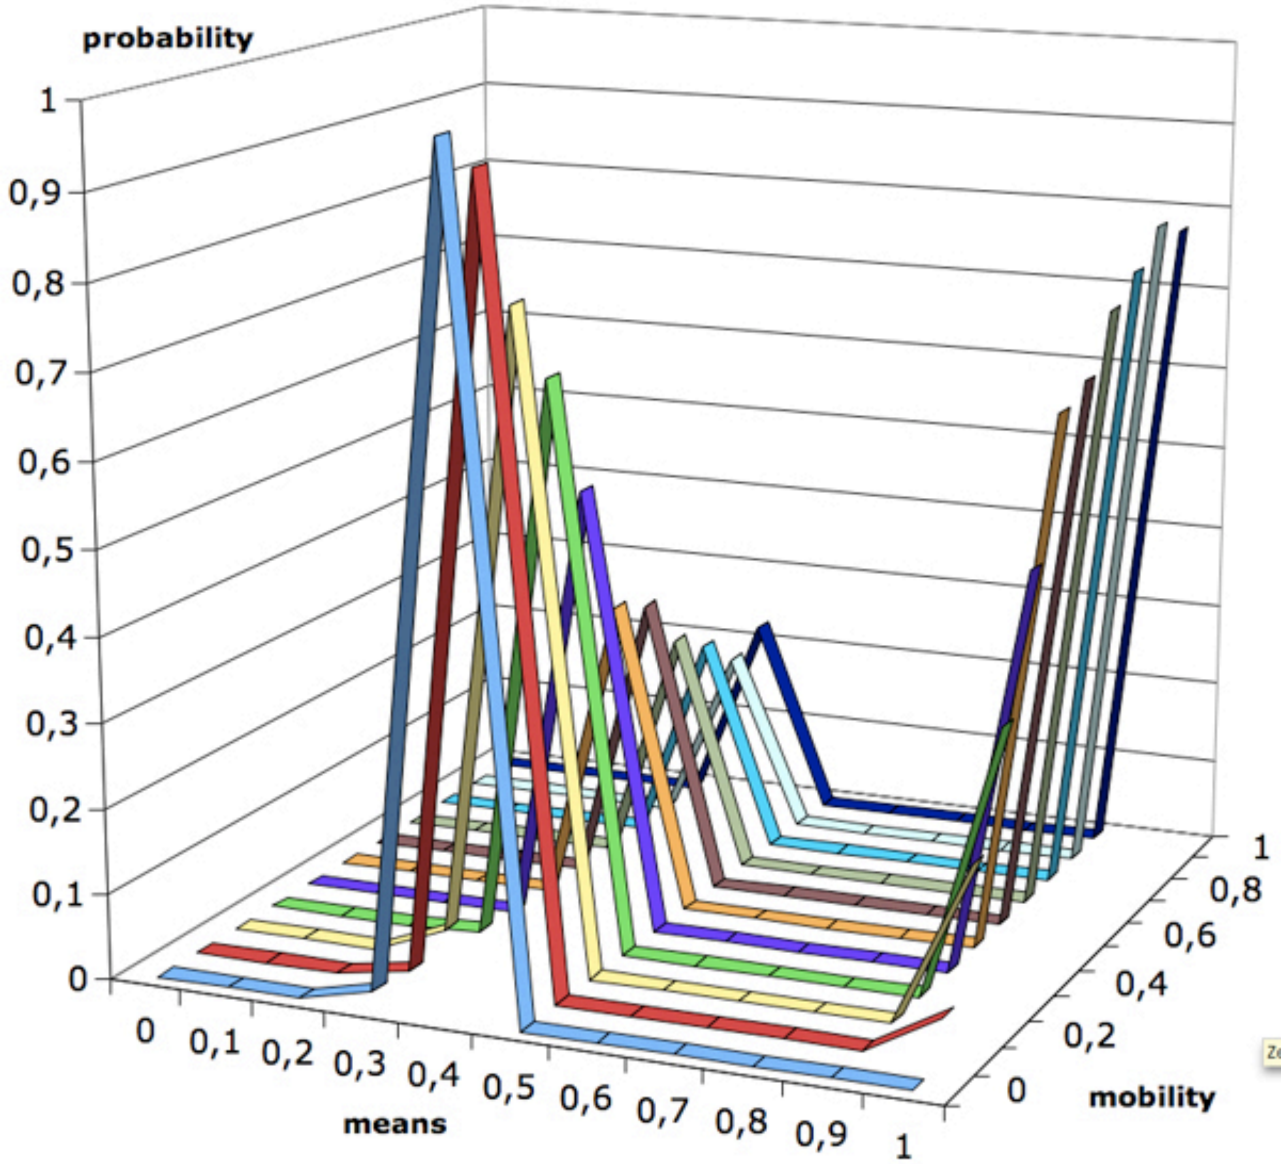

Figure 4. Distribution of trust in strangers by mobility rate in the original (top) and the replicated model (bottom), with N = 1,000, n = {10..100}, X = P, h = 1)

4.7

Figures 5 and 6 show the influence of social mobility on indicators of market interaction while Figures 7 and 8 display the relation between mobility rate and local interaction indicators. The curves of the original and the replicated model do not look very similar, but our results roughly verify Macy and Sato's finding that "local trust steadily declines, whereas trust in strangers increases" (Macy and Sato 2002, p.7218[2]). Similar to the analysis of the influence of social mobility on trust in strangers, we cannot find that "at very high mobility rates, both local and global trust collapse" (Macy and Sato 2002, p.7218[2]).

4.8

To sum up, it seems that our model captures the positive effect of an increasing social mobility rate on trust in strangers that is generated from the expansion in opportunity costs for parochial agents. What seems to be missing is the reverse effect of transaction costs that Macy and Sato assume to be raising when the rate of social mobility is very high.

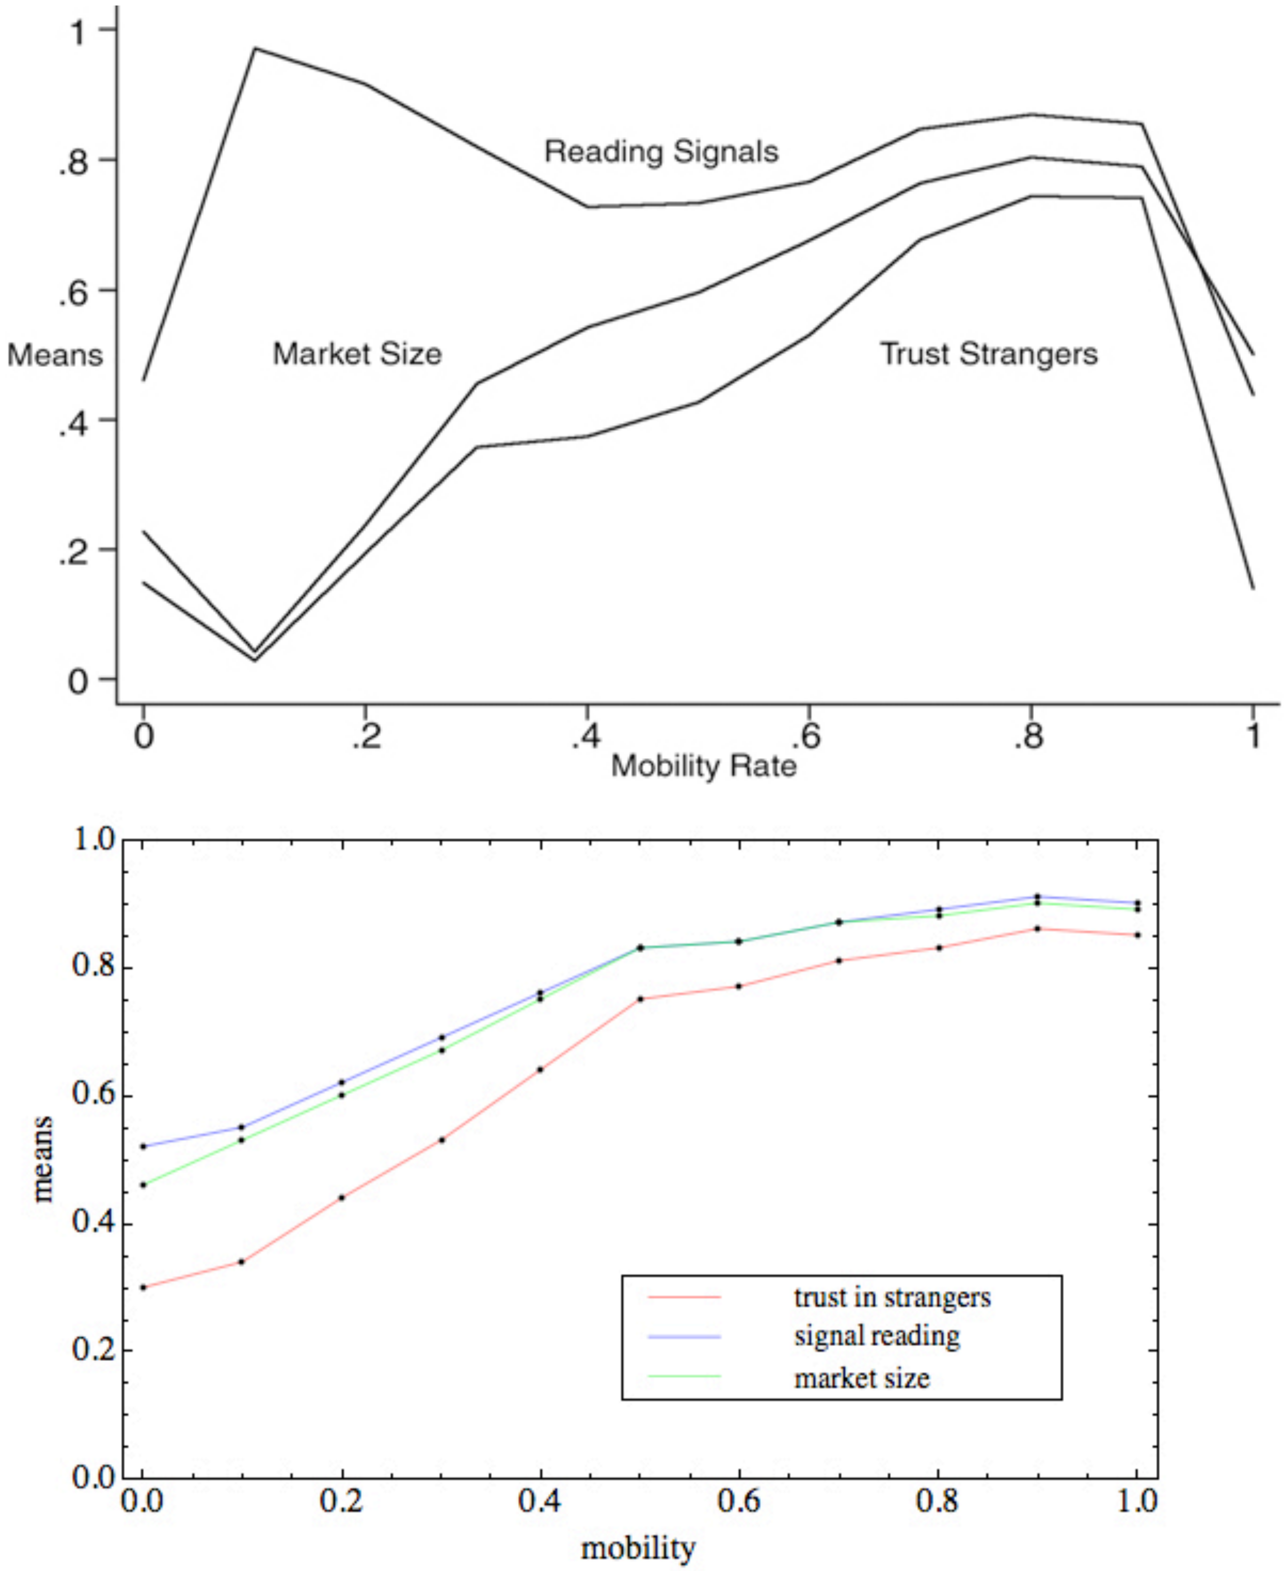

Figure 5. Expected value of three indicators of market interaction in original (top) and replicated model (bottom), based on 200 observations at each level of mobility, with N = 1,000, n = {10..100}, X = P, h =1

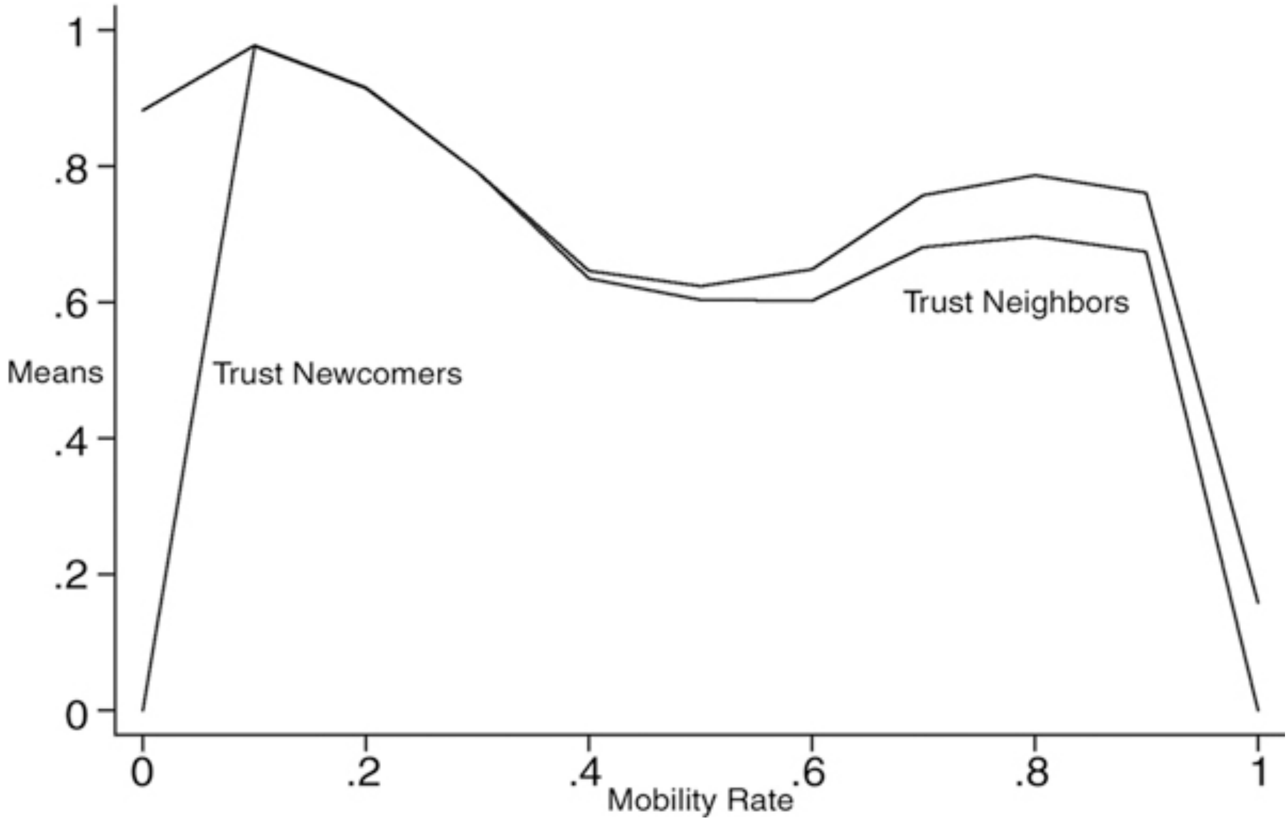

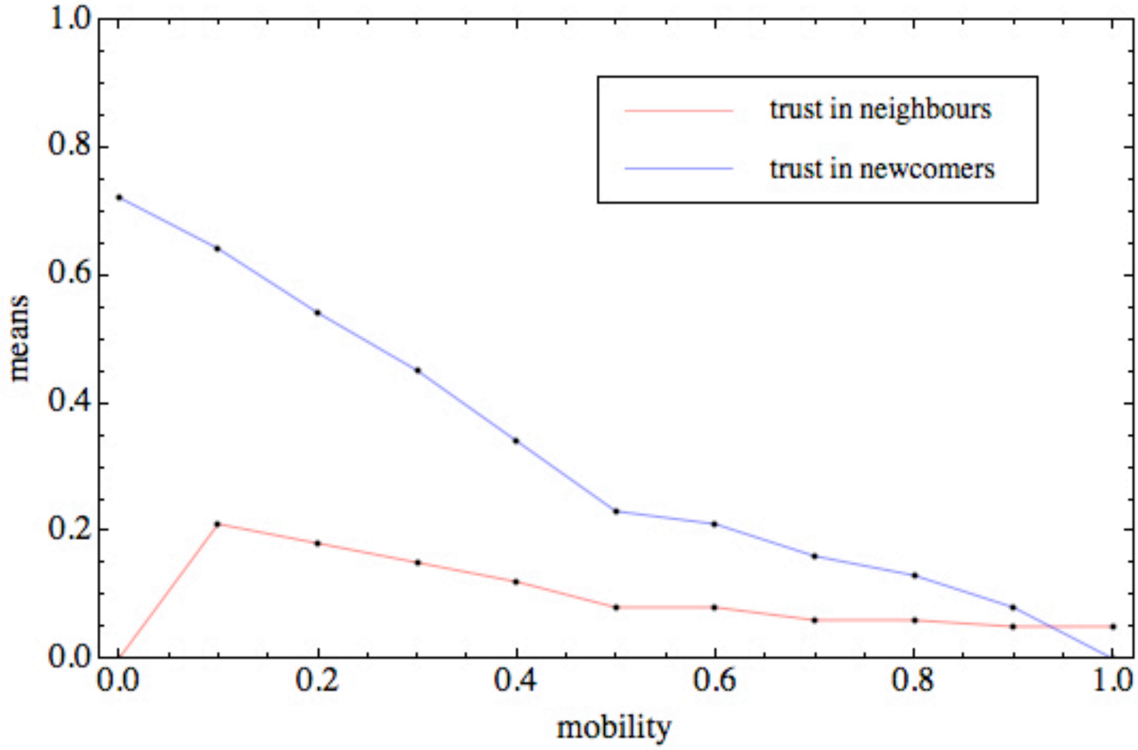

Figure 6. Expected value of three indicators of local interaction in original (top) and replicated model (bottom), based on 200 observations at each level of mobility, with N = 1,000, n = {10..100}, X = P, h = 1

Effects of neighbourhood size

4.9

Macy and Sato state, that "all five measures [of trust and market formation] increase quickly as neighborhood size increases from 10 to 20" but found "little or no effect of neighborhood size" for neighbourhood sizes larger than 20 (Macy and Sato 2002, p.7218[5]). Our results do not look very much alike and suggest a small positive effect on indicators of market interaction and a small negative effect on indicators of local interaction. Figures 7 and 8 show results from the original and the replicated model. It should be noted that the effect of neighbourhood size depends critically on the chosen concept of opportunity costs.

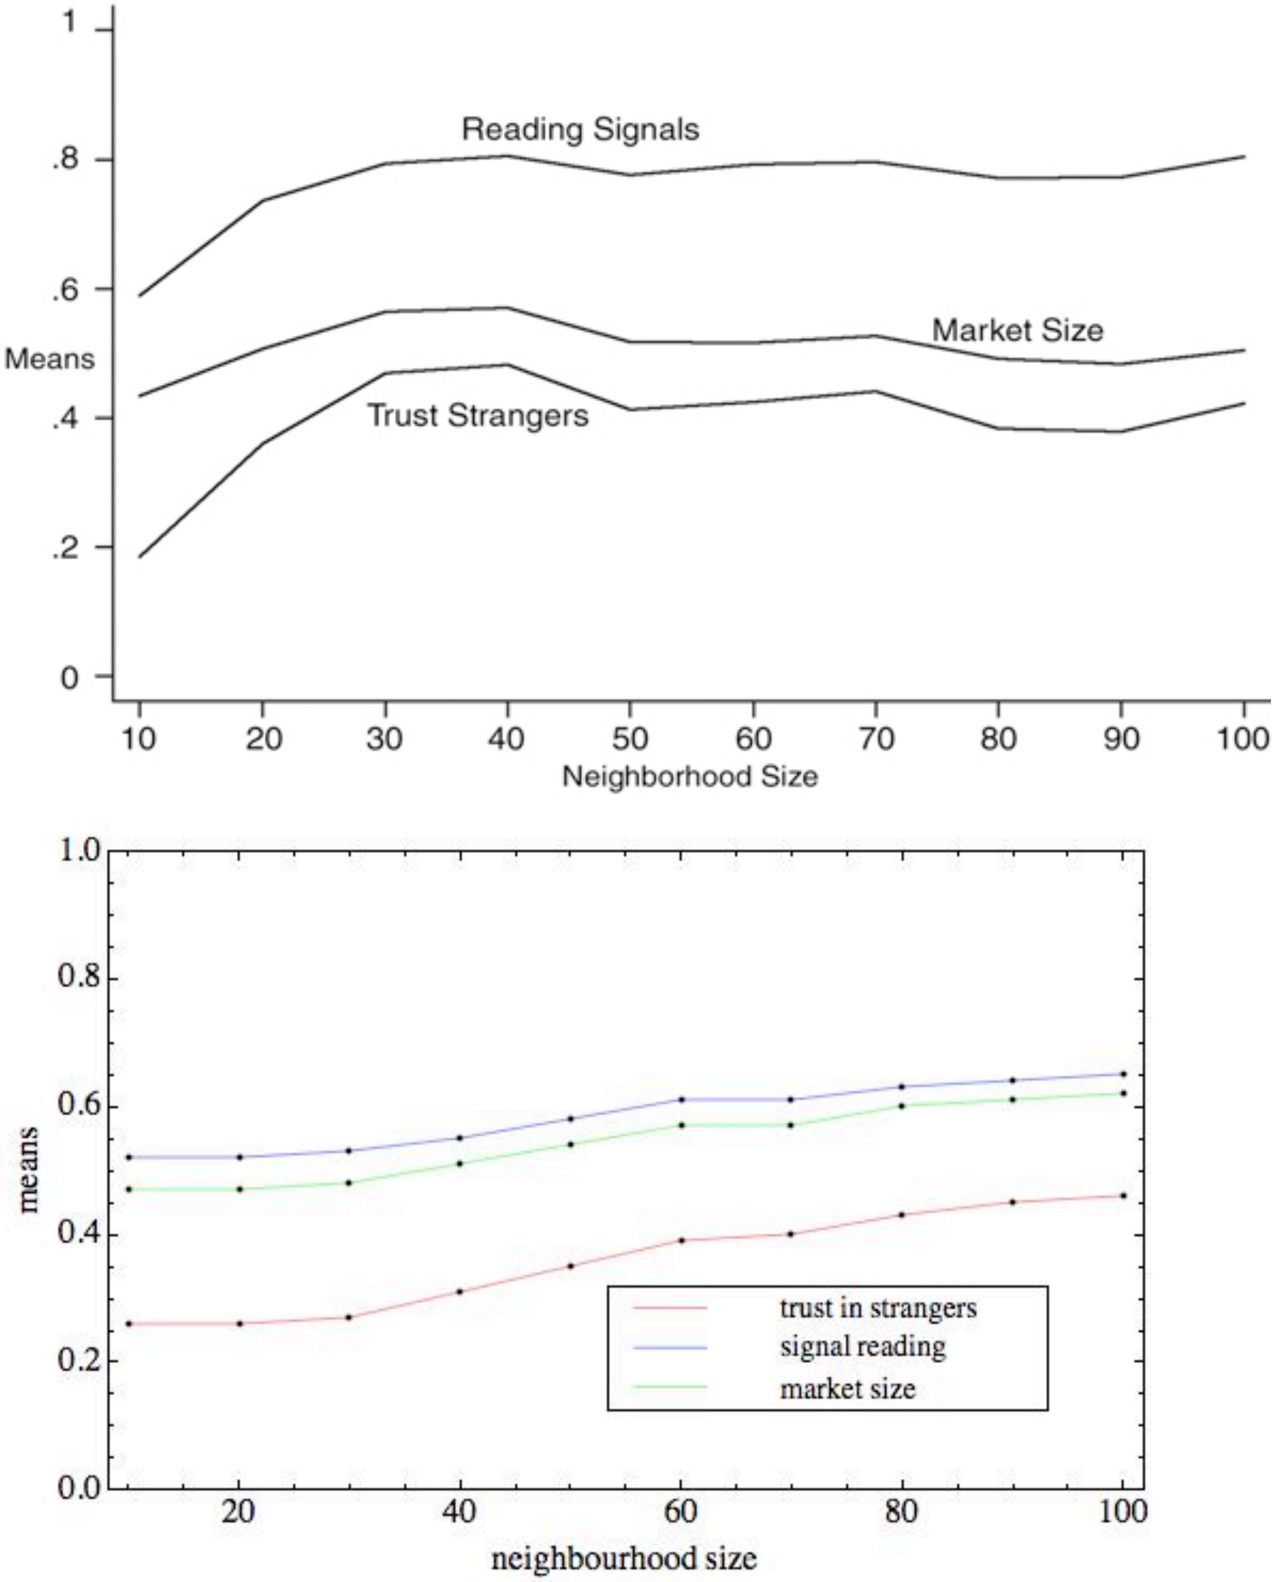

Figure 7. Expected value of three indicators of market interaction in original (top) and replicated model (bottom), based on 220 observations at each level of mobility, with N =

1,000, n = {10..100}, X = P, h = 1

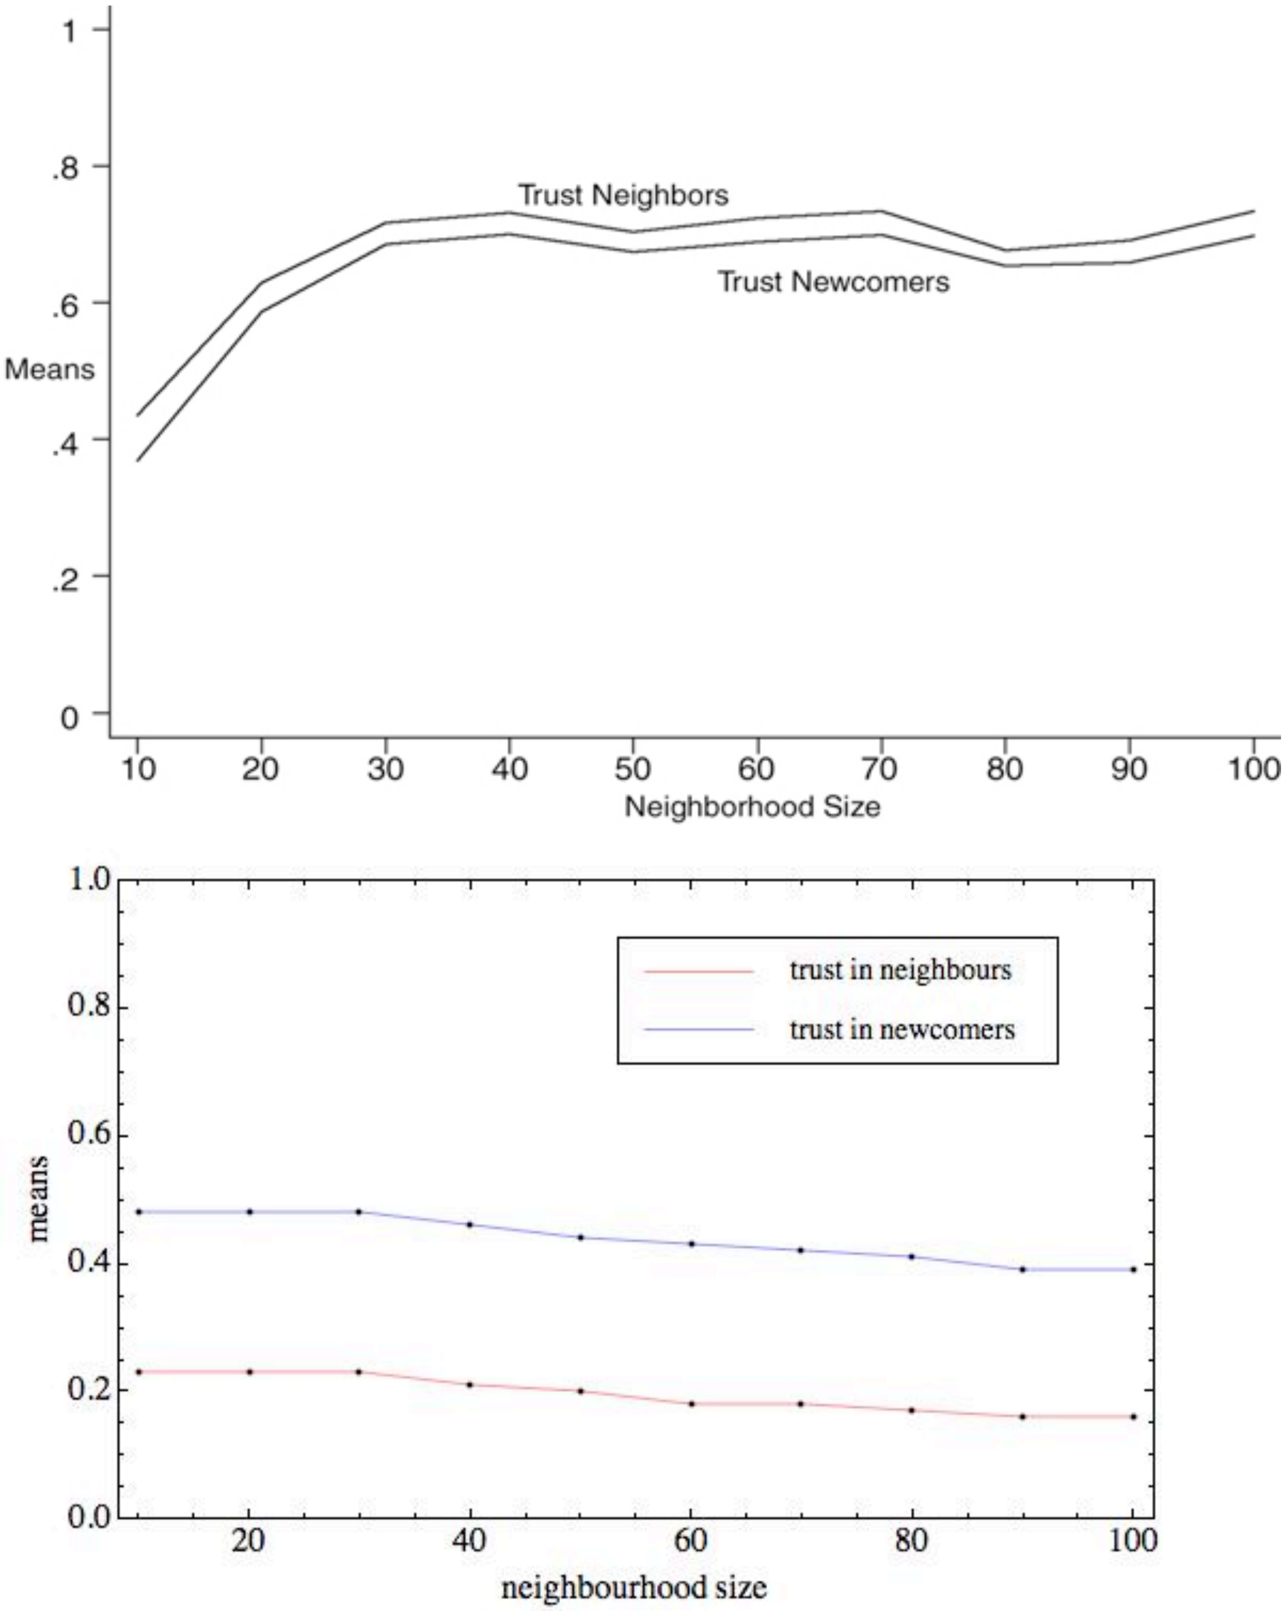

Figure 8. Expected value of three indicators of market interaction in original (top) and replicated model (bottom), based on 200 observations at each level of mobility, with N = 1,000, n = {10..100}, X = P, h = 1

4.10

The implementation of opportunity costs is straightforward. Though the formula is unambiguous, it is interesting to look at it carefully. Figure 9 illustrates the amount of opportunity costs in dependency on the size of the pool that the partner of an agent is drawn from and the heterogeneity of the agents demands and offers. Macy and Sato's main results stem from simulations with heterogeneity of one. For this amount of heterogeneity, we find that in comparison to the possible sizes of the global market, the applied neighbourhood sizes between 10 and 100 have little influence on the opportunity costs. This becomes even worse because in the initial situation there are about 500 agents on the global market. Thus, opportunity costs on the global market are about one fifth in the initial situation even if mean neighbourhood size is at its maximum value. It is, therefore, not surprising that Macy and Sato could were unable to find a significant influence of neighbourhood size on trust in strangers and we would hardly support Macy and Sato's conclusion that "These results suggest that differences in parochialism between the U.S. and Japan may reflect the higher rates of mobility in the U.S. but not differences in the structures of local interaction" (Macy and Sato 2002, p.7220[1]). One could of course think of a formula for opportunity costs that causes a much greater impact of differences in neighbourhood sizes between 10 and 100. Thus Macy and Sato's conclusion can only be justified if they give reasons why the chosen formula is the most adequate.

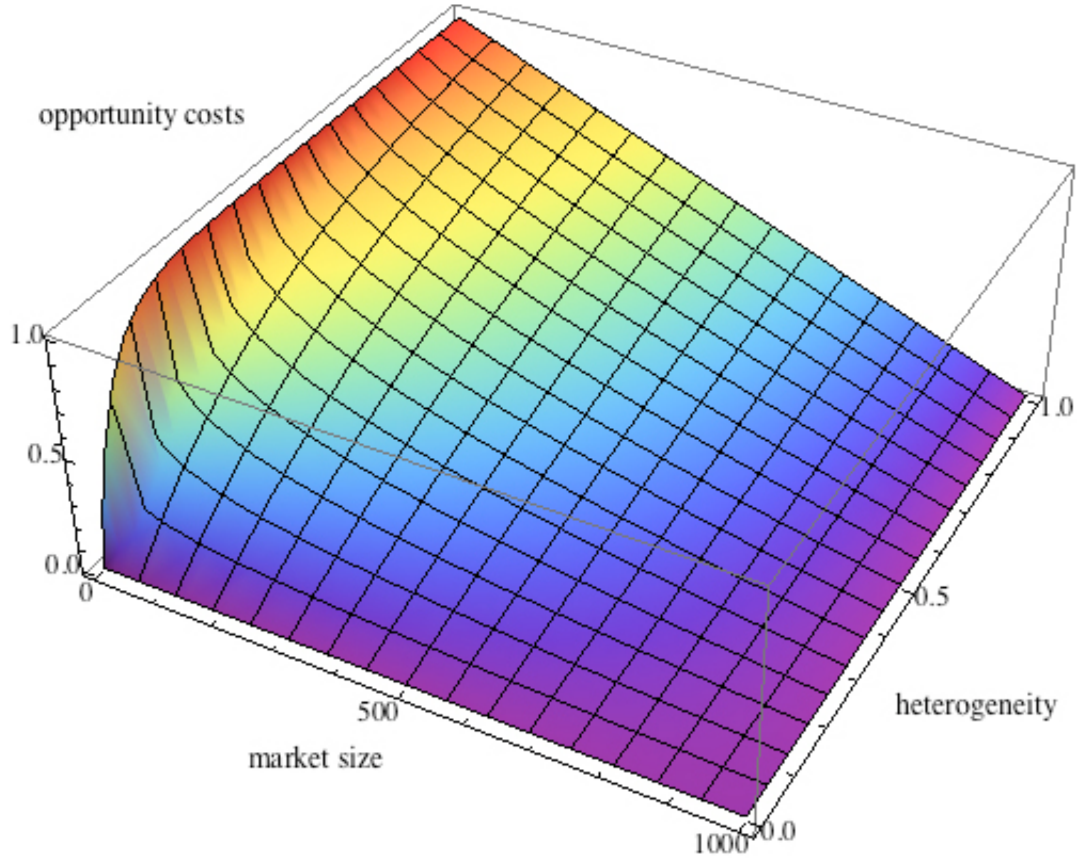

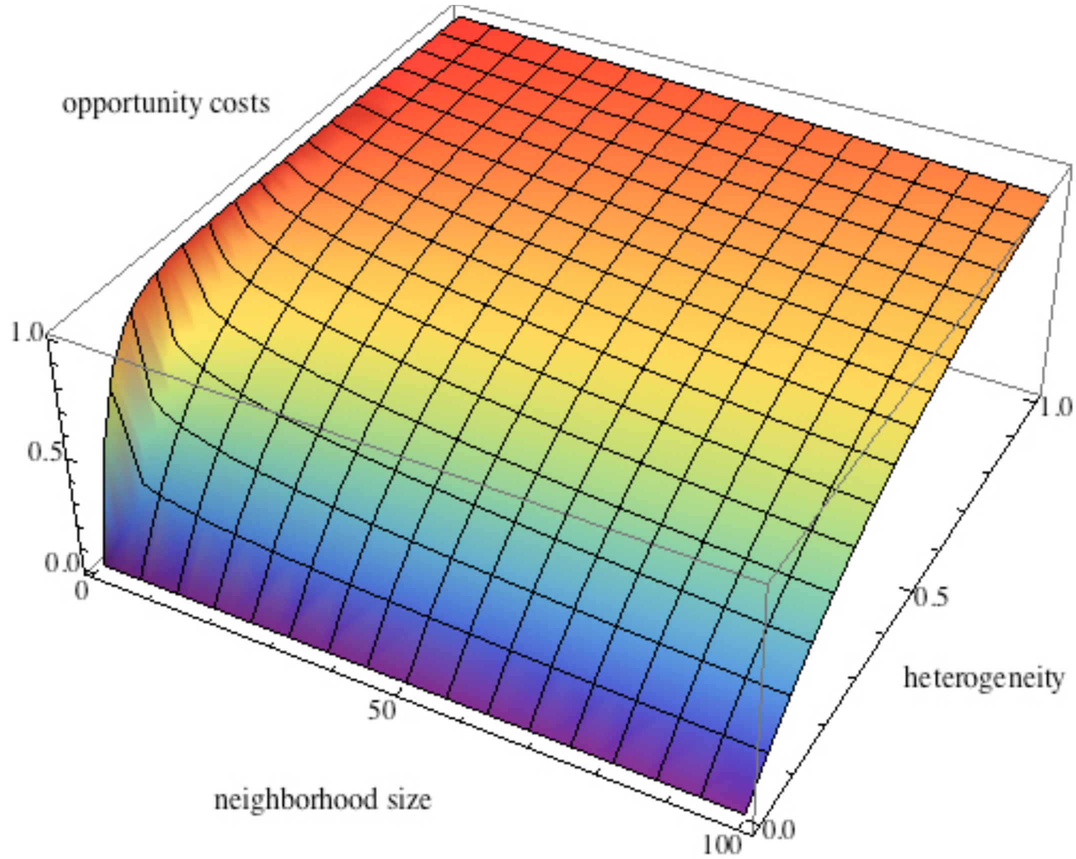

Figure 9. Opportunity costs in dependency on neighbourhood (top) and market size (bottom)

Concluding remarks

5.1

As seen, replication is not an easy task. It is especially messy when the description of the model to be replicated is fragmentary and no source code is provided. Axtel et al. pointed at three decreasing levels of replication: "numerical identity in which the results are reproduced precisely, distributional equivalence in which the results can not be distinguished statistically, and relational equivalence in which the qualitative relationships among the variables are reproduced" (Axtell 1996). To state it frankly: Our trial in replicating Macy and Sato's model did not even succeed in fully replicating the results qualitatively.

5.2

This is probably mainly due to problems with the transaction costs. Either there are unpublished hidden parameters in the model that affect the outcomes in the direction of the results that Macy and Sato published, (in that case we failed to re-invent them), or Macy and Sato think that the transaction costs are somehow built-in by the very design of model. In the latter case it does not make sense to look for a transaction costs component, which, for instance, is somehow similar to the opportunity cost component of the model. Though carefully trying to understand the interplay of the model's constituent components, we, then, would nevertheless have failed 'to see' that the transaction costs are already built-in. But if that were true, then — we think — it is even more alarming that Macy and Sato's main result is not replicable. It would imply that one can't 'repair' the situation and eventually save Macy and Sato's main result by introducing them. But what other candidates for a rescue operation are there? We do not know — as we don't know where we may have misunderstood assumptions or missed decisive points.

5.3

Regardless of the big effort it causes, replicating a model is nevertheless a worthwhile task. It strengthens the scientific character of the model. Replication is — in Robert Axelrod's words (1997) — "one of the hallmarks of cumulative sciences" and "is needed to confirm whether the claimed results of a given simulation are reliable." "Without this confirmation," he further claims, "it is possible that some published results are simply mistaken due to programming errors, misrepresentation of what was actually simulated, or errors in analyzing or reporting the results."

5.4

In our case, a computational model was used as evidence for an imminent danger, the possible break down of trust caused by too much social mobility. The computational evidence turned out not to be replicable. Obviously that does not imply that such an effect does not exist.

Notes

[1] As noted by one of our reviewers, Macy and Sato's use of the term social mobility is rather unusual. It generally refers to intergenerational mobility up or down the class hierarchy or income scale and not to mobility among different spatial partitions of the population. We nevertheless adopt Macy and Sato's expression.

[2] For Hume's contemporary and friend Adam Smith the empirical findings that motivated the Macy-Sato-model would probably not have come as a big surprise. Already in 1763 he stated important differences between the Dutch, the English and the Scots in his lecture The Influence of Commerce on Manners (Smith 1964). According to Smith's observations, the Dutch are the most faithful. “The English are more so than the Scotch, but much inferior to the Dutch.” The important point here is the explanation he gives. He explicitly rejects to impute the observed effects to different national characters. Instead he stresses that the Dutch are the “most commercial”, involved in lots of market exchanges — and under such circumstances learn that cheating does not pay off.

[3] It might be better to speak of “partition” because the “neighbourhoods” do not have a location themselves, i.e. we cannot say that a neighbourhood one is next to another neighbourhood. We keep Macy and Sato's jargon but this should be kept in mind.

[4] The effect of undiscounted payoffs of the first 1000 steps on the last 1000 steps would be avoided by a discount factor of 0.999

References

AXELROD R M (1997) Advancing the Art of Simulation in the Social Sciences. In Conte R, Hegselmann R, Terna P, editors. *Simulating Social Phenomena* (Lecture Notes in Economics and Mathematical Systems 456). Berlin: Springer-Verlag.

AXTELL R L, Axelrod R M, Epstein J M and Cohen M D (1996) Aligning Simulation Models: A Case Study and Results. *Computational and Mathematical Organization Theory*, 1 (2), pp. 123-141

CASSON M (1991) *The economics of business culture : game theory, transaction costs, and economic performance*. Oxford: Clarendon Press.

DORE R P (1987) *Taking Japan seriously : a Confucian perspective on leading economic issues*. Stanford, CA: Stanford University Press.

HAGEN JM and Choe S (1998) Institutional sanctions matter trust in japanese interfirm relations. *Academy of Management Review*, 23(3), pp. 589-600. 1998.

HUME D (1975a) *A Treatise of Human Nature*. Oxford: Clarendon Press.

HUME D (1975b) *Enquiry concerning the Principles of Morals*. Oxford: Clarendon Press.

MACY M and Sato Y (2002) Trust, cooperation and market formation in the US and Japan. *Proceedings of the National Academy of Sciences*, 99, pp. 7214-7220.

OUCHI W G (1981) *Theory Z : how American business can meet the Japanese challenge*. Reading, MA: Addison-Wesley Publishing Company, Inc..

PLATO (1961) "Protagoras". In Hamilton E and Huntington C (Eds.) *The collected Dialogues of Plato*. Princeton: Princeton University Press

SINGER P (1981) *The Expanding Circle: Ethics and Sociobiology*. New York: Farrar, Straus and Giroux.

SMITH A (1964) *Lectures on Justice, Police, Revenue, and Arms*. New York: August M. Kelley Publishers

YAMAGISHI T (1988) The provision of a sanctioning system in the united states and japan. *Social Psychology Quarterly*, 51, pp. 265-271.

YAMAGISHI T and Yamagishi M . (1994) Trust and commitment in the United States and Japan. *Motivation and Emotion*, 18, pp. 129-166.

Return to Contents of this issue