Trust and Social Cohesion
*An Agent-Based Model**

Lutz F. Krebs

International Conflict Research
ETH Zurich
8092 Zurich, Switzerland
krebs@icr.gess.ethz.ch

August 20, 2007

Work in progress. Please do not cite without permission.
Comments welcome!

Abstract.

This study introduces an agent-based model of the interactions between members of an ethnically heterogeneous society, using inter-ethnic trust and trust based on transaction history as the explaining factors. Preliminary tests of the model dynamics show that single shocks to social and inter-ethnic trust do not destabilize society sufficiently to explain the outbreak of violence. They do however provide a tool for group leaders to engage themselves as catalysts of distrust.

Keywords:
Agent-based modeling, heterogeneous societies; inter-ethnic trust; regime type transition; scapegoating; social dynamics.

*) Prepared for presentation at the 4th ECPR General Conference, Pisa, Italy, September 6-8, 2007. The author wishes to thank Luc Girardin and Nils Weidmann for their continued support in the development of the simulation, as well as Lars-Erik Cederman for conceptual feedback. This paper is based on research done in the project "Democratizing Divided Societies in Bad Neighborhoods" by the Swiss National Center of Competence in Research "Challenges to Democracy in the 21st Century" (NCCR Democracy), and funded by the Swiss National Science Foundation (SNSF). Financial support by the SNSF is gratefully acknowledged.

Introduction

Since the end of the Second World War, civil wars have consistently made up the largest share of conflicts around the world. With the end of the Cold War, they have captured the attention of the public by ending hopes for a more peaceful era. The most ferocious conflicts, such as the Balkan wars and the genocide in Rwanda, share ethnicity as a central theme, be it as a tool to pursue economic goals[1] or as the proper dimension for the articulation of grievances.[2] At the same time, numerous ethnically heterogeneous countries thrive on their diversity.

What can explain the decent of a society into ethnically structured conflict? How can ethnicity move from a non-issue to the predominant cleavage of a society? The present study is part of a two-pronged effort to understand the development of ethnic conflict in situations of regime type change. It formulates and tests a modeling framework for the analysis of inter-ethnic trust, a form of social cohesion that can inoculate societies against the wide-spread outbreak of ethnic violence.

The first chapter provides a brief overview of the relevant literature and details a possible path towards ethnic violence for societies experiencing regime type change. The second chapter discusses the method of choice, agent-based modeling (ABM), and the initial development of a model of inter-ethnic trust as a computer simulation. The third chapter discusses dynamics exhibited by the model during its first tests and the last chapter concludes with an outlook on future development.

1 Collier & Hoeffler 2004.
2 Gurr 1994; Wimmer 1997.

1 Background

In their study on the linkages between regime types and conflict, Mansfield and Snyder have identified that regime type transitions, and particularly the shift towards a more democratic regime, leaves countries vulnerable.[3] Based on this observation, a potential pathway towards ethnic violence can be derived.

A number of factors contribute to the vulnerability of young democracies to violence: Firstly, the mechanisms of control produced by authoritarian regimes are replaced with the democratic mechanism of communication.[4] Secondly, the political arena is opened to the largest possible set of political actors, leading to the foundation of many more parties and interest groups than will remain after a number of election cycles have been completed.[5] Whereas there used to be one voice with authority, there are now a great many voices, often with little authority.

This coincides with a crucial period of decisions, during which the members of a society need to determine the shape of their future regime and the level of access to power that individual groups will enjoy.[6] The outcome of these decisions will have a certain degree of finality – if one's rights have not been enshrined in the constitution, if one's group is not politically represented after the first national elections, it will be much harder to be heard.[7]

Even without considering the possibility of individual actors turning into spoilers, it should be clear that political actors have to fear long-term losses if they are unsuccessful during the early stages of a shift towards democracy. At the same time, competition for public attention is much harder than in more mature democracies. Following the logic of political survival, political actors will seize upon any topic or topics that will convey them to or keep them in power.[8]

3 Mansfield & Snyder 1995a/b, 2001, 2002a/b and 2005.
4 Vorrath et. al. 2007.
5 ibidem.
6 Bunce 2003.
7 Vorrath et. al. 2007.
8 de Mesquita et. al. 2004.

While ethnicity will not always be an important subject in political debates of ethnically heterogeneous societies, it will certainly be during a time when a new social structure is forged and the rights of individuals and groups are debated. However, ethnicity is a malleable concept. It is generally taken to be a group identity based on cultural, religious and linguistic factors. However, it is a socially constructed concept, guided and constrained by social attributes such as language and social pressures such as wealth differences.[9] As a result, it can be instrumentalized by political actors—indeed, it may already have been instrumentalized if the overthrow of the preceding regime had nationalist overtones. Like nationalism and other forms of identity, ethnicity contrasts an in-group, "us", with an out-group, "them". At the same time, it obscures differences and promotes collaboration between members of the in-group.

In an environment of heightened competition for power in a newly shaping society, political leaders may find ethnic identity a convenient concept to attract the attention of society, or if their standing within their own group is threatened, to secure their role as group leader. If multiple actors choose to play the ethnic card, an ethnic outbidding dynamic[10] may arise similar to the exchanges of Slobodan Milošević and Franjo Tuđman before the outbreaks of ethnic war in the Balkans in the early 1990s.[11]

While this process has been studied in detail for a number of prominent cases, there is no general model of the deterioration of a thriving multi-ethnic society to the point of inter-ethnic violence. In particular, the question remains how leaders are able to reframe identity to the point where preexisting trust between ethnicities dissolves. A functioning society, especially if it is heterogeneous, needs to evolve beyond solidarity that is merely based on similarity between actors[12] to a level of

9 Özkırımlı 2005.
10 Horowitz 1985.
11 Gagnon 1994.
12 "Mechanical solidarity" in Durkheim's terms.

solidarity that is constantly reaffirmed by mutually beneficial interaction.[13] [14] These regular interactions create a bond of trust both on the personal level and in terms of expectations towards the functioning of society. This trust facilitates transactions, increases social cohesion and in doing so, reaffirms itself.[15]

The present study is part of a two-pronged effort to understand the mechanisms that political leaders can use to breach social and inter-ethnic trust. The first part consists of case studies of situations of ethnic tension that provide the opportunity for political leaders to use ethnicity to their advantage. By tracing the process through individual events, the first part aims to identify how leaders can consciously or unconsciously create distrust. The second part, which is the subject of this study, aims to create a model of inter-ethnic trust that can then be used to simulate the effects of trust manipulations on the scale of a society.

The following chapter will describe the methodological choice and outline an initial version of the model.

2 Methodology and model building

In order to better understand the dynamics of inter-ethnic trust, a generative approach seems more suitable than an inductive approach. Fortunately, conflicts occur rather infrequently and the vast array of factors that played a role in the history of individual cases makes it hard to distill generally valid patterns that could then be applied to potential future cases.

This chapter discusses the implementation of such a generative approach. The next section presents the reasoning for the choice of method of this study. It is followed by a discussion of the assumptions that will underlie the model and an overview of an initial implementation as a computer simulation.

13 "Organic solidarity".
14 Soen 2003.
15 Cook 2001.

2.1 Strengths and weaknesses of Agent-Based Modeling

Given the purpose of understanding the development of inter-ethnic trust at the macro-level, a simplified recreation of these social dynamics seems a suitable approach. Agent-based modeling was chosen for this purpose since it provides social scientists with a set of tools for stylized experimentation. It is a computational approach like other forms of simulation, but unlike macro simulations, it requires few or no hard-to-guess systemic variables. Instead, it helps to understand how decentralized local interactions between heterogeneous agents create perceivable regularities at the macro-level. Moreover, it is flexible enough to deal with problems such as bounded rationality caused by limitations on the information available to actors and on the feasibility of the computational effort. As such, it requires weaker assumptions than its other close neighbor, rational choice theory.[16] It allows a more realistic depiction of complex systems based on interactions that cannot be predicted individually in the timing of their occurrence, their reasoning or their composition.[17]

However, agent-based modeling is a relatively young method and it has several disadvantages. Firstly, since it is a stylized, virtual experiment, it inherits problems with generalizability. At the same time, an agent-based model is based on designed algorithms instead of actual behavior, so the validity of its inner logic is not naturally given, either. This weakness is exacerbated by the fact that there is no easy way to search the universe of possible algorithms for social interaction.

While these caveats rule out the use of ABM for generalization and prediction and the method remains disputed due to its youth, it can be useful in the formation of theories by showing the consequence of different forms of micro-level interaction at the level of society,[18] and it has successfully been used in a number of contexts.[19]

16 Epstein 1999.
17 Jennings 1999.
18 Gilbert & Terna 1999.
19 Agent-based modeling has e.g. been used in the study of alliance formation (Cederman 1997), development and spread of epidemics (Epstein & Axtell 1996), military tactics (Ilachinski 1997), price distributions (Bak et. al. 1993), spatial distribution of unemployment (Topa 1997), spatial settlement patterns (Schelling 1971; Dean et. al. 1999), voting behavior (Kollman et. al. 1992), war sizes (Cederman 2003) and wealth distributions (Epstein & Axtell 1996).

The following section discusses the basic assumptions on which the model will be founded.

2.2 Founding assumptions

Agent-based modeling allows the creation of a virtual, highly-simplified society. In this case, the model simulation has to provide a society with multiple ethnic groups. Given a predetermined starting point, the members of this society need to interact, generating trust in each other.

Since artificial societies are entirely based on assumptions, they are at risk of becoming useless if any assumption is unrealistic or if the society that is envisaged is too complex. Too minimize this risk, the model was kept purposefully simple. The four broad assumptions for the present model are as follows:

1. *All actors in society belong to an ethnic group, and their ethnicity is visible to those around them.* This may not be true for a number of reasons. For one, individual actors may reject their ethnicity, e.g. because do not feel "at home" in their group or because they find that that their ethnicity should not be a defining characteristic for them. Alternatively, ethnicity may not be distinguishable through visible markers, allowing individuals to more easily swap sides. However, in a situation where ethnicity has become a relevant marker for a society, individuals may be forced to publicly align with one side regardless of their preferences. Therefore, this seems a reasonable assumption for the purposes of our model.
2. *Actors need to interact regularly with others, and geographic proximity is an important factor in the selection of trading partners.* While this statement will only ever be true to a certain degree, especially when individuals are tied to each other through friend- or kinship, it seems acceptable that a large share of any

individual's daily transactions is clustered close to his or her home and place of work.

3. *Actors will trade more with partners they trust.* Given rational behavior by actors and the absence of monopoly conditions, this would seem to be an acceptable assumption.

4. *Actors form trust based on similarity and experience.* This assumption also has limited applicability, since trust is a complex concept affected by a number of conscious and unconscious factors. Moreover, personal experience can be invalidated by new, possibly exogenous information. For the purposes of this model, we assume that all other factors are constant and deconstruct trust into three clear-cut factors.

    a. *Actors trust their own group.* This assumption is subject to the same caveats as assumption 1 about the association with a group. Moreover, if forced to associate with a group against their will, actors may actually trust their own group less. At the same time, during ethnic tensions in society, only limited information may be available to members of each group, and this information will likely cast their own group in a positive light. For the sake of simplicity, this assumption is accepted for the moment.

    b. *Actors trust others more if they have successfully interacted before.* This assumption forms a virtuous (or vicious) circle in combination with assumption 3 and is similarly subject to the rationality of behavior.

    c. *Trust can be influenced by exogenous information about the behavior of others.* Outside information, if it is sufficiently out of the ordinary, may lead individuals to adjust their current perception of other actors (both for better and for worse). Such information could e.g. take the form of public accusations towards a particular group, benevolent actions of individuals belonging to a certain group, etc. While this is a rough assumption, it

seems to be confirmed by daily observation e.g. of the public reaction towards minorities, the generation of rumors about outsiders and similar occurrences.

Based on these four assumptions, a computational model was written. The next section discusses the implementation in detail.

2.3 Model implementation

The model was realized using the Geographic Research on War Laboratory (GROWLab), a toolkit specifically designed for agent-based models in the social sciences[20]. It formalizes the setup of individual actors, provides a variety of different topologies in which the actors can interact and allows the efficient collection and visualization of data from individual simulations. The discussion of the model is divided into three parts: first the shape of the interaction space, then the behavioral rules for agents and finally the facilities for data collection.[21]

2.3.1 Interaction space

The "world" in which the simulated society is placed is a two-dimensional, square grid of dimensions that can freely be chosen by the user.[22] The actors in the simulation are "agents", autonomous computer programs with a limited perception of society and a predefined set of behaviors. Corresponding with assumption 1, each agent is associated with one ethnic group represented by a geometric shape and then placed in a free grid cell. The model provides parameters for the number of ethnic groups, their relative size and the presence of segregated settlement.[23]

[Illustration 1 about here.]

20 See Weidmann & Girardin (2006) for a detailed description of its features. GROWLab can be downloaded at http://www.icr.ethz.ch/research/growlab.

21 The model including its source code is available for inspection. To install the model, download to plug-in package from http://www.icr.ethz.ch/people/krebs/ethnictrust-0.2.0.zip and copy it to the GROWLab models folder. To view the source code, extract the contents of the plug-in package using freely available software such as 7Zip (Windows) or StuffIt Expander (Mac).

22 Simulations described below use a grid of 15×15 cells.

23 The influence of different values for these parameters is explored in the next chapter on model dynamics.

Each agent can only perceive his immediate neighborhood, i.e. the agents in the cells bordering his cell.[24] This yields eight neighbors for agents in the middle of the grid, five neighbors for those situated at the edge of the grid and only three neighbors for agents placed in a corner. Illustration 1 shows an agent belonging to the group of circles (black), and his field of vision (gray). Agents only interact with their neighbors and are unaware of their neighbors' interactions with others or trends at the level of society. They will attempt to "trade" with their neighbors as a basic requirement to social coexistence, and the decision whether they want to trade with any given neighbor will be based on their trust towards that neighbor's ethnic group and (optionally) their previous experiences with this individual agent.[25]

In order to test the influence of a shock to the trust on which social interaction is based, the model allows for the transmission of such a shock to all agents in society. The shock takes the form of a deduction (or if desired, an addition) to the level of trust each individual agent experiences.[26] Moreover, the model allows the choice whether the trust towards one particular group or all groups is affected. Finally, the timing of the shock can be chosen.

2.3.2 Behavior

As described above, each agent will attempt to trade with his neighbors based on his trust towards them. Therefore, assumptions 3 and 4 represent the two basic behaviors available to agents: they can try to interact and they can reassess the trust they feel towards individuals and groups in society.

During each time period, all agents will interact up to ten times with their neighbors. For each possible interaction, the initiating agent "A_I" will consider the receiving agent "A_R"'s ethnicity and (optionally) his trading record with A_R. The likelihood of A_I attempting to trade with A_R is proportional to A_I's trust towards A_R's

24 This represents the second assumption.
25 An implementation of assumptions 3 and 4.
26 This represents assumption 4c.

ethnic group (if only inter-ethnic trust is considered) or the average of inter-ethnic trust and the trust towards A_R as an individual. If A_I wants to trade with A_R, then A_R gets to make the same choice based on his trust towards A_I's ethnic group and A_I as an individual. Only if both find each other trustworthy will the transaction succeed. Based on this two-step process, any considered action will fall in one of three categories:

- avoided transaction (if A_I does not trust A_R),
- rejected transactions (if A_R does not trust A_I), or
- successful transactions (if both trust each other).

At the end of each time period after the trading, all agents reassess the trust they experience. At the start of each simulation run, all agents are seeded with an initial level of trust.[27] This starting value will be discarded after ten interactions, and replaced with a value based on the cumulative experiences with members of other groups and individual agents.[28] Trust is calculated individually for all groups and transaction partners, and is defined as

$$trust = \frac{successes}{attempts}$$

where attempted transactions equal the sum of all successful and rejected transactions.

It should be noted that agents only consider the transactions they initiated, so that each transaction is counted only once. Trust towards society as a whole is then defined as the sum of the trust towards each group weighed by that group's relative share of society.

27 These values are drawn randomly, but their upper and lower bounds can be set as parameters.

28 Following assumption 4a, trust towards the agents' own group is complete, while trust towards others is based on experience as argued by assumption 4b. Complete trust towards one's own group does not imply that agents cannot be distrustful of individual members of their own group if they have made bad experiences with them.

At this point, the model can inject a deduction (or addition) to the level of trust all agents experience either towards one particular group or towards all groups in society. This represents the shock to social trust that could e.g. be caused by political rhetoric directed against a scapegoat group or large-scale social events like violent riots. The trust of all agents towards the target will remain lowered by the shock value until the end of the simulation.

The newly calculated values including the shock value will then be used as the basis for trading in the next round.

2.3.3 Data recording

After the conclusion of each time step, the model gathers the trust values of all agents as well as their internal count of avoided, rejected and successful transactions and reproduces them in a variety of ways.

[Illustration 2 about here.]

Illustration 2 shows the user interface of the simulation toolkit. Of note are the spatial view in the upper-right corner, which shows the position of all agents (similar to Illustration 1) and colors them in accordance with their individual level of trust towards society[29]. This allows the visualization of inherently spatial effects, e.g. low trust at borders between segregated groups. The data view in the lower half of the screen visualizes the development of trust and transaction records over time as graphs and provides a table with the numerical data. Individual agents do not have access to any statistics at the level of society in accordance with their limited field of vision.

Given the general setup of the model described above, the next chapter discusses the model's behavior for different parameter sets.

29 High values are displayed in shades of green while red colors indicate low values.

3 Model dynamics

This section covers the dynamics of the model. First, the model's robustness to different values of basic parameters is tested, before the influence of manipulatable social factors such as as the conception of trust and segregated settlement of groups are discussed.

3.1 Basic choices and robustness

A range of parameters are responsible for the basic shape and behavior of the model: the random seed, the range of initial trust values, the number of ethnic groups as well as the timing, target and size of the shock.

Firstly, the random seed is the starting value for the random number generator. The model itself is deterministic; as long as the seed value and other parameters are unchanged, repeated runs of the simulation will always have the same outcome. Illustration 3 shows runs of a basic model with a small shock at time 0. The development of average trust in society is plotted for ten replications of the same model with different random seeds. The repetitions produce paths differing in intercept and slope, but retaining the same shape. This indicates that the model is both relatively robust towards different seeds and provides sufficient variation in its outcome. In all remaining graphs, the values will be averaged out over a number of test runs with different seeds. While this erases minor variations unique to each run, broader trends will be visible more clearly.

[Illustration 3 about here.]

The graph already shows a basic feature of the model: it is equilibrium-seeking and average trust in society asymptotically approaches its equilibrium value both at the beginning of the model, where only a small adjustment is needed, and after the exogenous shock at time 0. The shock itself plays out over a long time, with a cumulative effect much bigger than the initial value of the shock.

Next, the model is tested for its robustness towards different boundary values of initial trust. Agents are created with a starting value of trust randomly picked between these bounds, and perform their first ten interactions with other groups based on this value since they do not yet have experiences to rely on. Illustration 4 shows the development of average trust in society for low, medium and high starting values.[30] It should be noted that trust in society is generally higher than the initial value of trust since all agents have complete trust in their own group.

[Illustration 4 about here.]

The paths are virtually parallel, indicating that the choice of starting range will not have an important effect on the development of the model. For all following models, the initial values of trust are randomly generated from a range between 50-80%. This relatively high range of values was chosen to represent a society that is generally healthy and to give shocks of distrust enough room to develop their complete effect.

Next, the number of ethnic groups in the simulated society is addressed. For now, the model allows the inclusion of 2-3 ethnic groups, with one group playing the role of the scapegoat. Illustration 5 shows that the number of groups does not have a large effect. The development of average trust is shifted upwards for simulations with only two groups since all agents trust their own group completely and in a society that is split two ways instead of three, there simply are more group members to trust.

[Illustration 5 about here.]

Both Illustration 5 and 6 also show that the size of the post-shock adjustment is proportional to the size of the shock. Both graphs show runs with deductions of trust of 5% and 10%, with the latter having a much larger toll on society. Moreover, Illustration 6 shows the difference in effects of shocks targeted at only one instead of

30 Each line represents five replications of the model with different seeds.

all groups in society. For any given shock size, a shock targeted towards all groups yields a long-term loss much larger than a group-based shock.

[Illustration 6 about here.]

The shock size is proportionate to the difference between the size of the scapegoat group and the overall size of society. This implies that the consequences of a general shock are worse in comparison to a group-based shock even for the group that would suffer the brunt of a scapegoating campaign, as long as the scapegoat group does not represent the largest share of society.

Finally, Illustration 7 shows the influence of the timing of the shock, with each graph showing the adjustment path for simulations with "histories" of differing length before the incidence of the shock. Not surprisingly, the longer a society is allowed to interact peacefully, the less important an eventual shock becomes.

[Illustration 7 about here.]

This concludes the discussion of basic parameters. So far, the model does not exhibit macro-level patterns that seem unrealistic. The next section discusses initial tests on two factors that could facilitate the outbreak of violence in vulnerable societies: the conception of trust (or rather, whether individual trust is able to play a dampening role in ethnically charged situations) and the presence of segregated settlement.

3.2 Manipulatable factors

The initial tests described in the preceding section have indicated that the model does not behave unrealistically, and the results indicate a set of parameters that seems suitable for more detailed tests: three groups of similar size with initial trust values of 50-80% are subjected to a shock of -5% relatively early in the simulation.[31]

31 Complete specifications for replication are available upon request.

We use this basic set up to investigate the role of two factors that can be manipulated to some extent and that could lead a society towards ethnically structured violence. Firstly, the question whether impressions of an ethnic group as a whole outweigh individual trust. While political leaders cannot directly determine or influence how members of society value outside information over their personal experiences, they can use propaganda to give outside information more weight.

The model allows agents to either base their decisions entirely on their perception of ethnic groups (allowing for the full effect of a smear campaign) or to value their view of an ethnic group equally with their trust towards individuals. Due to the fact that agents exhibit complete trust towards their own group, enabling personal trust dramatically decreases overall trust in society (cf. Illustration 8). Moreover, individual bad experiences weigh more heavily when they are divided by the smaller number of individual transactions, further decreasing overall trust in society. Beyond this, however, individual trust limits the effects of scapegoating, as was to be expected. Leaders interested in manipulating inter-ethnic trust therefore have to exert additional effort to overcome positive personal experiences in the minds of society.

[Illustration 8 about here.]

Segregated settlement can help them in this. A segregated population distribution prevents the collection of positive experiences with members of other groups and at the same time gives bad experiences a bigger weight. Illustration 9 shows the spatial development of trust follows the predicted pattern. With high values of trust represented by shades of green, the highest values can be found in areas where actors interact mostly or exclusively with their own kind.

[Illustration 9 about here.]

Illustration 10 then shows that this effect can be detected at the macro-level. In situations of individual trust, the effect of segregation is positive, shifting overall

trust upwards. The map underlines that this is a dangerous dynamic. As long as people trust their own group more than others, a situation that can be created through a scapegoating campaign, an "ethnically clean" area will appear more attractive.

[Illustration 10 about here.]

If the agents in this model had a choice to move, a voluntary migration similar to the Schelling model[32] could set in. The combination of a low level of segregation and a moderate shock to trust towards one group have created a starting point for a further deterioration of inter-ethnic trust.

4 Outlook

This study has presented an initial look at an agent-based model of inter-ethnic trust. Based on four assumptions regarding the presence of ethnicity, the need for interaction and the reciprocal link between interaction and trust, the model shows sensible behavior for individual ranges of parameters. More importantly, the model shows how the existence of small pockets of ethnically homogeneous settlements combined with a moderate amount of distrust towards one group can create a situation where mixed settlement areas become unappealing. This may lead to further segregation and may strengthen tensions between different ethnic group.

However, the model is still in an early stage. The preliminary evaluation presented here will serve as the starting point for three series of more complex tests. Firstly, the algorithm representing trust will be exchanged with other possible algorithms to test the model's robustness to its own specification. Secondly, some of the more stylized assumptions will be dropped. E.g. the rule that agents trust their own group completely will be modified to a simple trust bonus, which seems more realistic. Finally, additional shocks can be implemented to recreate situations in which leaders of multiple groups conduct scapegoating campaigns.

32 Schelling 1971.

Once the model is fully operational, it will be used to systematically identify combinations of factors that lead to a break-down in inter-ethnic trust. These factors will then be tested in a series of case studies of ethnic tension in the Middle East to see whether break-down conditions were present in those cases that saw an eruption of ethnic violence, and to test whether political leaders were consciously or unconsciously engaged in the creation of such break-down conditions.

Bibliography

- Axelrod, R. (1997a), *The Complexity of Cooperation*, Princeton University Press.
- Axelrod, R. (1997b), *Advancing the Art of Simulation in the Social Sciences*, Springer, 21—40.
- Bak, P.; Chen, K.; Scheinkman, J. & Woodford, M. (1993), 'Aggregate Fluctuations from Independent Sectoral Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics', *Recherche Economiche* **47**, 3—30.
- Bunce, V. (2003), 'Rethinking Recent Democratization: Lessons from the Postcommunist Experience', *World Politics* **55**(2), 167—192.
- Casti, J. L. (1997), *Would-Be Worlds: How Simulation is Changing the Frontiers of Science*, John Wiley & Sons.
- Cederman, L. (2003), 'Modeling the Size of Wars: From Billiard Balls to Sandpiles', *American Political Science Review* **97**(1), 135—150.
- Cederman, L. (1997), *Emergent Actors in World Politics*, Princeton University Press.
- Cook, S. A. (2006), 'The Promise of Pacts', *Journal of Democracy* **17**(1), 63—74.
- Dean, J.; Gumerman, G.; Epstein, J.; Axtell, R.; Swedlund, A.; Parker, M. & McCarroll, S. (1999),Understanding Anasazi Culture Change Through Agent-Based Modeling, *in* GJ Gumerman & T Kohler, ed.,'Modeling Small Scale Societies', Oxford University Press, New York.
- Epstein, J. M. (1999), 'Agent-Based Computational Models and Generative Social Science', *Complexity* **4**(5), 41—60.
- Epstein, J. M. & Axtell, R. (1996), *Growing Artificial Societies: Social Science from the Bottom Up*, MIT Press.

- Gagnon Jr., V. P. (1994), 'Ethnic Nationalism and International Conflict: The Case of Serbia', *International Security* **19**(3), 130—166.
- Gilbert, N. & Terna, P. (2000), 'How to build and use agent-based models in social science', *Mind & Society* **1**(1), 57—72.
- Holland, J. H. (1995), *Hidden Order: How Adaption Builds Complexity*, Perseus Books.
- Horowitz, D. (1985), *Ethnic Groups in Conflict*, University of California Press, Berkeley.
- Ilachinski, A. (1997), 'Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare', Memorandum, Center for Naval Analyses Research.
- Jennings, N. R. (1999), Agent-Based Computing: Promise and Perils, *in* 'Proceedings of 16th Int. Joint Conf. on Artificial Intelligence'.
- Kollman, K.; Miller, J. & Page, S. (1992), 'Adaptive Parties in Spatial Elections', *American Political Science Review* **86**, 929—937.
- Mansfield, E. D. & Snyder, J. (2005), *Electing to Fight: Why Emerging Democracies go to War*, MIT Press, Cambridge, MA.
- Mansfield, E. D. & Snyder, J. (2002a), 'Incomplete Democratization and the Outbreak of Military Disputes', *International Studies Quarterly* **46**(4), 529—549.
- Mansfield, E. D. & Snyder, J. (2002b), 'Democratic Transitions, Institutional Strength, and War', *International Organization* **56**(2), 297—337.
- Mansfield, E. D. & Snyder, J. (2001),Democratic Transitions and War, *in* Fen Osler Hampson Chester A. Crocker & Pamela Aall, ed.,'Turbulent Peace', United States Institute of Peace Press, Washington D.C., 113—126.

- Mansfield, E. D. & Snyder, J. (1995a), 'Democratization and the Danger of War', *International Security* **20**(1), 5—38.
- Mansfield, E. D. & Snyder, J. (1995b), 'Democratization and War', *Foreign Affairs* **74**(3), 79—97.
- de Mesquita, B. B.; Smith, A.; Siverson, R. M. & Morrow, J. D. (2004), *The Logic of Political Survival*, MIT Press.
- Özkırımlı, U. (2005), *Rethinking Nationalism: A Social Constructivist Approach*, Palgrave Macmillan.
- Schelling, T. (1971), 'Dynamic Models of Segregation', *Journal of Mathematical Sociology* **1**, 143—186.
- Soen, D. (2003), 'A Binational Society: The Jewish-Arab Cleavage and Tolerance Education in the State of Israel', *Israel Affairs* **9**(1-2), 97—120.
- Topa, G. (1997), 'Social Interactions, Local Spillovers and Unemployment', Manuscript, Department of Economics, New York University.
- Vorrath, J.; Krebs, L. & Senn, D. (2007),'Linking Ethnic Conflict & Democratization: An Assessment of Four Troubled Regions', Working paper No. 6, NCCR Challenges to Democracy in the 21st Century, Zurich.
- Weidmann, N. B. & Girardin, L. (2006),'GROWLab: A Toolbox for Social Simulation', First World Congress on Social Simulation, August 21-25, 2006, Kyoto, Japan.
- Tesfatsion, L. & Judd, K., ed. (2006), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Vol. 2, North Holland Publishing, Amsterdam.

Illustrations

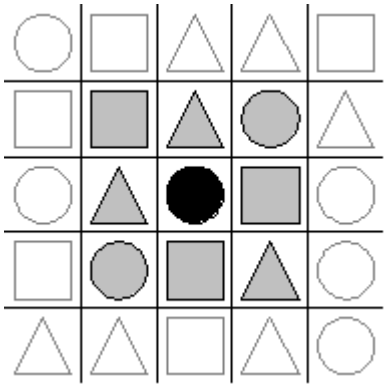

Illustration 1: An agent and his neighborhood

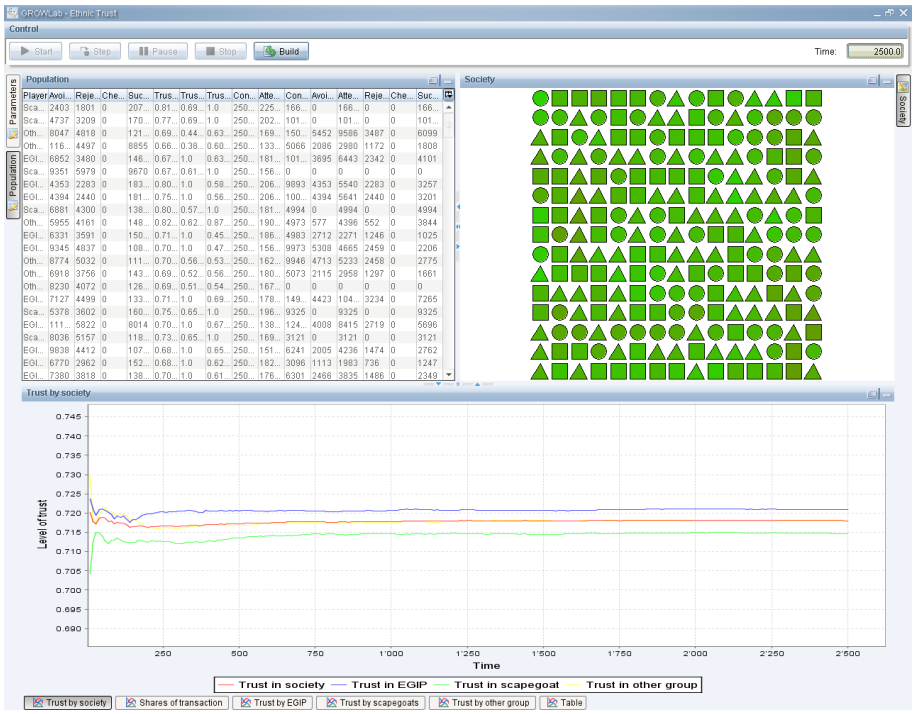

Illustration 2: User interface with sample run (three groups, no shock)

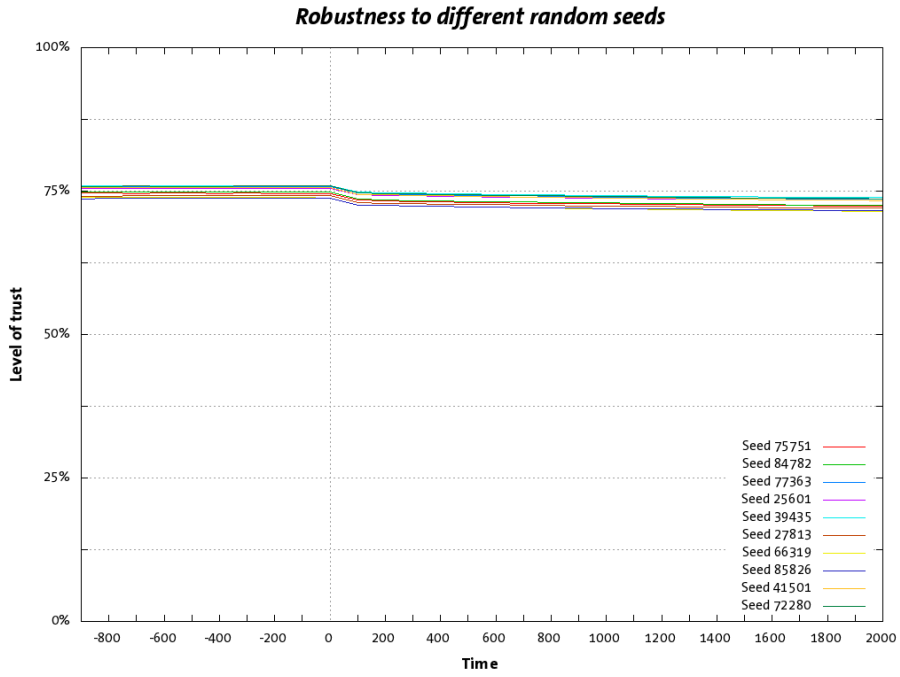

Illustration 3: Robustness to different random seeds

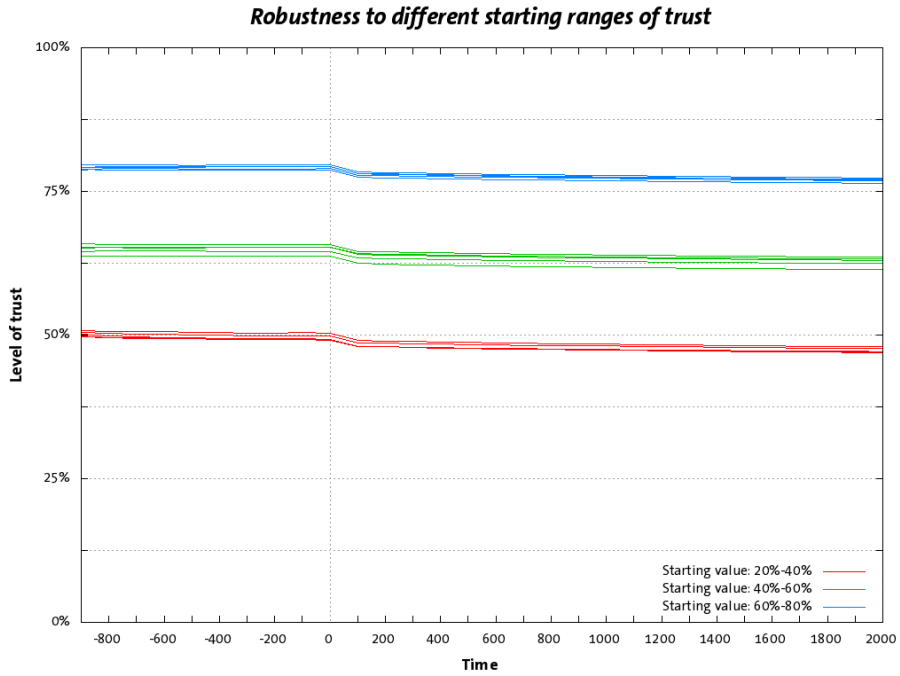

Illustration 4: Robustness to different starting ranges of trust

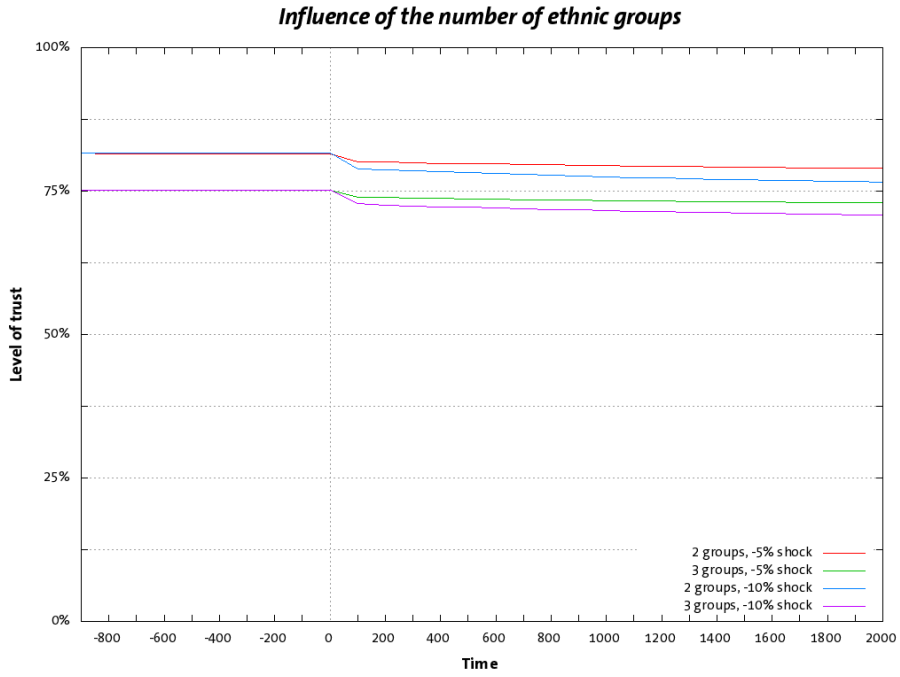

Illustration 5: Influence of the number of ethnic groups

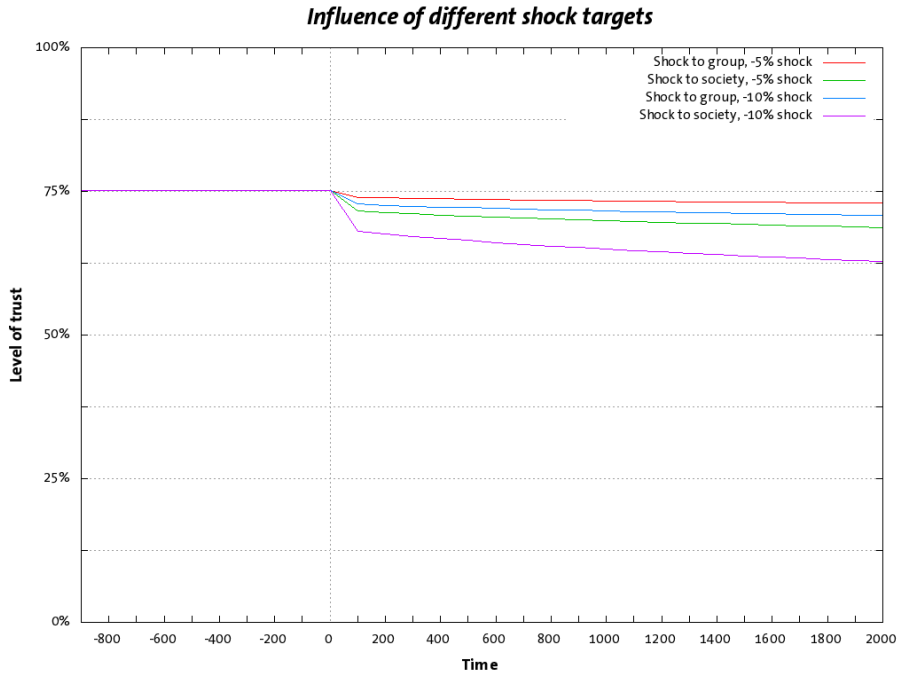

Illustration 6: Influence of different shock targets

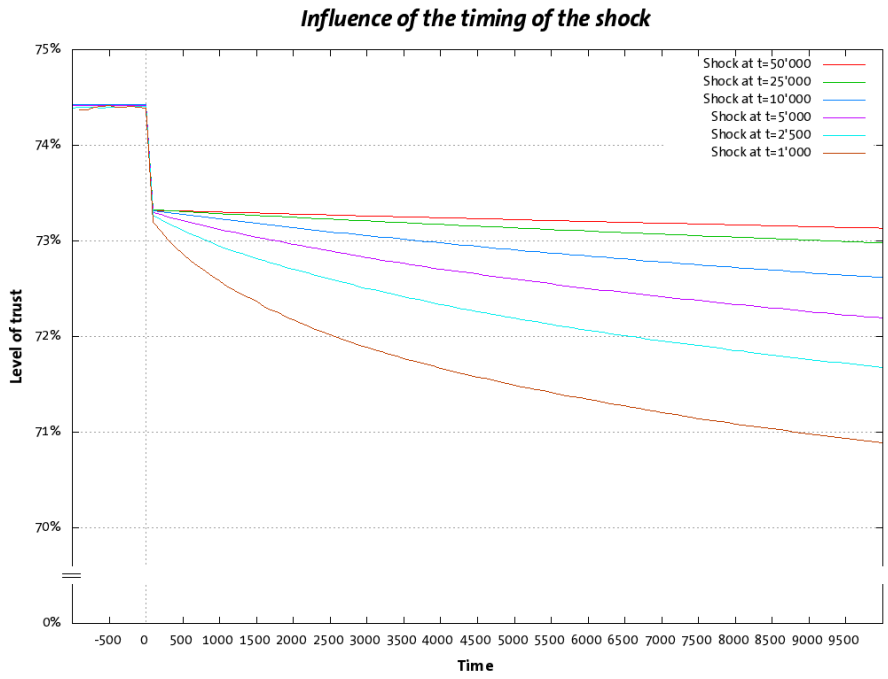

Illustration 7: Influence of the timing of the shock

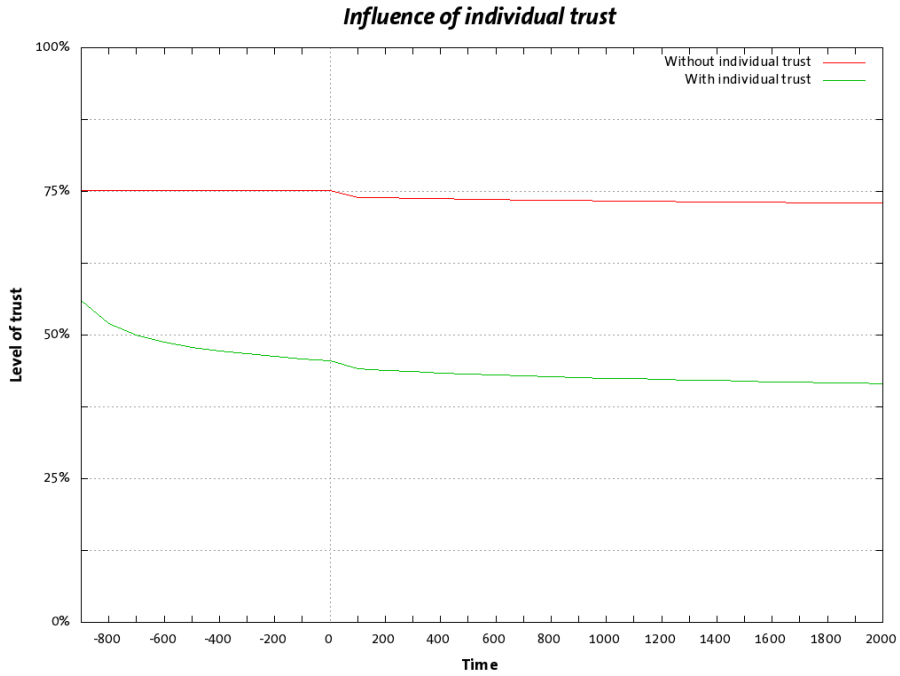

Illustration 8: Influence of individual trust

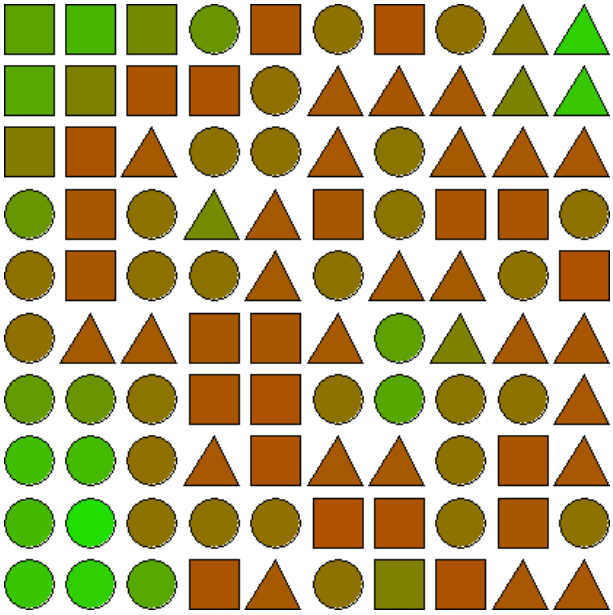

Illustration 9: Influence of segregation and individual trust

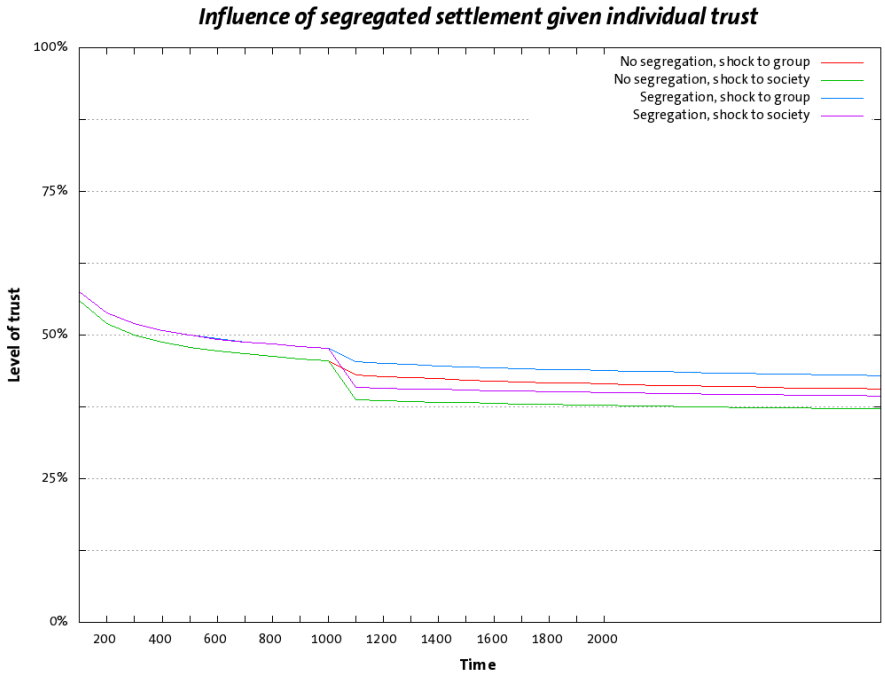

Illustration 10: Influence of segregated settlement given individual trust