Article

Sociological Methods & Research
00(0) 1-23
© The Author(s) 2013
Reprints and permission:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0049124113508093
smr.sagepub.com

The Signal Importance of Noise

Michael Macy[1] and Milena Tsvetkova[1]

Abstract

Noise is widely regarded as a residual category—the unexplained variance in a linear model or the random disturbance of a predictable pattern. Accordingly, formal models often impose the simplifying assumption that the world is noise-free and social dynamics are deterministic. Where noise is assigned causal importance, it is often assumed to be a source of inefficiency, unpredictability, or heterogeneity. We review recent sociological studies that are noteworthy for demonstrating the theoretical importance of noise for understanding the dynamics of a complex system. Contrary to widely held assumptions, these studies identify conditions in which noise can increase efficiency and predictability and reduce diversity. We conclude with a methodological warning that deterministic assumptions are not an innocent simplification.

Keywords

agent-based models, analytical sociology, complex systems, computer simulation, determinism, stochastic models, game theory, random error

Introduction

> *O! many a shaft, at random sent, Finds mark the archer little meant; And many a word, at random spoken, May soothe or wound a heart that's broken!*
>
> Sir Walter Scott (1771–1832)

[1] Cornell University, Ithaca, NY, USA

Corresponding Author:
Michael Macy, Cornell University, Ithaca, NY 14853, USA.
Email: mwm14@cornell.edu

Science looks for nonrandom patterns that signal the workings of an underlying causal process. Those patterns are often overlaid with noise that needs to be peeled away to reveal the underlying signal. In statistics, it is the central tendency in the underlying population that motivates the inquiry, not the random variations introduced by the luck of the draw. In linear models, the error term is the unexplained variance, which analysts usually ignore so long as the errors are random. In measurement models, we expect and tolerate some amount of noise, so long as it does not obscure the signal that we are looking for. The signal is the object of inquiry, not the noise.

We take a different viewpoint. The "signal importance of noise" is that it can be decisively important for understanding outcomes. Strip away the noise and you may strip away the explanation. This is not because the world that we are modeling is noisy. There is no problem assuming away random deviation from the central tendency in a distribution of observed outcomes, whether that variation corresponds to an inherently random process or to a systematic process that is simply not yet understood. Rather, we seek to call attention to the flip side of the issue: There is a problem assuming away random deviation from the behavior that is assumed to generate an outcome, whether that deviation corresponds to an inherently random process or to a systematic process that is simply not yet understood. Simply put, our argument is not about whether the world is noisy (although we happen to believe that it is). Our argument is about the theoretical reliability of the results derived from deterministic theoretical models in which noise is assumed away.

Noise does not mean "unpredictable," it means "uncertain." Uncertainty need not imply that the events occur rarely, unintentionally, or unpredictably. Uncertainty can characterize an event with any likelihood, large or small, so long as the probability is not zero or one. Noisy events may be intended (as when people toss a coin to choose a restaurant) or unintended (e.g., mistakes and misperceptions). Noisy events can be predictable, as when the probability distribution over possible outcomes is known, or unpredictable when the distribution is unknown. Whether rare or typical, intended or accidental, expected or surprising, a noisy event cannot be predicted with certainty.

Consider a cage filled with a finite number of red and blue balls, where every ball is uniquely numbered and all but one ball is blue. We then draw a ball at random. If we know the number of red and blue balls, we then know the probability to draw a blue ball, and as the proportion of blue balls increases, that probability approaches one asymptotically. However, no matter how close that probability comes to one, we do not know with certainty what color will be drawn, even in the case where blue is almost (but not quite) certain to be chosen. That is because every ball in the cage has an equal probability to be chosen, even though the selection of a ball of a given color is not equally likely. Because every ball has an equal probability to be chosen, the sequence of draws will be random, meaning that the events have zero autocorrelation and are independent and identically distributed.

It could be the case that we do not know the number of red and blue balls and hence we do not know the probability to select blue, but this condition is unrelated to whether every ball in the cage has an equal chance to be selected. In either condition—known or unknown probability—the process is stochastic, not deterministic, and the outcome of every draw is inherently uncertain.

The presence or absence of noise is the distinguishing feature of stochastic and deterministic models. Since noise is unbiased, it is tempting to assume it away in order to capture the central tendency or latent pattern that is the object of inquiry. Analytical models based on ordinary differential equations have often relied on a simplifying assumption that the processes under investigation are deterministic, an assumption that makes the models more mathematically tractable. Computational models also tend to be deterministic which makes them easier to debug and allows for dynamics in which equilibrium outcomes can be strictly identified as the absence of change.

Nevertheless, we caution against this simplification, not because deterministic models are empirically less plausible (a problem that we leave for a separate investigation), but because noise can have important consequences for system dynamics that cannot be safely ignored. Historically rooted in physics, this idea is fundamental to the kinetic theory of gasses, thermodynamics, and statistical mechanics (e.g., Prigogine 1976), but is less prominent in the social sciences, and in particular, sociology. Nevertheless, growing interest in analytical sociology—the study of causal processes in dynamical and complex social systems—has led to increasing awareness of the possibility for very small perturbations to behavioral rules or local structure to have highly nonrandom consequences for the dynamics of a complex system. In particular, while noise is often associated with inefficiency, unpredictability, and diversity, under certain conditions, random perturbations can lead to outcomes that are more efficient, more predictable, and less diverse. We review studies in which noise has these highly counterintuitive effects. In these applications, a deterministic model might not only be less empirically plausible, it would be theoretically misleading, in pointing to logical implications of a set of assumptions that do not follow when the determinism is relaxed.

We start with a brief methodological overview of analytical and computational approaches to modeling noise in social interaction, beginning with game theory. Game theorists have long been aware of the explanatory importance of noise. Game theorists model random deviations from a best response strategy caused by miscalculation, misperception (a "trembling eye" that misperceives which button was pressed), or misimplementation (a "trembling hand" that presses the wrong button). The best response is the strategy that provides the optimal outcome for a player, given the strategies that the other players have chosen. When all players choose their best response, a Nash equilibrium is said to obtain. Players might deviate from the best response for several reasons. First, perfectly rational players may nevertheless guess incorrectly about others' intentions, as when two people try to pass in a doorway (van de Rijt and Macy 2009). Second, they may misperceive or misinterpret the strategies of others (Axelrod and Dion 1988), due to imperfect monitoring of others' behavior or mistakes in processing the information. Finally, even if an actor correctly calculates the best response, mistakes can be made in implementation, characterized by Selten (1975) as a random "tremble" of the hand.

Since noise can represent different behaviors, it can be modeled in different ways (Helbing 2010). Most commonly, the random deviations are assumed to be independent in the population and unbiased with respect to the underlying behavior. This is implemented by assigning an identical error probability to all actors and by drawing the nonoptimal action uniformly at random from all other possible actions. If, instead, the deviations tend to be biased, the nonoptimal action can be drawn using a different distribution (compare, for instance, the effect of uniformly distributed noise in Pineda, Toral, and Hernández-García (2009) to the effect of normally distributed "white" noise in Mäs, Flache, and Helbing 2010). Further, a deviation from the best response strategy may be more or less likely under different conditions. For example, when there is a lot at stake, actors are hypothesized to be more likely to deliberate carefully and less prone to make costly mistakes. This has been implemented by modeling the probability to deviate from the best response as a function of the payoff (e.g., Binmore and Samuelson 1994; McFadden 1973; Montanari and Saberi 2010; Seymour 2000; Young 2011).

These advances in game theory have led to increasing awareness of the importance of noise in other applications as well. In particular, computational modeling has allowed researchers to systematically investigate the effects of noise in models of complex systems. Computational methods allow researchers to relax simplifying assumptions needed for mathematical tractability. In particular, agent-based computational (ABC) models have enabled researchers to relax the simplifying assumption of a noise-free world. In contrast to equation-based analytical approaches, ABC modeling replaces a single model of the population with a population of models, each corresponding to an autonomous but interdependent actor (or "agent"). As in game theory, the modeler supplies a set of micro-level assumptions about the properties of heterogeneous agents who then interact under constraints corresponding to a set of contextual assumptions such as network structure. The properties of the population emerge out of the interactions among the agents. This method allows investigators to identify the logical implications at the population level of a set of micro-level behavioral assumptions. As with game theory, these implications are often highly counterintuitive and motivate empirical studies to test the hypotheses suggested by the model (Clark 1991; He and Wong 2004; Macy 1995; Ruoff and Schneider 2006; Strang and Still 2004; Valente 1996; Willer, Kuwabara, and Macy 2009).

Although ABC models can be purely deterministic, most models are stochastic. The outcomes of a stochastic model are typically a probability distribution rather than a single ineluctable result. In this review, we survey stochastic models in which the central tendency of the distribution of outcomes turns out not to resemble the unique outcome in an otherwise identical deterministic model. Examples include paradoxical effects in which noise:

- increases both the diversity of new technologies and organizational practices and their average success,
- allows cultural diversity to collapse and cultural diffusion to diversify,
- makes residential integration more likely in populations that are less tolerant of ethnic minorities,
- makes collective action more likely to occur but less likely to diffuse, and
- accelerates the spread of viral information and disease, yet retards and obstructs the diffusion of costly or risky innovations.

We have grouped these studies around three surprising implications of noise that have been identified in social science applications of computational models: Noise can increase efficiency, improve predictability, and decrease diversity.

Noise and Efficiency

Noise is often assumed to interfere with efficiency and optimality, but recent studies show that it can also have the opposite effect. We begin by examining models in which noise is socially inefficient and then consider some paradoxical situations in which noise can be individually and/or collectively beneficial.

Even small amounts of randomness can drive a system away from an optimal equilibrium, especially if this equilibrium is fragile. The ultimate example is an equilibrium between nuclear superpowers based on a strategy of instant massive retaliation (known as mutually assured destruction [MAD]). Although the world has so far survived the danger, the equilibrium is vulnerable to the inability to entirely eliminate the possibility of an accident with fatal consequences for human civilization.

A more familiar example is an accidental affront that triggers an endless cycle of recrimination between feuding neighbors. Citing Ghandi's original insight that an "an eye for an eye leaves everyone blind," Kollock (1993) showed how small amounts of noise can lead to the collapse of cooperation in repeated play of the Prisoner's Dilemma. In this game, players have a choice between two actions, "cooperate" and "defect." Although cooperation yields the maximum mutual benefit, each actor is always individually better off by defecting in a single play of the game. That is not true, however, if the game is ongoing. In a famous study, Robert Axelrod showed how the cumulative benefit from mutual cooperation in future interactions may outweigh the temptation to defect in the current moment, which Axelrod calls "the shadow of the future" (Axelrod 1984; Axelrod and Hamilton 1981). Axelrod invited leading experts in game theory to submit strategies to compete against each other in a computerized tournament. The winner was TIT FOR TAT, a simple strategy that starts by cooperating but thereafter responds in kind to the partner's previous action. The strategy is successful because it is nice (it does not defect unless provoked) and also forgiving (it will resume cooperation if the opponent stops defecting), but it is not naive (it always retaliates when provoked).

However, Kollock (1993) demonstrated a potentially devastating weakness. TIT FOR TAT is vulnerable to random errors that can trigger an endless cycle of retaliation (Molander 1985; Reeves and Pitts 1996). Using computer simulation, he showed that more generous and forgiving strategies (e.g., TIT FOR TWO TATS) can perform better in the presence of small amounts of noise. Signorino (1996) identified a related vulnerability. Not only is TIT FOR TAT insufficiently forgiving, it also lacks contrition, that is, it is unable to accept punishment after its own unprovoked defection. While forgiveness is required to prevent endless cycles of recrimination triggered by the partner's "trembling hand" or one's own misperception, contrition is needed to correct for one's own miscue or the partner's misperception.

Helbing and Yu (2009) demonstrated another way that noise can undermine cooperation in the Prisoner's Dilemma. In a noise-free world in which players imitate the strategy of their most successful neighbor, cooperation can be sustained when cooperators are clustered on a spatial network. However, cooperation collapses when a random mutation causes a cooperator to defect. The higher payoff to defectors then causes defection to quickly spread.

Noise can also be incorporated in games as a mixed strategy equilibrium. For example, in the game of Chicken, two drivers speed toward each other, and each has two choices: to swerve to avoid the oncoming driver or stay the course. The collectively optimal equilibrium outcome is for one player to swerve to avoid a collision and the other to stay the course. This equilibrium is collectively optimal in that there is no alternate outcome that is as good or better for both players simultaneously. Game theorists refer to this as a *pure strategy* equilibrium because each player is committed to just one of the two possible choices. There is also a mixed strategy equilibrium in which each player assigns a positive probability to swerving and staying, such that both players are indifferent between the two actions. However, there is now the risk that both players will randomly choose to stay the course, resulting in a collision that both would prefer to avoid. Thus, both players will be better off to coordinate their actions on a pure strategy equilibrium in which one driver gives the right of way to the other and nothing is left to chance.

These studies illustrate the danger of assuming noise away in a deterministic model of social interaction. In the absence of noise, cooperation can be sustained by strategies of reciprocity/retaliation, by clustering of cooperators, and by the willingness of one side or the other to avoid an action that might lead to a catastrophe. However, the deterministic assumption turns out to be a surprisingly fragile simplification. Any nonzero error rate is sufficient to trigger the collapse of social order under conditions in which cooperation would otherwise thrive.

The vulnerability to random error demonstrated by Kollock and by Helbing and Yu resonates with our naive intuition that noise—like snakes, forest fires, and the common cold—is generally something that we would prefer to avoid. Yet, like these other seeming annoyances (which are themselves not necessarily noisy), it turns out that noise can also be indispensable, as shown in the studies we turn to next.

Although random errors usually prevent rational actors from making optimal choices, game theorists have identified a situation in which deliberate randomization is the individually optimal strategy—when we need to keep other players guessing about our intentions. A good example is the penalty kick in soccer (Chiappori, Levitt, and Groseclose 2002). If the kicker favors one direction and the goalie knows this, then the goalie will have an advantage, and vice versa if the kicker knows that the goalie has a bias. Thus, both sides are better off to choose directions randomly, by playing a mixed strategy, as in "mixing it up." By keeping the other player guessing, mixed strategies are the unique best response in zero-sum games like the penalty kick, "matching pennies," and "rock, paper, scissors." In zero-sum games, one player's loss is another player's gain. A mixed strategy can also be a best response in games in which there is an opportunity for mutual gain and mutual loss.

There are other situations in which random errors can improve the collective as well as individual outcome. A simple (and somewhat trivial) example is the flip side of the problem identified by Kollock (1993). If people defect in iterated Prisoner's Dilemma only because they do not trust others to cooperate, a random act of cooperation can generate an endless cycle of reciprocation in the same way that random defection can trigger an endless cycle of retaliation (Bendor, Kramer, and Swistak 1996). Helbing and Yu (2009) discover similar effects of noise in a spatial Prisoner's Dilemma game with imitation and migration. Players imitate successful neighbors and migrate to empty locations where their strategy would be more successful. Under these assumptions, the authors found that noise can trigger a transition from a collectively suboptimal to optimal equilibrium. In a world where everyone almost always defects, random mutation eventually creates a small cluster with enough cooperators to make cooperation viable, despite the vulnerability to exploitation by defectors. "Success-driven migration" attracts defectors but also allows cooperators to protect themselves from invasion by becoming densely packed. In this way, cooperation can spontaneously emerge through a chance event that triggers a cascade. The authors conclude: "The level of cooperation decreases with the noise strength [...] but moderate values [of noise] can even accelerate the transition to predominant cooperation" (p. 7).

More generally, noise can disrupt premature lock-in on a suboptimal equilibrium and allow agents to discover a superior solution. This possibility lies behind the process of annealing in metallurgy—heating the metal introduces noise to the movement of atoms, which allows them to rearrange, which in turn results in fewer imperfections in the crystal structure. Processes like this led computer scientists and mathematicians to develop search algorithms based on stochastic optimization. If the search stalls near a suboptimal solution, random perturbation spontaneously pushes the algorithm toward uncharted areas of the search space (Spall 2003).

Biological evolution is a compelling example of stochastic optimization in nature, in which random errors (in the form of genetic mutations) contribute to the search for adaptive solutions. Random mutations can increase or decrease an organism's fitness and therefore have no inherently adaptive function. It is natural selection, not mutation, that allows evolution to find better solutions to adaptive problems. However, without mutations, natural selection will eventually exhaust the genetic diversity of the population. Without heterogeneity, recombination is unable to build on partial solutions to find more adaptive strategies. Mutation restores heterogeneity in the face of selection pressures that tend to reduce it. Too much mutation of course is also suboptimal, but a small amount is essential to allow evolutionary exploration of an uneven fitness landscape that can trap the population on a "false peak."

Principles of evolutionary search have been applied to technological innovation and organizational learning, such as March's model of the trade-off between exploitation and exploration (March 1991; see also Helbing, Treiber, and Saam 2005; Lazer and Friedland 2007; Miller, Zhao, and Calantone 2006). March argues that although "the exploitation of old certainties" has short-term benefits, "the exploration of new possibilities" improves organizational performance in the long term. Exploitation involves copying best practices, analogous to the process of natural selection in biological evolution and with the same problem of premature lock-in on suboptimal solutions. By allowing for random search as well as systematic independent inquiry, exploration increases the diversity of possible solutions competing for selection/exploitation.

Noise and Predictability

In linear models, the slope parameters allow prediction of an outcome based on the state of one or more covariates, while the error term represents random deviation of the observations from the predicted values. Similarly, random measurement error obstructs accurate identification of an object of investigation. More generally, if an event is random, then a particular occurrence of the event cannot be predicted, even if the probability distribution is known.[1] Intuition therefore suggests that the outcome of a process will become less predictable as the level of noise in that process increases. Nevertheless, in this section, we review studies of social dynamics in which the introduction of small amounts of noise makes the outcomes more rather than less predictable.

Granovetter's (1978) threshold model of collective behavior provides a compelling illustration. Granovetter modeled a threshold as the critical number of participants at which an individual becomes willing to join a collective behavior. The weaker the individual's interest in a successful mobilization, the greater the tendency to wait to see how many others are willing to participate. Depending on the distribution of individual thresholds, cascades are possible in which each additional participant triggers participation by others. Granovetter showed how a cascade can stall if it reaches a gap in the distribution of thresholds. To take an extreme example, suppose everyone has a threshold of $N - 1$, that is, no one will participate unless everyone else does. Not surprisingly, collective behavior in this population is very likely to fail. Equally unsurprising, if everyone has a threshold of zero, the collective behavior is guaranteed to succeed. But suppose everyone has a threshold of one. Although the level of interest in a successful mobilization is nearly as strong as the threshold-zero case, the threshold-one outcome is identical to the case of threshold $N - 1$—everyone looks around to see if someone else will be the first mover and therefore no one participates. These stylized illustrations demonstrate a key insight of Granovetter's article: The equilibrium level of participation need not correspond to the level of interest in successful mobilization.

The outcome in each of these examples is a deterministic Nash equilibrium. It is a Nash equilibrium because the level of participation is such that each individual has chosen their best response given the actions of others. It is deterministic because the outcome is completely controlled by the values assigned to the model parameters and hence will always be the same, no matter how many times the experiment is repeated. Hence, the equilibrium level of participation can never change.

Now suppose instead that the model is stochastic rather than deterministic. Specifically, suppose everyone in the threshold-one population has some epsilon probability to participate, no matter how many others have already done so. Again, everyone is waiting around for someone to go first. But this time, someone eventually acts by chance rather than by threshold, triggering a cascade. The outcome in the threshold-one case is now almost indistinguishable from the threshold-zero case and very different from the case of threshold $N - 1$. This shows that Granovetter's insight, while theoretically interesting, rests on a deterministic assumption that is empirically implausible.

The stochastic version of the threshold model yields a counterintuitive insight of its own. For very small epsilon, each individual has only a very small probability of experiencing a random event, but the probability that someone will trigger a cascade increases exponentially with *N*. (Technically, the probability of the trigger event is $1 - (1 - \varepsilon)^N$). This insight illustrates one of the most important discoveries in the theoretical literature on collective action—"the paradox of group size" (Oliver and Marwell 1988). The larger the population, the higher the probability that there will be a chance outlier (or perhaps a small group of outliers) who is willing to go first.

The introduction of noise into Granovetter's deterministic model has an additional surprise: System behavior becomes *more* predictable, not less. If we know everyone's threshold, the deterministic model can predict the outcome of the mobilization perfectly. The problem arises when thresholds are unknown. Although Granovetter's aim was not to model predictability, a key result is that, depending on the location of the gaps in the distribution of thresholds, every possible outcome—from zero to *N* participants—is a potential equilibrium, regardless of the average level of interest in the mobilization. It is not hard to imagine that observers might have a reasonable estimate of the average interest in a successful mobilization, from which we can estimate whether the average threshold is closer to 0 or to $N - 1$. In contrast, it is much more difficult to locate in advance the gaps in the distribution of thresholds that might cause a cascade to stall. In the deterministic model, the outcome is highly sensitive to those gaps, and therefore difficult to predict. But if we assume that the process is stochastic, these deterministic equilibria can be disturbed by a random, idiosyncratic decision to participate, allowing the cascade to proceed. As a result, the long-term tendency of the stochastic process can be expected to correspond with the level of interest in the outcome. In short, adding noise makes the outcome more predictable, given prior knowledge of the level of interest in the outcome among the members of the population. Or to put it the other way around, the assumption that the world is noisy allows us to have greater confidence that the likelihood of a successful collective action increases with the level of interest in success. If we lived in Granovetter's deterministic world, that confidence would be unwarranted, for the reason that Granovetter identifies.

Granovetter did not relax the deterministic assumption and did not note the possibility that noise can make outcomes more predictable. Game theorists, in contrast, have developed this paradoxical effect of noise into an important line of theoretical research, beginning with the pioneering work of the evolutionary biologists Smith and Price (1973). Smith and Price used game theory to model an evolutionary process in which Nature selects strategies based on their average payoffs. The evolutionary application confronts an important limitation of classical game theory—the assumption that players exercise rational foresight. This assumption may be warranted for games played by expert strategists in business, politics, and the military who have learned how to use backward induction to identify their best initial move by reasoning back from the last move in a sequence. However, this assumption introduces a teleological error for optimizing processes that rely on trial and error, not rational foresight. Building on biological insights, Smith and Price (1973) proposed "evolutionary stability" as a refinement that eliminated any Nash equilibrium based on a strategy that, once adopted by every player, could still be invaded by a randomly mutated alternative. For example, in repeated play of Prisoner's Dilemma, TIT FOR TAT is a Nash equilibrium but it is not evolutionarily stable because it can be invaded by a random mutant variation that always cooperates, which in turn can make the population vulnerable to invasion by another mutant that always defects.

At about the same time as the discovery of evolutionary stability, Selten (1975) proposed another refinement—the "trembling hand"—that eliminates any equilibrium that cannot withstand the effects of random perturbation. Suppose two friends enjoy meeting at a restaurant for dinner more than eating alone at home, but they really hate eating alone at a restaurant. This game has two Nash equilibria in which the friends eat together, whether at home or at a restaurant. However, eating at the restaurant is not trembling hand perfect (assuming a world without cell phones). That is because there is always a small chance that the friend will not show up due to some unforeseen chance event.

Foster and Young (1990) proposed a related refinement of the Nash equilibrium called *stochastic stability*. They started by supposing that the system is *continually* exposed to perturbations. Then, a stochastically stable equilibrium is one that prevails in the long run as the rate of perturbation approaches zero. In the eating out example, eating at home is the risk-dominant strategy. Foster and Young showed that, as time goes to infinity and the probability of error slowly vanishes, the risk-dominant strategy will be chosen almost all the time.

Although the abstract representation of these problems in game theory may make the solutions appear artificial to empirical social scientists, all three equilibrium refinements—evolutionary stability, trembling hand, and stochastic stability—manifest a common underlying principle that is intuitively appealing: *an equilibrium must withstand the effects of random perturbation.* As Granovetter's threshold model clearly illustrates, an equilibrium that is not robust to noise does not have much predictive power. By removing these fragile equilibria from the set of possible solutions, the predictive power of the model can be increased, often dramatically.

Noise, Diversity, and Diffusion

We have seen how noise can increase (as well as decrease) efficiency, and increase (as well as decrease) predictability. We now turn to studies in which noise can increase (as well as decrease) diversity. In these models, people come to be more similar to their neighbors in a dynamic social network, either because they change their attributes to match their neighbors, change their neighbors to match their attributes, or both. We begin with studies of cultural assimilation, then turn to ethnic segregation, and conclude with models of diffusion.

Cultural Assimilation

The earlier discussion of noise and optimization called attention to a fundamental principle of evolution by natural selection—the need for random mutation to restore genetic diversity in the face of selection pressures that tend to reduce it. The diversifying effect of random variation is highly intuitive, yet it turns out that noise can also have the opposite effect—causing diversity to decline or completely disappear (Flache and Macy 2011).

This was demonstrated recently by Klemm et al. (2003a, 2003b; cf. De Sanctis and Galla 2009) in an article that extended Axelrod's (1997) earlier research on cultural assimilation. Axelrod's model addresses the paradox that cultural diversity is both persistent and precarious. Although Greig (2002) stresses the robustness of cultural minorities to the forces of assimilation, many others accept the more pessimistic outlook about the future of diversity expressed by the United Nations Educational, Scientific, and Cultural Organization (UNESCO) in a 2001 white paper. For example, there are about 6,800 languages worldwide, but about half that number are expected to become extinct (Crystal 2000). Axelrod investigated the dynamics of assimilation and diversity using agent-based models that combined two widely observed social mechanisms: homophily (the tendency to interact with similar others) and social influence (the tendency to become more similar to those with whom we interact). These processes imply a self-reinforcing dynamic that might be expected to drive populations toward cultural homogeneity. However, Axelrod showed that this is not necessarily the case, because of what Flache and Macy (2011) call "a cultural analog of biological speciation." If two people are dissimilar on every salient cultural dimension, Axelrod argues, they can no longer interact, analogously to the inability of sexually reproducing organisms with a common ancestor to mate once they differentiate beyond a critical threshold. As a consequence, homophily and social influence can "speciate" stable cultural regions that are protected from assimilation. In short, the same process that generates local convergence also preserves diversity at the global level.

Although Axelrod's explanation is elegant and compelling, it turns out to depend decisively on a deterministic assumption—that cultural traits are not susceptible to random mutation. Klemm and collaborators (2003a, 2003b; cf. De Sanctis and Galla 2009) introduced the empirically more plausible possibility for random perturbation of cultural traits. Their model confirms what intuition would suggest—that random mutation increases local diversity. However, the increase in local heterogeneity can cause the collapse of cultural diversity at the global level. To see why, consider a simple case of a perfectly polarized population composed of two groups whose members are identical in every respect to other in-group members and exactly the opposite of every member in the out-group. This situation precludes cultural contact and influence between the groups that might otherwise lead to mutual assimilation. In other words, the polarized population is an equilibrium. However, as noted in the discussion of predictability, a deterministic equilibrium can be extremely brittle to random perturbation. Small amounts of cultural noise can create the common ground that makes interaction possible across otherwise impermeable cultural boundaries. This interaction, in turn, reduces remaining differences, leading to even more interaction. "Thus, formerly dissimilar neighbors become increasingly similar until no differences remain and a new cultural boundary forms around a larger region. Eventually this boundary too will be breached by a perturbation that creates a common trait between otherwise dissimilar neighbors, and so on, until no differences remain" (Flache and Macy 2011:972–973).

If the mutation rate is sufficiently high, the resulting cultural turbulence precludes the formation of stable cultural regions, a pattern Centola et al. (2007:918) characterize as "cultural anomie" to indicate the absence of conformity to local conventions. As Flache and Macy note, cultural anomie should not be confused with diversity. "The latter requires cultural convergence within internally homogenous regions that differ from one another sufficiently that they maintain their distinctiveness over time. Anomie means distinct regions cannot coalesce, due to the evanescence of individuals who differentiate themselves from one another more than they are attracted to one another" (Flache and Macy 2011:973).

In short, noise attacks diversity from two sides, according to Flache and Macy. If the noise rate is sufficiently low, perturbations breach the boundaries between cultural regions, promoting assimilation. If the rate is sufficiently high, perturbations preclude local convergence. Thus, cultural diversity is only viable within a narrow window of perturbation rates that closes asymptotically as population size increases (Flache and Macy 2011).

Ethnic Segregation

The effects of noise on cultural differentiation have also been investigated in the context of ethnic segregation of residential networks. In a seminal article in the *Journal of Mathematical Sociology*, Schelling (1971) showed how a tipping process can lead to complete segregation even in populations that are highly tolerant of ethnic diversity. However, like Granovetter and Axelrod, Schelling assumed that decision making was based on a deterministic threshold function. So long as the proportion of coethnic neighbors is below a critical threshold, no one ever moves, and above the threshold, they never stay.

In a widely celebrated article, Bruch and Mare (2006) argued that Schelling's results were not robust: The tipping process at the population level was an artifact of Schelling's assumption of a tipping process at the individual level, such that individuals do not react to changes in the ethnic composition of their neighborhoods except at a single threshold, such as 50 percent coethnic neighbors. Bruch and Mare's model captures the more empirically plausible assumption that people are far more sensitive to small changes in the ethnic composition of their neighborhood than Schelling assumed. When the authors replaced Schelling's threshold function with a continuous function, segregation largely disappeared. More precisely, segregation levels were lower in a population that notices every additional out-group neighbor, compared to a population that tolerates ethnic diversity up to a critical threshold.

This highly counterintuitive discovery led van de Rijt, Siegel, and Macy (2009) to carefully replicate Bruch and Mare's model. However, they found the opposite of what Bruch and Mare had reported, and Bruch and Mare (2009) eventually acknowledged that their original finding was an artifact of a coding error. Using a model identical to Bruch and Mare's, but without the "bug," van de Rijt et al. found that sensitivity to small changes in ethnic composition generally leads to more segregation, not less. The lower levels of segregation that Bruch and Mare report are caused by the additional assumption that the decision whether to move is highly random. In Schelling's model, decisions to move are completely determined by ethnic composition of the neighborhood. At the other extreme, Bruch and Mare assumed that decisions are heavily influenced by chance. It is this high level of noise, not the sensitivity to ethnic composition, that causes segregation to decline. Even a population of racists will tend to integrate if the level of random mixing is sufficiently high. By examining the full range of noise levels, from purely deterministic to purely random, van de Rijt et al. found that the effects of noise are amplified by sensitivity to ethnic composition. When people notice every additional out-group neighbor, random moves can trigger "error cascades" whose logic is identical to the tipping process revealed by Schelling. Just as a single decision to move *out* of a neighborhood with too many out-group neighbors can precipitate a cascade of "white flight," a random move *into* a neighborhood with too many out-group neighbors can trigger an error cascade that leads to ethnic integration. Sensitivity to small changes in ethnic composition can promote cascades in both directions—toward segregation when cascades are triggered by nonrandom decisions to leave a changing neighborhood, and toward integration when cascades are triggered by random decisions to move in. Thus, as the noise level increases, it becomes possible for error cascades to integrate members of an intolerant population, whose dissatisfaction with their current location increases with each additional out-group neighbor. Paradoxically, van de Rijt et al. showed how these cascades may be less likely in a tolerant population that is largely indifferent to small changes in ethnic composition.

Diffusion of Innovation

Random moves on a residential lattice, as modeled by Bruch and Mare and by van de Rijt et al., can also be modeled as a dynamic network in which ties between neighbors are randomly formed and broken. Random ties facilitate contact between otherwise distant clusters, which affects not only the dynamics of cultural assimilation and ethnic segregation but also the spread of disease, information, fads, fashions, and social movements.

The importance of these network bridges is illustrated in one of the most-cited articles in sociology—Granovetter's (1973) study of the "strength of weak ties." According to Granovetter, compared to ties between close friends, ties between acquaintances tend to involve lower trust, less frequent interaction, and weaker commitment. However, the strength of these weak ties is that they tend to bridge between network clusters, in contrast to the strong ties that interconnect a cluster of close friends. These bridges provide distant regions of a social network with access to information, which facilitates diffusion, promotes social integration, and explains the famously observed "six degrees of separation" between any two randomly chosen people on the planet. But how can this "small world" be reconciled with evidence that social networks are highly clustered, such that most people are embedded in a tightly knit "small circle of friends?"

The answer was discovered by Watts and Strogatz (1998; Watts 1999). They began with a highly clustered network and then introduced structural noise by replacing a randomly selected tie with a tie to a randomly chosen node. They then measured the time required for a contagion (e.g., a biological or cultural virus) to spread from one node to the entire network of connected nodes. As expected, with no random rewiring, contagions spread much more slowly than in a network in which every tie was random. The surprise was what happens in a "small world" network—one that remains highly clustered because only about 1 tie in 10 has been randomly perturbed. Watts and Strogatz found that contagions spread nearly as fast as in a completely random network. In other words, small amounts of randomness in the structure of a highly clustered network can significantly accelerate the spread of contagions that readily transfer from one individual to another, such as a contagious pathogen, news about a job opening, a viral rumor, or a catchy phrase.

This "small world" effect could be taken to suggest that it is relational randomness, and not the tightly patterned relations among close friends, that forms the "glue" connecting the social world. However, a follow-up study by Centola and Macy (2007) showed that network noise can also have an effect that is the opposite of what Watts and Strogatz discovered. In their small-world experiment, Watts and Strogatz assumed that contagions have a threshold of one, meaning that the contagion can spread through contact with a single infected neighbor. However, Centola and Macy remind us that many social contagions have thresholds greater than one. For example, hearing the same news from two or more friends is redundant, but hearing the same advice from two or more friends is reinforcing. The same is true for adopting a costly innovation, joining a risky social movement, or participating in a controversial practice such as using contraceptives in a deeply religious community. For these social contagions, the larger the number of others who have already adopted the behavior, the greater the sense of confidence and legitimacy conveyed to their network neighbors.

Centola and Macy (2007) began by replicating the original Watts and Strogatz experimental setup, randomly perturbing the structure of a highly clustered network, and observing the rate of propagation. They then repeated the contagion experiment, except that they increased the threshold of adoption from one infected neighbor (which they term a "simple contagion") to two or more (a "complex contagion"). Not only did random rewiring not accelerate diffusion, it completely prevented complex contagions from spreading. While simple contagions like information or disease benefit from random ties, the spread of complex contagions requires the redundant structure found in local clusters, which random rewiring erodes. Just as a small amount of network noise allows simple contagions to attain the maximum possible rate of propagation, it takes only a small amount of random rewiring to preclude the takeoff of a complex contagion.

Conclusion

Social scientists typically regard noise as a residual category—the unexplained variance in a linear model or the random disturbance of a systematic pattern. Accordingly, noise can appear to be nothing more than a meaningless or irrelevant behavior that obscures the underlying causal mechanisms in social life and can therefore be safely removed in order to more accurately identify and measure the object of inquiry.

Nevertheless, recent decades have witnessed increasing awareness of the explanatory importance of random error in social interactions. Our review of these studies showed how small perturbations to behavioral rules or local structure can lead to dramatic changes in the dynamics and the equilibria of a social system. In complex systems of interdependent action, mistakes can get amplified. A single error can trigger a cascade of behavioral changes or a migration into a new solution space and thus steer the dynamics in an entirely new and unexpected direction. Contrary to widely held assumptions, the studies we reviewed demonstrate that, under certain conditions, random perturbations can have highly paradoxical effects, making outcomes more efficient, more predictable, and less diverse:

- Noise can disturb a local equilibrium that is globally suboptimal, leading to a collectively preferred outcome.
- Eliminating fragile equilibria reduces the size of the solution set, making outcomes more predictable.
- While noise increases local heterogeneity, this can in turn facilitate social interactions that reduce global diversity.
- While random mobility can reduce neighborhood ethnic segregation by triggering "error cascades," these cascades are more likely in an ethnically sensitive population than in one where people are more tolerant of ethnic diversity.
- A few random ties can shrink our world for the spread of information and disease, but these weak ties are not conducive to the early spread

of social contagions such as controversial beliefs or high-risk collective behaviors.

The take-home message from these studies is that deterministic models that assume away noise should be approached with caution. If the predictions of a deterministic model are not robust to noise, they may conceal important dynamics, especially in multidimensional nonlinear systems. Noise can have highly counterintuitive effects that are difficult to recognize in models expressed in natural language. Game theoretic and ABC models have proved highly useful in the effort to reveal, identify, and analyze effects that often defy intuition. Nonetheless, a formal model 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.

Acknowledgment

The authors wish to acknowledge the contributions of Michael Mäs, Andreas Flache, Dirk Helbing, and members of the Social Dynamics Laboratory at Cornell.

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) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This article was prepared while the authors were supported by grant SES 1260348 from the U.S. National Science Foundation.

Note

1. While all random events are unpredictable even when the probability is known, unpredictable events are not necessarily random. For example, an observer who does not know the encryption key is unable to predict the next value in a nonrandom sequence of encrypted numbers.

References

Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, R. 1997. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41:203-26.

Axelrod, R., and D. Dion. 1988. "The Further Evolution of Cooperation." *Science* 242:1385-90.

Axelrod, R., and W. D. Hamilton. 1981. "The Evolution of Cooperation." *Science* 211:1390-96.

Bendor, J., R. Kramer, and P. Swistak. 1996. "Cooperation under Uncertainty: What is New, What is True, and What is Important." *American Sociological Review* 61: 333-38.

Binmore, K., and L. Samuelson. 1994. "An Economist's Perspective on the Evolution of Norms." *The Journal of Institutional and Theoretical Economics* 150: 45-63.

Bruch, E. E., and R. D. Mare. 2006. "Neighborhood Choice and Neighborhood Change." *American Journal of Sociology* 112:667-709.

Bruch, E. E., and R. D. Mare. 2009. "Preferences and Pathways to Segregation: Reply to Van de Rijt, Siegel, and Macy." *American Journal of Sociology* 114: 1181-98.

Centola, D., J. C. González-Avella, V.M. Eguíluz, and M. San Miguel. 2007. "Homophily, Cultural Drift, and the Co-evolution of Cultural Groups." *Journal of Conflict Resolution* 51:905-29.

Centola, D., and M. W. Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113:702-34.

Chiappori, P. A., S. Levitt, and T. Groseclose. 2002. "Testing Mixed-strategy Equilibria When Players are Heterogeneous: The Case of Penalty Kicks in Soccer." *American Economic Review* 92:1138-51.

Clark, W. A. V. 1991. "Residential Preferences and Neighborhood Racial Segregation: A Test of the Schelling Segregation Model." *Demography* 28:1.

Crystal, D. 2000. *Language Death*. Cambridge, England: Cambridge University Press.

De Sanctis, L., and T. Galla. 2009. "Effects of Noise and Confidence Thresholds in Nominal and Metric Axelrod Dynamics of Social Influence." *Physical Review E* 79:046108.

Flache, A., and M. W. Macy. 2011. "Local Convergence and Global Diversity: From Interpersonal to Social Influence." *Journal of Conflict Resolution* 55:970-995.

Foster, D., and P. Young. 1990. "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38:219-32.

Granovetter, M. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360-80.

Granovetter, M. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420-43.

Greig, J. M. 2002. "The End of Geography? Globalization, Communications, and Culture in the International System." *Journal of Conflict Resolution* 46:225-43.

He, Z.-L., and P-.K. Wong. 2004. "Exploration vs. Exploitation: An Empirical Test of the Ambidexterity Hypothesis." *Organization Science* 15:481-94.

Helbing, D. 2010. "The Future of Social Experimenting." *Proceedings of the National Academy of Sciences* 107:5265-66.

Helbing, D., M. Treiber, and N. J. Saam. 2005. "Analytical Investigation of Innovation Dynamics Considering Stochasticity in the Evaluation of Fitness." *Physical Review E* 71:067101.

Helbing, D., and W. Yu. 2009. "The Outbreak of Cooperation among Success-driven Individuals under Noisy Conditions." *Proceedings of the National Academy of Sciences* 106:3680-85.

Klemm, K., V. M. Eguíluz, R. Toral, and M. San Miguel. 2003a. "Global Culture: A Noise-induced Transition in Finite Systems." *Physical Review E* 67: 045101.

Klemm, K., V. M. Eguíluz, R. Toral, and M. San Miguel. 2003b. "Nonequilibrium Transitions in Complex Networks: A Model of Social Interaction." *Physical Review E* 67:026120.

Kollock, P. 1993. "'An Eye for an Eye Leaves Everyone Blind': Cooperation and Accounting Systems." *American Sociological Review* 58:768-86.

Lazer, D., and A. Friedman. 2007. "The Network Structure of Exploration and Exploitation." *Administrative Science Quarterly* 52:667-94.

Macy, M. W. 1995. "PAVLOV and the Evolution of Cooperation: An Experimental Test." *Social Psychology Quarterly* 58:74-87.

March, J. G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2:71-87.

Mäs, M., A. Flache, and D. Helbing. 2010. "Individualization as Driving Force of Clustering Phenomena in Humans." *PLoS Computational Biology* 6:e1000959.

McFadden, D. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." Pp. 105-42 in *Fronitiers in Econometrics*, edited by P. Zarembka. New York: Academic.

Miller, K. D., M. Zhao, and R. D. Calantone. 2006. "Adding Interpersonal Learning and Tacit Knowledge to March's Exploration-Exploitation Model." *Academy of Management Journal* 49:709-22.

Molander, P. 1985. "The Optimal Level of Generosity in a Selfish, Uncertain Environment. *Journal of Conflict Resolution* 29:611-18.

Montanari, A., and A. Saberi. 2010. "The Spread of Innovations in Social Networks." *Proceedings of the National Academy of Sciences* 107:20196-201.

Oliver, P. E., and G. Marwell. 1988. "The Paradox of Group Size in Collective Action: A Theory of the Critical Mass. II." *American Sociological Review* 53:1-8.

Pineda, M., R. Toral, and E. Hernández-García. 2009. "Noisy Continuous-opinion Dynamics." *Journal of Statistical Mechanics* 2009:P08001.

Prigogine, I. 1976. "Order through Fluctuations: Self-organization and Social System." Pp. 93-130 in *Evolution and Consciousness: Human Systems in Transition*, edited by E. Jantsch and C. H. Waddington. MA: Addison-Wesley.

Reeves, E. B., and T. C. Pitts. 1996. "Cooperative Strategies in Low-noise Environments." *American Sociological Review* 61:338-41.

Ruoff, G., and G. Schneider. 2006. "Segregation in the Classroom." *Rationality & Society* 18:95-117.

Schelling, T. C. 1971. "Dynamic Models of Segregation." *The Journal of Mathematical Sociology* 1:143.

Selten, R. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory* 4:25-55.

Seymour, R. M. 2000. "Stationary Distributions of Noisy Replicator Dynamics in the Ultimatum Game." *The Journal of Mathematical Sociology* 24:193-243.

Signorino, C. S. 1996. "Simulating International Cooperation under Uncertainty: The Effects of Symmetric and Asymmetric Noise. *Journal of Conflict Resolution* 40: 152-205.

Smith, M. J., and G. R. Price. 1973. "The Logic of Animal Conflict." *Nature* 246: 15-18.

Spall, J. C. 2003. *Introduction to Stochastic Search and Optimization*. Hoboken, NJ: John Wiley.

Strang, D., and M. C. Still. 2004. "In Search of the élite: Revising a Model of Adaptive Emulation with Evidence from Benchmarking Team." *Industrial and Corporate Change* 13:309-33.

UNESCO (United Nations Educational, Scientific, and Cultural Organization). 2001. *Universal Declaration on Cultural Diversity*. Adopted by the 31st Session of the General Conference of UNESCO, Paris, November 2, 2001. Accessed on October 16, 2013. http://unesdoc.unesco.org/images/0012/001271/127160m.pdf.

Valente, T. W. 1996. "Social Network Thresholds in the Diffusion of Innovations. *Social Networks* 18:69-89.

van de Rijt, A., and M. W. Macy. 2009. "The Problem of Order: Self-interest or Autonomy?" *Advances in Group Processes* 26:25-51.

van de Rijt, A., D. Siegel, and M. W. Macy. 2009. "Neighborhood Chance and Neighborhood Change: A Comment on Bruch and Mare." *American Journal of Sociology* 114:1166-80.

Watts, D. J. 1999. "Networks, Dynamics, and the Small–world Phenomenon." *American Journal of Sociology* 105:493-527.

Watts, D. J., and S. H. Strogatz. 1998. "Collective Dynamics of 'Small-world' Networks." *Nature* 393:440.

Willer, R., K. Kuwabara, and M. W. Macy. 2009. "The False Enforcement of Unpopular Norms." *American Journal of Sociology* 115:451-90.

Young, H. P. 2011. "The Dynamics of Social Innovation." *Proceedings of the National Academy of Sciences* 108:21285-91.

Author Biographies

Michael Macy is Goldwin Smith Professor of Arts and Sciences and Director of the Social Dynamics Laboratory at Cornell. His research uses computational models, online laboratory experiments, and digital traces of device-mediated interaction to explore familiar but enigmatic social patterns, such as circadian rhythms, the emergence and collapse of fads, the spread of self-destructive behaviors, cooperation in social dilemmas, the critical mass in collective action, the spread of high-threshold contagions on smallworld networks, the polarization of opinion, segregation of neighborhoods, and assimilation of minority cultures. These studies have been published in leading journals, including *Science, PNAS, American Journal of Sociology, American Sociological Review*, and *Annual Review of Sociology*.

Milena Tsvetkova is a doctoral candidate in the field of Sociology at Cornell University. Her research involves online experiments, computational modeling, and network analysis. Her current projects investigate the contagiousness of prosocial behavior and the emergence of segregation and inequality.