



Delft University of Technology

Random Deviations Improve Micro–Macro Predictions An Empirical Test

Mäs, Michael; Helbing, Dirk

DOI

[10.1177/0049124117729708](https://doi.org/10.1177/0049124117729708)

Publication date

2020

Document Version

Final published version

Published in

Sociological Methods and Research

Citation (APA)

Mäs, M., & Helbing, D. (2020). Random Deviations Improve Micro–Macro Predictions: An Empirical Test. *Sociological Methods and Research*, 49(2), 387–417. <https://doi.org/10.1177/0049124117729708>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Random Deviations Improve Micro–Macro Predictions: An Empirical Test

Sociological Methods & Research
2020, Vol. 49(2) 387–417
© The Author(s) 2017



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0049124117729708
journals.sagepub.com/home/smr



Michael Mäs¹ and Dirk Helbing^{2,3}

Abstract

Many sociological theories make critically different macropredictions when their microassumptions are implemented stochastically rather than deterministically. Deviations from individuals' behavioral patterns described by microtheories can spark cascades that change macrooutcomes, even when deviations are rare and random. With two experiments, we empirically tested whether macrophenomena can be critically shaped by random deviations. Ninety-six percent of participants' decisions were in line with a deterministic theory of bounded rationality. Despite this impressive micro-level accuracy, the deterministic model failed to predict the observed macrooutcomes. However, a stochastic version of the same microtheory largely improved macropredictions. The stochastic model also correctly predicted the conditions under which deviations mattered. Results also supported the hypothesis that nonrandom deviations can result in fundamentally different macrooutcomes than random deviations. In conclusion, we echo the warning that deterministic microtheories can be misleading. Our findings show that taking into account deviations in sociological theories can improve explanations and predictions.

¹ Department of Sociology/ICS, University of Groningen, Groningen, the Netherlands

² ETH Zurich, Zürich, Switzerland

³ Delft University of Technology, Delft, the Netherlands

Corresponding Author:

Michael Mäs, Department of Sociology/ICS, University of Groningen, Grote Rozenstraat 31, Groningen 9712 TG, the Netherlands.

Email: m.maes@rug.nl

Keywords

randomness, noise, micro–macro problem, experiment, evolutionary game, theory, formal modeling, coordination game, network

Contemporary sociology has an ambivalent relationship with the concept of randomness. On the one hand, randomness is an integral part of statistical models, where commonly the biggest part of a dependent variable's variance remains unexplained and is therefore assumed to be random. Sociological theory, on the other hand, has been criticized for often neglecting that at least parts of human behavior deviates from the general patterns that theories describe (Macy and Tsvetkova 2015). This deterministic approach to sociological theory is problematic because a theory “in which individual behavior is completely determined, flawlessly executed, entirely knowable, and perfectly predictable is not only empirically implausible, it can also be highly misleading” (Macy and Tsvetkova 2015:324). Deviations from general behavioral patterns of individuals are of potentially critical importance whenever individuals do not act in isolation but react to each others' action. In such settings, deviations can lead to behavioral reactions in actors' social environment, which can, in turn, motivate further behavior changes. Formal models demonstrated that such chains of reaction can have profound impact on the collective level, even though they might have been sparked by a random and rare event on the individual level.

Unlike in sociology, random deviations are basic building blocks of theoretical models in other disciplines that study systems with a micro–macro structure, such as physics (Nicolis et al. 1977), chemistry (Pomeau 1986), biology (Camazine et al. 2001), traffic research (Treiber, Hennecke, and Helbing 2000), cognitive science (Rabinovich et al. 2008), and economics (Binmore and Samuelson 1999; Selten 1975; Harsanyi and Selten 1988). A very prominent example are theories of biological evolution, where random mutations in the genes of individuals generate the biological variation that species need to adapt to changes in their environment. Similarly, the random microscopic motion of particles (temperature) is a key ingredient of models of turbulent fluid flows and the kinetic theory of gases (Nicolis et al. 1977).

A growing body of formal modeling work demonstrates that deviations can also change predictions of theories explaining sociological phenomena such as social norms (Young 2015), conventions (Young 1993), opinion polarization (Mäs, Flache, and Helbing 2010; Pineda, Toral, and Hernandez-Garcia 2009), cultural assimilation (Huckfeldt, Johnson, and

Sprague 2004; Klemm et al. 2003), residential segregation (van de Rijt, Siegel, and Macy 2009), social classes (Axtell, Epstein, and Young 2000), collective behavior (Granovetter 1978), and the production of public goods (Kollock 1993). One intriguing finding was, for instance, that Schelling's famous model of residential segregation predicts higher levels of ethnic segregation when random deviations from Schelling's deterministic microassumptions are included (van de Rijt et al. 2009). According to Schelling's original model, even highly tolerant populations segregate because every move of an actor implies that his ethnic group becomes less represented in her old neighborhood and more represented in the new neighborhood. As a consequence, former neighbors of the same ethnic group and new neighbors of the other ethnic group may also decide to move, which can, in turn, motivate further moves. Thus, a single actor's decision to move can precipitate a cascade of further movings that lead the population to surprisingly high degrees of ethnic segregation. It turned out that random moves can intensify segregation because they can spark moving cascades that would not have occurred in Schelling's deterministic model.

So far, deviation effects have mainly been investigated from a theoretical perspective. Empirical support for the macroeffects of microdeviations was provided in an laboratory experiment by Goeree, Holt and Palfrey (2007) who found that information cascades are much shorter than a deterministic rational choice model predicts. Information cascades can arise when individuals with limited information make decisions in a sequence. At some point in the sequence, decision makers rationally ignore their private information and only respond to the decisions of those actors who decided earlier. This can lead to a situation where all individuals rationally make the wrong choice (Bikhchandani, Hirshleifer, and Welch 1992). Contrary to this deterministic prediction, Goeree et al. observed that information cascades broke down because participants sometimes responded to their private signal rather than rationally following the herd.

We report here results from two empirical studies testing hypotheses about the effects of microlevel deviations on macrooutcomes. Study 1 had two aims. First, we tested the central notion that deviations matter, studying a social setting where a microtheory that abstracts from deviations makes very different macropredictions than the same microtheory with deviations. Second, we tested whether the theoretical model with deviations accurately identifies the conditions under which deviations matter for macrooutcomes and when macrooutcomes are unaffected by deviations. Study 2 challenged the assumption that deviations are random. On the one hand, the prediction that microdeviations have macroeffects is most surprising when deviations

are assumed to occur randomly. On the other hand, we show that sometimes the theoretical assumption that deviations are random leads to very different macropredictions than assumptions about systematic deviations from the prevalent patterns of individual decision-making, even when deviations are rare. Study 2, therefore, empirically tested whether it can be problematic to treat deviations as random phenomena when they actually follow a pattern.

Empirically testing hypotheses about the effects of microlevel deviation on macrooutcomes is challenging for various reasons. First, theories predict that macroeffects of deviations are not immediate. In contrast, deviations trigger off behavioral cascades, processes that may or may not take very long until they translate into macroeffects. This lack of precision makes many predictions immune to empirical falsification. Another problem is that deviations are difficult to quantify, in particular when they might be random and rare. What is more, it is very difficult to determine whether a given action constitutes a deviation or is in line with the general behavior pattern. For instance, a white person leaving a white neighborhood and moving into a black one does deviate from Schelling's model assumptions. However, a less abstract micromodel that also takes into account actors' budget constraints might perfectly capture this moving decision. Thus, it often depends on the assumed micromodel whether a given behavior is considered a deviation or not.

Confronted with these methodological challenges, we decided to study microdeviations and their macroeffects in two controlled laboratory experiments. This had two crucial advantages. First, strategically providing participants with limited information about the behavior of others and their social environment, we managed to create a setting where our participants made decisions based on the so-called best-response heuristic, a simple rule of boundedly rational decision-making (e.g., Alós-Ferrer and Netzer 2010; Blume 1993; Fudenberg and Levine 1998; McFadden 1973; McKelvey and Palfrey 1995; Montanari and Saberi 2010; Young 1993). We were also able to unequivocally determine for each observed decision of our participants whether it was a best response or a deviation. It turned out that 96 percent of all observed decisions were best responses. Second, we were able to implement a setting where a deterministic version of the best-response model makes different macropredictions than a version of the same microtheory that adds random deviations. This allowed us to empirically test whether the best-response model with deviations makes more accurate macropredictions than a deterministic version.

Despite the artificial context of our laboratory studies, we contribute to a debate of broad sociological relevance because any theory of *social action* in the sense of Max Weber (1978) may make different macropredictions, when

deviations are included. Whenever individuals do not act in isolation but respond to their social environment, it cannot be excluded that deviations spark behavioral cascades that change macropatterns. This is independent of how complex or simple the micromodel is and does not only apply to the best-response heuristic. In fact, the literature provides multiple examples of prominent deterministic micromodels that were believed to explain given macrooutcomes but actually failed to provide valid explanations when random deviations were included (e.g., Kollock 1993; Klemm et al. 2003; Mäs et al. 2010). Likewise, there are examples where including assumptions about random deviations from individuals' behavioral patterns (e.g., Mäs et al. 2010; Pineda et al. 2009) provided new solutions to long-standing theoretical puzzles (e.g., Abelson 1964). In the light of these strong theoretical reasons to take into account deviations, it is important to empirically test the notion that deviations matter.

Random behavior is often falsely conceived as a residual component that needs to be stripped away in order to reveal systematic behavior that allows one to explain and predict the behavior of humans and collectives. In contrast to this view, we argue that taking into account deviations in sociological theories will improve explanations and predictions. Even though microlevel deviations might be random and unpredictable, their macroeffects can be systematic and, thus, predictable. Obviously, when individual deviations are random, it is not possible to predict when they occur. However, there is a rich theoretical toolbox that allows one to identify the structural conditions under which deviations have no macroconsequences and when they have the potential to spark behavioral cascades that change macrooutcomes (Foster and Young 1990; Freidlin and Wentzell 2012). Identifying these conditions promises to improve explanations and predictions of collective phenomena.

The remainder of this article is structured as follows. In the following section, we illustrate the effects of deviations in a very simple, stylized example and show how predictions about the conditions under which deviations matter for macropredictions can be derived. The two subsequent sections summarize the hypotheses, the design, and the results of the two empirical studies. We conclude with a summary of the results and an agenda for future research.

A Stylized Example

The notion that microlevel deviations can have profound and systematic effects on the collective level is counterintuitive, in particular when deviations are random and rare. However, the following stylized example




Figure 1. Stylized example.

illustrates that deviations can have macroeffects even in a setting that is highly simplistic and easy to analyze. To be sure, this example has not been chosen with the motivation to study the impact of deviations on a specific sociological phenomenon, even though very similar settings have been explored to study the emergence and diffusion of norms and conventions (Bicchieri 2005; Cialdini, Reno, and Kallgren 1990; Ehrlich and Levin 2005; Montanari et al. 2010; Opp 2001; Young 1993, 2015). We chose it because it illustrates the main hypotheses that we tested and because it is possible to implement it in a laboratory experiment, which in turn makes it possible to directly test these hypotheses.

Consider the social network shown in Figure 1, a circle network where each of the 20 nodes is connected to the two closest neighbors to the right and to the left. Assume furthermore that all nodes simultaneously choose between two options, labeled “red” and “blue” and that the position on the circle determines the type of the network node. There are two types, which we denote here metaphorically as “cats” and “dogs.” Cats receive $P = 1$ money units (MU) when choosing blue. Dogs receive the same payoff for selecting red. Furthermore, participants receive 1 MU for each network neighbor who chose the same color. These payoff rules create a game resembling the battle of the sexes

in game theory, a decision problem where actors seek to coordinate with their interaction partners but have opposing preferences (if $P > 0$).

Assume furthermore that nodes make decisions based on a simple heuristic. When they decide between red and blue for the first time and, thus, cannot condition their choice on the past choices of their network neighbors, they simply choose the option that corresponds to their type (cats: blue; dogs: red). After that, nodes choose the so-called myopic best response, which is a simple heuristic of boundedly rational decision-making (e.g., Alós-Ferrer and Netzer 2010; Blume 1993; Fudenberg et al. 1998; McFadden 1973; McKelvey and Palfrey 1995; Montanari et al. 2010; Young 1993). It assumes that participants choose the color that promises the highest payoff if the four neighbors choose the same color as in the previous round. In other words, actors are not hyperrational maximizers that take into account all possible future consequences of their behavior. Actors are also not perfectly informed about all potentially relevant aspects of the decision problem, such as the network structure and the types of their neighbors. Actors just assume that their neighbors will stick to their previous choices and then choose the color that promises a higher payoff in the present round.

These microassumptions imply that dynamics can reach three Nash equilibria, collective rest points where no individual can increase her payoff by unilaterally changing color. The first equilibrium is called “anomic coexistence” and obtains when all individuals choose the color that corresponds to their type (cats choose blue and dogs choose red, as in Figure 1). In this situation, all actors expect to receive a payoff of $P + 2 = 3$ MU when they stick to their choice. Switching to the other color would lead to a payoff of only 2 MU. Thus, the best response for all actors is to not change color. Second, when all individuals choose the same color (either all red or all blue), a *descriptive norm* has emerged. This is a rest point, because actors receive a payoff of either 5 MU or 4 MU, depending on whether the network coordinated on the color that the respective actor prefers or not. Choosing the opposite color would lead to a payoff of either 0 or 1, which is always smaller. Thus, in a world where all actors choose the best response, all actors would follow the descriptive norm. Third, when $P = 1$ also “segmented coexistence” is a rest point. Segmented coexistence means that the circle is split into multiple internally coordinated but mutually anticonordinated segments of at least five players.

Strikingly, the best-response model makes different prediction about which equilibrium is selected by the dynamics, depending on whether deviations are included or not. First, when all actors deterministically follow the heuristic described above, then every actor chooses in the very first round the color that corresponds to her type. In the second round, all actors realize that

they have two neighbors who chose the same color and, thus, expect to receive $P + 2 = 3$ MU when they stick to their choice and only 2 MU when they change color. As a consequence, best responders will not change color and the collective rest point of anomic coexistence is reached.

Contrary to the deterministic best-response model, however, the stochastic best-response model predicts the emergence of a descriptive norm when $P = 1$. In fact, two deviations from the best-response rule suffice to generate a rapid shift from anomic coexistence to coordination on a descriptive norm if $P = 1$, as Figure 2 illustrates. The figure shows that cats choose blue and dogs choose red at the outset. This aggregates to the collective state of segmented coexistence, which would be an equilibrium in a deterministic world. Let us add, however, that in period 1, actors 2 and 20 happen to deviate from the best-response rule, choosing red instead of blue. All subsequent decisions are best responses. Nevertheless, the system shifts from anomic coexistence to a descriptive norm. In period 2, actors 2 and 20 will stick to red, as this promises a payoff of 3 MU instead of only $P + 1 = 2$. Furthermore, actors 4 and 18 also choose option red, expecting to receive 3 MU. This is 1 MU more than for option blue. In period 3, actors 6 and 16 face the same situation as 4 and 18 in the previous period and will also switch to red. This process continues until all actors have chosen red. Thus, the two initial deviations have sparked a cascade of color changes that moves along the circle until the whole population coordinated. Two deviations are the minimum requirement for coordination to emerge. However, we show in the Online Appendix that similar dynamics can unfold when the two deviations occur at maximally distant positions in the network and at different points in time.

The collective state of a descriptive norm is more robust to deviations than anomic coexistence. Assume, for instance, that the network coordinated on color blue and again actors 20 and 2 happen to deviate from the best-response rule. In the next round, the two actors will switch back to red as they expect to receive $P + 3 = 4$ MU for red and only 1 MU for blue. Actor 1, however, will be affected by the two initial deviations. Two of her neighbors chose red in the previous round, making her expect $P + 2 = 3$ MU for red and only 2 MU for blue. However, in the following round, all of her neighbors will have chosen blue again, motivating also actor 1 to switch back to blue. As a consequence, the system returned to the collective state of a descriptive norm, despite the initial deviations.

However, deviations do not always spark cascades like those illustrated in Figure 2 and, therefore, do not always alter collective outcomes. For instance, deviations fail to affect collective outcomes when the payoff P that actors receive for choosing the color that corresponds to their type is




Figure 2. Example for a cascade sparked off by two deviations ($P = 1$).

increased to 3 MU. This is a rather subtle change in the payoff structure, in particular because anomic coexistence and descriptive norms remain rest points. In a state of anomic coexistence, all actors expect to receive $P + 2 = 5$ MU for sticking to the color of their type, which is more than the expected 2 MU for switching. Likewise, actors will not change color if a descriptive norm has emerged as switching would decrease their expected payoff to either 0 or P , which is always less than the expected payoff of following the norm (either $P + 4$ or 4).

Under $P = 3$, actors who deterministically follow the simple heuristic will choose the color of their type in the first round and will then stick to this choice forever, which aggregates to the collective pattern of anomic coexistence. This macroprediction does not change when deviations are included. To illustrate this, assume the same scenario as in Figure 2 where actors 20 and 2 happen to deviate. When $P = 3$, both actors will immediately switch back to the color of their type, as this promises a payoff of $P + 1 = 4$ MU rather than only 3 MU. What is more, their network neighbors will be unaffected by the two deviations, as choosing the color of their type remains their best response. More deviations do not change this prediction. For instance, if in addition to actors 20 and 2, 18 and 4 happen to deviate, then actors 20 and 2 have four red neighbors, which makes red the best response. However, actors 18 and 4 have only three red neighbors. Thus, their best response is to switch back to their blue. As a consequence, actors 20 and 2 also have only three red neighbors in the next round and will, thus, also switch back to blue. In sum, the collective will return to the pattern of anomic coexistence after only two rounds.

This simple example illustrates the two theoretical insights (Macy and Tsvetkova 2015) that our studies put to the test. First, sometimes very few deviations from the otherwise dominant behavioral patterns of individuals can have profound effects on macrooutcomes. In these cases, even a micro-theory that perfectly describes the otherwise prevalent behavioral patterns of individuals will make false macropredictions. Second, microdeviations do not always have macroeffects, but it is possible to derive testable hypotheses about the conditions under which they have the potential to spark off cascades that lead the system into another state.

Study I

Study 1 put two theoretical predictions to the test. First, we tested whether deviations on the microlevel can affect collective outcomes. Second, we tested whether a theoretical model that includes random deviations correctly

predicts under what conditions deviations matter and when they do not affect macro-outcomes.

To this end, we conducted a laboratory experiment that implemented the core aspects of the stylized example from the previous section at the Decision Science Laboratory at ETH Zurich (<https://www.descil.ethz.ch>). That is, we arranged a computer network with the structure shown in Figure 1, assigned participants to the network nodes, and confronted them with the same decision problem as the actors in the example. Furthermore, it was critical to create a setting for the participants that would lead them to decide based on the best-response heuristic, the decision rule that we assumed in the previous section. As the best-response heuristic is intuitive and simple, providing participants with all information needed to make a best response and withholding crucial information necessary to form decisions based on alternative decision rules was sufficient to create a setting where 96 percent of participants' decisions were best responses. In this setting, we studied the two experimental treatments that we explored in the previous section. Under $P = 1$, the stochastic version of the best-response model makes opposite macropredictions than its deterministic counterpart, which allowed us to test whether the model that takes into account deviations makes more accurate macropredictions. Under $P = 3$, both micromodels make the same macroprediction. This allowed us to test whether or not the stochastic model is able to correctly predict when deviations matter and when they do not change macrooutcomes.

Procedures

We invited groups of 20 participants to the laboratory, where they sat in separate cubicles and interacted in a computer network. Every participant of an experimental session was randomly assigned to one of the positions in the circle network shown in Figure 1 and was instructed that per round she would earn one MU for each interaction partner who chose the same color and P MUs for selecting the color of their type. All experimental sessions consisted of 150 interaction rounds, but participants were not informed about this.

During the experiment, we always updated participants about the color choices of their four network neighbors in the previous round, which is the only information needed to identify the best response. We withheld further information, to ensure that it was impossible to apply other typical decision rules such as the imitation of successful others (Judd, Kearns, and Vorobeychik 2010; Kirchkamp and Nagel 2007; McCubbins, Paturi, and Weller

2009; Selten and Apesteguia 2005) or forward-looking rules of expected payoff maximization (Frey, Corten, and Buskens 2012). In particular, participants were not informed about their neighbors' payoff rules and types, others' payoffs, the structure of the network, and the number of interaction periods (Hart et al. 2003).

It turned out that most of participants' decisions were best responses. As we studied only values of the payoff parameter P where participants cannot be indifferent between the two color options, we could determine for every observed decision whether it was a best response or a deviation. It turned out that 96 percent of all participants' decisions in studies 1 and 2 were in line with the best-response rule, which is in line with similar experiments (Lim and Neary 2016). In other words, for the setting of our experiments, only 4 percent of the choices deviated from the deterministic best-response rule. The simple heuristic that we studied in the previous section furthermore assumes that dogs would choose red and cats would choose blue in the first round, as this guaranteed them a payoff of P . It turned out that 78 percent of the participants did this.

Participants were recruited from a general pool of student volunteers from ETH Zurich and the University of Zurich that was set up for behavioral experiments where deception of participants is strictly forbidden. For each of the two experimental treatments, we conducted three independent replications with 20 participants per replication. In total, there were 120 participants in study 1. At the very beginning of the experiment, participants were informed about the payoff rules and that, at the end of the experiment, the computer would randomly pick three rounds of the experiment to determine each participant's payoff. Participants received two Swiss Francs for every MU earned in these three rounds. On average, participants earned 33 Swiss Francs. All experimental sessions were finished within one hour.

Hypotheses

Figure 3 visualizes the hypotheses that follow from the deterministic and the stochastic versions of the best-response model, showing ideal-typical predictions for each of the two experimental treatments. The figure shows only typical predictions of the stochastic model, but in section 1.2 of the Online Supplemental Material, we report results from large-scale simulation experiments, demonstrating the dynamics shown in Figure 3 are indeed typical outcomes of the stochastic model. The color of the markers in Figure 3 shows actors' color choices in the 150 periods. Squares represent best-response




Figure 3. Hypotheses derived from the deterministic and the stochastic version of the best-response model for the two treatments of study I.

decisions and circles identify deviations. Furthermore, circle size indicates the deviation cost, the difference in payoff for choosing the best response and the payoff the actor actually received when deviating. Bigger circles indicate higher deviation costs.

The top graph of each panel visualizes the hypothesis derived from the deterministic best-response model (no deviations). As described in the previous section, the deterministic model predicts that actors will choose the color that corresponds to their type in the first period and will then stick to this color for the remainder of the experiment. Thus, the deterministic model predicts the collective pattern of anomic coexistence.

To generate predictions for the best-response model with *random* deviations, we adopted the standard logit-response model (Alós-Ferrer and Netzer 2010; Blume 1993; Fudenberg et al. 1998; Montanari et al. 2010; McFadden 1973; McKelvey and Palfrey 1995; Young 1993). This model assumes that there is always a positive chance that agents deviate from the best-response rule and that deviations are more likely when the payoff difference between the two options is small. In other words, this model assumes that deviations occur with a higher probability when they imply small deviation costs. Formally, the probability p_i^{blue} that actor i chooses option blue in period t is given by:

$$p_i^{\text{blue}} = \frac{1}{1 + e^{-(\beta(U_i(\text{blue}) - U_i(\text{red})))}}, \quad (1)$$

where $U_i(\text{blue})$ represents the expected payoff of individual i when she chooses blue and her network neighbors choose the same color as in the previous period. $U_i(\text{red})$ is the expected payoff of choosing red. Parameter β models the degree to which individuals choose the best response or deviate. For instance, $\beta = 0$ models entirely random choices and $\beta = \infty$ implies deterministic best-response decisions. Based on the observed decisions during the two studies, we statistically estimated a β of 1.5 ($SE = 0.02$). This parameter value was also used to simulate the typical dynamics of the stochastic model shown in Figure 3. However, in the Online Supplemental Material, we report predictions of the stochastic model with less or more deviations from the best-response model. We also compare model predictions for the duration of the experiment (150 rounds) and much longer time frames as well as different initial color distribution.

Supporting the assumption that deviations occur more frequently when they imply smaller costs, Figure 4 shows the share of red and blue choices depending on the deviation costs. The gray line is the estimated logistic function ($\beta = 1.5$), showing that the logit-response model is supported by our data.¹ We used this estimate to generate the ideal-typical dynamics of the stochastic model shown in Figure 3.

Figure 3 (center graph of panel A) shows that under $P = 1$, deviations change macropredictions even when deviations are random. The network




Figure 4. Share of red and blue choices depending on deviation costs (data from both studies).

started in a state of anomic coexistence but within only 20 rounds all actors coordinated on one color, which contradicts the prediction of the deterministic model. Deviations keep occurring also after the descriptive norm has emerged, but the descriptive norm remains stable.

However, deviations do not always spark cascades. The center graph of panel B in Figure 3 shows that under $P = 3$, anomic coexistence remains stable for the 150 periods of the experiment, despite the deviations ($\beta = 1.5$). Thus, for $P = 3$, the deterministic and the stochastic best-response model make the same macroprediction.

Several aspects of our experimental design make the hypothesis that a descriptive norm will emerge under $P = 1$ intuitively puzzling. First, unlike in earlier research (Judd et al. 2010; Kearns, Suri, and Montfort 2010), participants were not rewarded for reaching certain collective outcomes and were never informed about the distribution of behavior in the population (Judd et al. 2010; Kearns et al. 2010; McCubbins et al. 2009). This excludes that any collective patterns obtained in the laboratory because participants intended to create it. Second, participants were randomly assigned one of the two types (cats or dogs), which were given monetary incentives for opposite behavior. This makes coordination on one color more surprising than in earlier experiments where all participants received the same monetary payoff for a given behavioral option (Berninghaus, Ehrhart, and Keser 2002; Cassar 2007; Frey et al. 2012). Third, the two descriptive norms that could emerge (all red, or all blue) implied the same collective welfare and were equally risky. Thus, contrary to the design of earlier studies (Berninghaus et al. 2002; Cassar 2007; Frey et al. 2012), our design made both norms equally focal (Schelling 1960), which excluded that a norm emerged because participants

independently from each other chose the most focal option. Fourth, each participant was involved in only one realization of the experiment in order to avoid that previous collective outcomes affected decisions and that participants learned to create or avoid a certain collective outcome. Finally, the networks that we studied were perfectly symmetric, excluding that some participants had a greater impact on the collective dynamics, for example, because of a more central position in the network (Frey et al. 2012; Judd et al. 2010; Kearns et al. 2010).

What is more, the hypothesis that a descriptive norm will emerge under $P = 1$ is risky, as it is built on two implicit but crucial assumptions. First, the model specified in Equation (1) assumes that actors do not condition their behavior on whether the choices of their network neighbors were deviations or not. It appears plausible, however, that a participant who observes a color change by a network neighbor reasons that this color change resulted, for example, from a mistake and might therefore decide to not change color herself even if this would have been the best response. Such behavior would challenge the prediction of the stochastic model, as it would prevent that deviations spark cascades that change macrooutcomes. Second, our predictions will be challenged when deviations are not random but follow specific patterns. The fact that each respondent had a financial interest in one of the two colors implies that respondents may strategically rather than randomly deviate from the best-response model. For instance, participants on cat positions who realize that their neighbors tend to choose red, the color that dogs prefer, might strategically deviate, choosing blue even though red would be the best response. These participants would give up short-term payoffs in order to harvest long-term payoffs arising from the collective pattern that is most profitable for themselves. Such systematic deviations would make the coordination on one color unlikely and, thus, challenge the hypothesis derived from the stochastic model.

Results

The empirical results of study 1 clearly confirmed the hypotheses derived from the stochastic model. For comparison with the hypotheses derived from the two theoretical models, the bottom graphs of the two panels in Figure 3 show a typical experimental session from the respective treatment. The Online Appendix contains the same graphs for the remaining sessions.

In all three replications with $P = 1$, anomic coexistence was found only at the beginning of the dynamics. Contrary to the predictions of the deterministic best-response model, all experimental populations coordinated quickly




Figure 5. Test of the hypotheses derived from the deterministic and the stochastic model.

on one of the two colors. These descriptive norms were stable until the experiment ended, as predicted by the stochastic model. The experiment also confirmed that anomic coexistence is much more stable under $P = 3$. In all three replications under this condition, anomic coexistence was found throughout the 150 periods of the experiment, showing that deviations do not always change macropatterns and that the stochastic model correctly predicted when the collective pattern is unstable or not.

Figure 5 reports the statistical tests of the macropredictions. To quantify which macropattern emerged, we measured the size of the biggest cluster in the network. A cluster was defined as a set of nodes with adjacent positions on the circle that chose the same color at a given point in time. This measure adopts a value of 1 when the system is in a state of anomic coexistence and adopts its maximal value of 20 when a descriptive norm has emerged.

The green bars in Figure 5 show the average size of the biggest cluster in the network as predicted by the deterministic best-response model. We have shown in Figure 3 that the deterministic best-response model predicts the pattern of anomic coexistence when all participants initially choose the color of their type. However, when some individuals happen to choose the opposite color in the very first round, then the deterministic best-response model may make slightly different macropredictions. For instance, it is possible that the system happens to start in a collective state of segmented coexistence, which is also an equilibrium (see section 2). To exclude that our statistical test is affected by this, we simulated the deterministic best-response dynamics that follow from the initial color choices observed during the respective sessions of the experiment. The size of the green bars, thus, shows the average size of the biggest cluster in the final 100 rounds of the simulated dynamics. The

error bars were obtained with linear multilevel models to control for the nestedness of observations in the three experimental sessions. The red bars show average size of the biggest cluster in the network at the final (150th) round in 1,000 independent simulations with the stochastic model, assuming $\beta = 1.5$. In section 1.2 of the Online Supplemental Material, we show that predictions of the probabilistic model are robust to changes in β . Also these simulations started with the initial color distribution observed during the experiment.

Gray bars show the average outcome in the final 100 periods of the experiment. Estimates and 95 percent confidence intervals were obtained with linear multilevel models, controlling for the nestedness of observations in experimental sessions.²

Figure 5 shows that the experiment confirmed the macrohypothesis derived from the stochastic best-response model for the treatment with $P = 1$ and, thus, did not support the deterministic model. This supports the theoretical notion that deviations on the level of individuals can have a decisive impact on collective outcomes. A micromodel that takes into account deviations makes, thus, more accurate macropredictions than a deterministic version of the same model, even when deviations are assumed to be random. What is more, deviations do not always change macro-outcomes, but the stochastic model correctly predicted when deviations matter.

Study 2

Study 1 provided empirical support for the notion that the macropredictions of a microtheory can become more accurate when deviations are taken into account. On the one hand, this finding challenges intuition because we assumed that deviations occur randomly. On the other hand, the assumption that deviations are random might also be problematic, as Figure 6 suggests.

Figure 6 shows examples of seemingly systematic deviation patterns that we observed in the two studies. In example a, participant 15 attempted to enforce her preferred color (see Online Appendix Figure S19). Even though blue would have been the best response, this participant kept choosing red, the color of her type. In example b, participant 6 seemed to have attempted to disrupt the descriptive norm that had emerged, systematically deviating several times (see Online Appendix Figure S21). In example c, participant 9 attempted to conserve her preferred color (see Online Appendix Figure S19). Furthermore, example d suggests that participants who observed that a network neighbor changed color might have deviated with an increased probability, generating small deviation chains (see Online Appendix Figure S14).




Figure 6. Observed deviation patterns.

Study 2 was designed to test the hypothesis that systematic, nonrandom deviations can lead to very different macropatterns than random deviations, even when deviations are rare. In study 1, deviation patterns occurred but did not affect the dynamics to a degree that macrooutcomes differed from the predictions of the micromodel with random deviations, which suggests that adding stochasticity suffices to improve macropredictions even when deviations are not entirely random but follow patterns that are not captured in the microtheory. However, there might be conditions under which systematic deviations generate dynamics that differ from those generated by random deviations. We therefore tested whether it may be misleading to assume that deviations are random if they are actually systematic.

Study 2 differed from study 1 only in one aspect, the structure of the interaction network. In study 1, participants were linked to the two closest neighbors to their left and to their right, which created a mixed network where nodes have two neighbors of the same type and two neighbors of the opposite type. For study 2, however, we linked actors to their two closest neighbors on the circle and the




Figure 7. Network used in study 2.

two neighbors with a distance of two, as Figure 7 shows. The resulting network is very similar to the network of study 1, but actors are connected only to actors of the opposite type. Like in study 1, we studied two treatments with different values of parameter P . We conducted three independent replications with $P = 1$ and four with $P = 3$. In total, there were 140 participants.

Predictions

Similar to the settings of study 1, coordination on a descriptive norm and segmented coexistence are also rest points in study 2, but there is a third very important stable pattern, which we call “dynamic oscillation.” This is a dynamic pattern where cats and dogs always choose opposite colors but also swap color in every period. This pattern is stable in that switching is always the best response to the choices of one’s interaction partners in the previous round.

The fact that best-response behavior generates dynamic oscillation is important, as it allowed us to study collective outcomes when decision makers deviate systematically from the best-response rule. When a participant applies the best-response rule, she assumes that the interaction partners will stick to their previous choices, an assumption that is obviously false when the

system is the dynamic-oscillation phase and color choices alternate every round. As the participants of our experiment were always informed about the past choices of their network neighbors, we expected that they would *systematically* deviate from the best-response model whenever the system was in a state of dynamic oscillation.

Figure 8 visualizes the predictions of the deterministic and the stochastic ($\beta = 1.5$) versions of the best-response model for the two treatments of study 2. For the treatment with $P = 1$, both theoretical models predict that the populations will end up in a state of dynamic oscillation. When dogs choose red and cats choose blue in the first round, then nobody coordinates with any of her neighbors, leading to a payoff of only P for everybody. Switching color in Round 2, however, promises a payoff of four MU, which motivates all individuals to change color. In the subsequent round, however, decision makers realize again that they countercoordinated with their neighbors and that their payoff is even smaller (zero MU). As another color change promises a payoff of $P + 4$, everybody changes back to the initial color. Thus, the deterministic best-response model predicts that individuals will switch color until infinity. With random deviations, as assumed by the stochastic best-response model (equation [1]), the state of collective oscillation is stable for the duration of the experiment when $P = 1$, as the center graph of panel A shows. That is, even when actors sometimes happen to deviate from the best-response rule, the collective pattern of dynamic oscillation remains stable.

However, even though both the deterministic and the stochastic best-response model predict the collective pattern of dynamic oscillation when $P = 1$, we expected this collective pattern to be very fragile and to disappear quickly. As just described, it is not plausible that participants apply the best-response rule when the system is in a state of dynamic oscillation. Therefore, we expected participants to systematically deviate whenever the system is in a state of dynamic oscillation and that dynamics would lead the systems into one of the other equilibrium candidates, either coordination on a descriptive norm or on a segmented coexistence. Segmented coexistence is stochastically very unstable, however. In fact, a single deviation by any member of a segment who has a network connection to another segment member who is of the type that does not prefer the color of the segment and who is connected to a node outside the segment will motivate these contacts to also change color. Coordination on one of the two colors is stochastically much more stable. If, for instance, the network coordinated on red and a dog happens to deviate, then even the four connected cats (who prefer blue) will stick to red, as red promises a payoff three and blue only a payoff of $P + 1 = 2$. As a consequence, the system will return to the state of coordination.




Figure 8. Predictions and typical results for the two experimental treatments of study 2.

In sum, for the treatment with $P = 1$, the deterministic and the stochastic version of the myopic best-response model predict the pattern of dynamic oscillation. However, if our intuition is correct that sometimes systematic deviations can generate different macro-outcomes than random deviations,

then we would observe coordination on a descriptive norm in the treatment with $P = 1$.

For the treatment with $P = 3$, the deterministic and the stochastic best-response model make opposite predictions, as Figure 8 illustrates. The deterministic model predicts the pattern of dynamic oscillation, like under $P = 1$. According to the stochastic model, however, random deviations lead to recurrent phase transitions from dynamic oscillations to segmented coexistence to coordination and back (see center graph of panel B in Figure 8). If an actor deviates in a period with an even number, then all of her contacts will stop alternating, as this promises a payoff of $1 + P = 4$, rather than only 3 MU. If the actor who just deviated sticks to his last color choice, which would also be a deviation, then the system has reached a new equilibrium (segmented coexistence). However, segmented coexistence is also stochastically very unstable, as a deviation by any segment member will motivate other segment members of also change color. Segments can also grow until the whole network coordinated on one color, however, as a result of further deviations.

Strikingly, also coordination is stochastically unstable when $P = 3$. For instance, if everybody coordinated on option red and a dog happens to deviate, then her four neighbors will also change color. If the dog who just deviated sticks to his past choice, which would be a deviation, then a new equilibrium is reached (segmented coexistence). If, however, the dog who deviated chooses the best response after the deviation (red), then the segment of five actors will start the oscillation pattern.

The recurrent phase transitions that the best-response model with random deviations generates are intriguing. However, if our intuition is correct that participants will deviate systematically from the best-response rule whenever the system is in a state of dynamic oscillation, then the following prediction can be formulated. The system will start in dynamic oscillation, but systematic deviations will drive it into another equilibrium (either coordination or segmented coexistence). But also these equilibria are unstable and the system will keep switching from coordination to segmented coexistence. Phases of dynamic oscillation will be very short, as they lead to systematic deviations.

Results

The two bottom graphs in the two panels of Figure 8 show typical dynamics from the experiment, showing that our experimental results supported our prediction that nonrandom deviations can generate different collective patterns than random deviations. Challenging the predictions of both the




Figure 9. Test of the hypotheses derived from the deterministic and the stochastic model.

deterministic and the stochastic best-response model, participants coordinated quickly on one of the two colors in all three replications for $P = 1$.

Likewise, we observed only very short phases of oscillation in the four replications for $P = 3$. In line with the predictions of the stochastic best-response model, participants coordinated on one color. However, this happened much faster than the stochastic model predicted, which provides further support for our assumption that dynamic oscillation leads to systematic deviations. Nevertheless, in agreement with the stochastic best-response model, coordination was unstable when $P = 3$. In one of the four replications (the one shown in Figure 8), the descriptive norm broke down as a result of a deviation and the system entered a state of segmented coexistence.

Figure 9 reports the same statistical tests as those from study 1, but the result is very different. Unlike in study 1, both best-response models fail to explain the macropatterns observed in the two treatments of the experiment. In addition, the analyses reported in section 1.2 of the Online Supplemental Material show that also a stochastic model with more or less deviations would not have explained the empirically observed dynamics. In a nutshell, the results of study 2 challenged the standard assumption that deviations are random. Systematic deviation has the potential to lead systems into different collective states than random deviations.

Discussion

The notion that the behavior of individuals can lead to complex and unexpected macropatterns when individuals do not act in isolation but interact

with each other is, in a nutshell, sociology's contribution to the understanding of human behavior. Formal models of social processes and theories from other disciplines suggest that in these settings, even rare and random deviations can have decisive impact on how microbehavior translates into macro-patterns. Under conditions that theoretical methods allow to identify (Foster and Young 1990; Freidlin and Wentzell 2012), deviations may spark cascades that have a substantial impact on collective behavior. To test this notion, we conducted two laboratory experiments.

Our empirical studies support the notion that microdeviations matter for macrooutcomes. We draw three conclusions from our empirical results. First, our findings support the notion that even micromodels that very well capture the prevalent patterns of individual behavior can make false macropredictions. Deviations from the patterns that microtheories describe can spark behavioral cascades that profoundly change collective outcomes. Second, the stochastic model correctly predicted the conditions under which deviations have the potential to alter macrooutcomes, even when deviations are assumed to be random and, thus, unpredictable. Third, when individuals deviate in a systematic way from otherwise prevalent patterns of behavior, the resulting macrooutcomes can differ substantially from those generated by random deviations. In other words, even in settings where deviations are rare, it can be misleading to assume that deviations are random if they actually follow a pattern. Including random deviations may thus fail to improve macropredictions if the modeler failed to take into account an important deviation pattern. Stochasticity improves macropredictions of accurate microtheories, but it does not necessarily fix the predictions of false micromodels.

Future research on the stochastic component of individual behavior is crucial. Two main challenges lie ahead. First, more theoretical and empirical work on the conditions of stochastic instability is needed in order to understand better why and when the stochastic component of individual behavior alters collective dynamics. Crucial methodological tools to analyze the stochastic stability of social systems (Foster and Young 1990; Freidlin and Wentzell 2012) have been developed, and their usefulness has been demonstrated on social processes such as the evolution of conventions and norms (Young 2015, 1993), the spread of innovation (Montanari et al. 2010; Young 2011), the emergence of hierarchies (Axtell et al. 2000), the polarization of opinions (Mäs et al. 2010; Pineda et al. 2009), and residential segregation (van de Rijt et al. 2009). Empirical studies are needed to test whether stochastic models outperform deterministic theories also in these contexts. Our studies have demonstrated that laboratory experiments are a useful approach

to testing hypotheses about deviation effects. The core advantage of laboratory studies is that they allow the researcher to measure whether a given action was a deviation or not. An important challenge for future empirical research will be to study deviation effects in the field (an inspirational example is the study by Chadeaux 2016).

Second, far too little is known about the conditions of deviations (Butler, Isoni, and Loomes 2012; Freidlin and Wentzell 2012; Goeree et al. 2005; McKelvey and Palfrey 1995; Loomes 2005; Lim and Neary 2016; Wilcox 2008; Naidu, Hwang, and Bowles 2010; Hwang et al. 2014; Pradelski and Young 2012). Deviations have been argued to have multiple sources including mistakes, misperceptions, inertia, or trial-and-error experiments. These forms of deviations have been incorporated in very different ways into formal models, which can matter for macropredictions. One important assumption, for instance, is that deviations occur with a constant rate (Kandori, Mailath, and Rob 1993; Young 1993), which might be a good model of deviations resulting from mistakes. Deviations resulting from trial-and-error experiments, however, might be better captured by the logit-response model (Alós-Ferrer and Netzer 2010; Blume 1993; Ellison 1993; Young 1998, 2011), as it seems plausible that actors experiment more likely when it involves low costs. Our studies support the logit-response model (see Figure 4), but more research is needed to identify the conditions under which alternative deviation models offer more or less accurate descriptions of real deviations (Mäs and Nax 2016). It is important to point out, however, that from the view of a sociologist interested in macrophenomena, it is certainly not necessary to develop complex deviation theories seeking to predict individual deviations. Study 1 illustrated that often even very simple deviation models, such as the logit-response model, suffice to improve macropredictions. What is needed, however, are theories that point to the conditions under which deviations are more or less likely.

Our experimental design was very much tailored to the best-response heuristic, which is just one of the many candidates for a micromodel. Nevertheless, there is no reason to expect that deviations do not have similar macroeffects when actors apply even simpler or more complex decisions rules. For instance, theoretical models predict deviation effects when actors also use simple reinforcement-learning heuristics (Pradelski and Young 2012) or imitate successful others (Kandori et al. 1993). In fact, cascades can be sparked whenever decision makers are influenced by the behavior of others. One can be certain that deviations are irrelevant for collective outcomes, only when actors act in perfect isolation.

Acknowledgments

Mäs and Helbing acknowledge support by the European Commission through the ERC Advanced Investigator Grant “Momentum” (Grant 324247). The authors would like to thank Stefan Wehrli and Silvana Jud for helping organize the experiments. For helpful comments, we thank Ryan O. Murphy, Thomas Grund, Vincenz Frey, Andreas Flache, Thomas Chadeaux, Olivia Woolley, Steve Genoud, Chris Snijders, Siegwart Lindenberg, Peter Hedström, Michael Hechter, and, in particular, Heinrich Nax.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Supplemental Material

Supplemental material for this article is available online.

Notes

1. A detailed analysis of the microlevel deviations during our two experiments has been published by Mäs et al. (2016). These analyses provided support for the logit-response model.
2. The simulated dynamics of the deterministic and the stochastic model that started with the initial color choices observed during the experiment are also reported in the Online Appendix.

References

- Abelson, Robert P. 1964. “Mathematical Models of the Distribution of Attitudes under Controversy.” Pp. 142-60 in *Contributions to Mathematical Psychology*, edited by N. Frederiksen and H. Gulliksen. New York: Rinehart Winston.
- Alós-Ferrer, Carlos and Nick Netzer. 2010. “The Logit-response Dynamics.” *Games and Economic Behavior* 68:413-27.
- Axtell, Robert, Joshua M. Epstein, and H. Peyton Young. 2000. “The Emergence of Classes in a Multi-agent Bargaining Model.” Pp. 191-211 in *Social Dynamics*, edited by S. N. Durlauf and P. Young. Cambridge, MA: MIT Press.
- Berninghaus, Siegfried K., Karl-Martin Ehrhart, and Claudia Keser. 2002. “Conventions and Local Interaction Structures: Experimental Evidence.” *Games and Economic Behavior* 39:177-205.

- Bicchieri, Cristina. 2005. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, United Kingdom: Cambridge University Press.
- Binmore, Ken and Larry Samuelson. 1999. "Evolutionary Drift and Equilibrium Selection." *The Review of Economic Studies* 66:363-93.
- Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural-change as Informational Cascades." *Journal of Political Economy* 100:992-1026.
- Blume, Lawrence E. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387-424.
- Butler, D., A. Isoni, and G. Loomes. 2012. "Testing the 'Standard' Model of Stochastic Choice under Risk." *Journal of Risk and Uncertainty* 45: 191-213.
- Camazine, S., J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. 2001. *Self-organization in Biological Systems*. Princeton, NJ: Princeton University Press.
- Cassar, Alessandra. 2007. "Coordination and Cooperation in Local, Random and Small World Networks: Experimental Evidence." *Games and Economic Behavior* 58:209-30.
- Chadefaux, Thomas. "The Triggers of War: Disentangling the Spark from the Powder Keg." Retrieved May 1, 2016 (http://www.thomaschadefaux.com/files/Chadefaux_Triggers.pdf).
- Cialdini, Robert B., Raymond R. Reno, and Carl A. Kallgren. 1990. "A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places." *Journal of Personality and Social Psychology* 58: 1015.
- Ehrlich, Paul R. and Simon A. Levin. 2005. "The Evolution of Norms." *PLoS Biology* 3:e194.
- Ellison, Glenn. 1993. "Learning, local Interaction, and Coordination." *Econometrica: Journal of the Econometric Society* 61:1047-71.
- Foster, Dean and Peyton Young. 1990. "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38:219-32.
- Freidlin, Mark I. and Alexander D. Wentzell. 2012. *Random Perturbations of Dynamical Systems*. New York: Springer Science & Business Media.
- Frey, Vincenz, Rense Corten, and Vincent Buskens. 2012. "Equilibrium Selection in Network Coordination Games: An Experimental Study." *Review of Network Economics* 11(3).
- Fudenberg, Drew and David K. Levine. 1998. *The Theory of Learning in Games*. Vol. 2. Cambridge, MA: MIT Press.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. 2005. "Regular Quantal Response Equilibrium." *Experimental Economics* 8:347-67.

- Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey. 2007. "Self-correcting Information Cascades." *Review of Economic Studies* 74:733-62.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420-43.
- Harsanyi, J. C. and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press
- Hart, Sergiu and Andreu Mas-Colell. 2003. "Uncoupled Dynamics Do Not Lead to Nash Equilibrium." *The American Economic Review* 93:1830-1836.
- Huckfeldt, Robert, Paul E. Johnson, and John Sprague. 2004. *Political Disagreement. The Survival of Diverse Opinions within Communication Networks*. Cambridge, United Kingdom: Cambridge University Press.
- Hwang, S. H. and J. Newton. 2014. "A Classification of Bargaining Solutions by Evolutionary Origin." University of Sydney. Economics Working Paper Series. Retrieved August 1, 2015 (<http://hdl.handle.net/2123/9993>).
- Judd, Stephen, Michael Kearns, and Yevgeniy Vorobeychik. 2010. "Behavioral Dynamics and Influence in Networked Coloring and Consensus." *Proceedings of the National Academy of Sciences* 107:14978-82.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica: Journal of the Econometric Society* 61:29-56.
- Kearns, Michael, Siddharth Suri, and Nick Montfort. 2006. "An Experimental Study of the Coloring Problem on Human Subject Networks." *Science* 313:824-7.
- Kirchkamp, O. and R. Nagel. 2007. "Naive Learning and Cooperation in Network Experiments." *Games and Economic Behavior* 58:269-92.
- Klemm, K., V. M. Eguiluz, R. Toral, and M. S. Miguel. 2003. "Global Culture: A Noise-induced Transition in Finite Systems." *Physical Review E* 67:045101(R).
- Kollock, P. 1993. "'An Eye for an Eye Leaves Everyone Blind': Cooperation and Accounting Systems." *American Sociological Review* 58:768-86.
- Lim, Wooyoung, and Philip R. Neary. 2016. "An Experimental Investigation of Stochastic Adjustment Dynamics." *Games and Economic Behavior* 100: 208-19.
- Loomes, G. 2005. "Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data." *Experimental Economics* 8:301-23.
- Macy, M. and M. Tsvetkova. 2015 "The Signal Importance of Noise." *Sociological Methods & Research* 44:306-28.
- Mäs, Michael, Andreas Flache, and Dirk Helbing. 2010. "Individualization as Driving Force of Clustering Phenomena in Humans." *PLoS Computational Biology* 6: e1000959.
- Mäs, Michael and Heinrich H. Nax. 2016. "A Behavioral Study of 'Noise' in Coordination Games." *Journal of Economic Theory* 162:195-208.

- McCubbins, Mathew D., Ramamohan Paturi, and Nicholas Weller. 2009. "Connected Coordination Network Structure and Group Coordination." *American Politics Research* 37:899-920.
- McFadden, Daniel. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." Pp. 105-142 in *Frontiers in Econometrics*, edited by P. Zarembka. New York: Academic.
- McKelvey, Richard D. and Thomas R. Palfrey. 1995 "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10:6-38.
- Montanari, Andrea, and Amin Saberi. 2010. "The Spread of Innovations in Social Networks." *Proceedings of the National Academy of Sciences* 107:20196-201.
- Naidu, S., H. S. Hwang, and S. Bowles. 2010. "Evolutionary Bargaining with Intentional Idiosyncratic Play." *Economics Letters* 109:31-33.
- Nicolis, G. and I. Prigogine. 1977. *Self-organization in Nonequilibrium Systems: From Dissipative Structures to Order through Fluctuations*. New York: Wiley-Interscience.
- Opp, Karl-Dieter. 2001. "How do Norms Emerge? An Outline of a Theory." *Mind & Society* 2:101-28.
- Pineda, Miguel, Raul Toral, and Emilio Hernandez-Garcia. 2009. "Noisy Continuous-opinion Dynamics." *Journal of Statistical Mechanics: Theory and Experiment* 2009:P08001.
- Pomeau, Y. 1986. "Front Motion, Metastability and Subcritical Bifurcations in Hydrodynamics." *Physica D* 23:3-11.
- Pradelski, B. S. R. and H. P. Young. 2012. "Learning Efficient Nash Equilibria in Distributed Systems." *Games and Economic Behavior* 75:882-97.
- Rabinovich, M. I., R. Huerta, P. Varona, and V. S. Afraimovich. 2008. "Transient Cognitive Dynamics, Metastability, and Decision Making." *PLoS Computational Biology* 4.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Selten, Reinhard and J. Apesteguia. 2005. "Experimentally Observed Imitation and Cooperation in Price Competition on the Circle." *Games and Economic Behavior* 51:171-92.
- Selten, Reinhard. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory* 4:25-55.
- Treiber, M., A. Hennecke, and D. Helbing. 2000. "Congested Traffic States in Empirical Observations and Microscopic Simulations." *Physical Review E* 62: 1805-24.
- van de Rijt, A., D. Siegel, and M. Macy. 2009. "Neighborhood Chance and Neighborhood Change: A Comment on Bruch and Mare." *American Journal of Sociology* 114:1166-80.

- Weber, Max. 1978. *Economy and Society: An Outline of Interpretive Sociology*. Berkeley, CA: University of California Press.
- Wilcox, Nathaniel. 2008. "Stochastic Models for Binary Discrete Choice under Risk: A Critical Primer and Econometric Comparison." Pp. 197-292 in *Research in Experimental Economics Vol. 12: Risk Aversion in Experiments*, edited by J. C. Cox and G. W. Harrison. Bingley, UK: Emerald.
- Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica: Journal of the Econometric Society* 61:57-84.
- Young, H. Peyton. 1998. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press
- Young, H. Peyton. 2011 "The Dynamics of Social Innovation." *Proceedings of the National Academy of Sciences* 108:21285-91.
- Young, H. Peyton. 2015 "The Evolution of Social Norms." *Annual Review of Economics* 7:359-87.

Author Biographies

Michael Mäs is an assistant professor at the Department of Sociology and the ICS at the University of Groningen. Before he was a senior scientist at ETH Zurich. His general research interest concerns collective action and social integration in social networks. He uses computational modeling techniques to derive new hypotheses and tests them with network data and laboratory experiments.

Dirk Helbing is a professor of Computational Social Science at ETH Zurich. Before he was working as Professor of Sociology at ETH Zurich and as Managing Director of the Institute for Transport & Economics at TU Dresden. Dr. Helbing is an elected member of the German Academy of Sciences "Leopoldina" and leads the PhD school on "Engineering Social Technologies for a Responsible Digital Future" at TU Delft.