The Emergence of Local Norms in Networks

Mary A. Burke

Gary M. Fournier

Kislaya Prasad

Florida State University

August 17, 2004

Abstract

We develop an explanation of the emergence of local norms, and the associated phenomenon of geographical variation in behavior. Individuals are assumed to interact locally with neighbors in an environment with a network externality. Although many patterns of behavior are possible, the dispersed interactive choices of agents are shown to select behavior that is locally uniform but globally diverse. The range of applications of the theory includes regional variation in the practice of medicine, technology choice, and corruption. The framework is also useful for further developing our understanding of important phenomena like lock-in, critical thresholds, and contagion.

1 Introduction

Geographical variations in behavior are a persistent puzzle. People in seemingly similar situations often choose to do different things. Their choices depend upon where they live and the people they associate with. These circumstances give rise to uniformity of behavior within groups, together with global diversity across groups. Such differences are attributed, almost axiomatically, to differences in "culture" or "norms," as in matters of dress, speech, or driving, for example. For our purposes the term "norm" refers to an established behavior that is widespread, if not universal, within a particular community. It is also self-reinforcing: once the norm is in place it is in each individual's interest to conform to it, even though *ex ante* an alternative norm might have emerged.[1] In a given region, or within a given social group, there is the appearance of a socially agreed upon way of responding to situations. Across groups, however, differences in behavior persist even when the groups are not isolated. The challenge is to understand how such locally uniform, globally diverse, patterns emerge. To paraphrase Morris (2000), we want to know when and why we might expect "coexistent conventions."

Geographical variation in norms is pervasive. For example, the medical treatment that a patient receives depends, to an inordinate extent, on where he lives. Geographical variations in medical procedure choice were first documented by Glover (1938), and subsequent studies (e.g. Wennberg and Gittelsohn 1973, Wennberg and Cooper 1999) have confirmed the presence of "small area variations." Furthermore, choices appear relatively uniform at the local level (Burke, Fournier, and Prasad, 2004, hereafter BFP). In sharecropping, contracts between tenants and landlords often take a simple form where the tenant keeps a fixed fraction of the produce. Studies have shown that the specific fraction (e.g. half or two-thirds) tends to be uniform within regions, but can vary consid-

[1] In this definition a norm is the same thing as a convention, as defined in Young (1993), and we will use the terms interchangeably.

erably across regions (Young and Burke, 2001; Burke and Young, 2003). The competition between alternative standards in technology (Arthur, 1989; David, 1985) is another example. Locally uniform convergence to a standard often arises, together with differences in choices by different groups. Likewise, there is considerable evidence of differences in norms of corruption of government officials across regions with similar governance structures, of courtesy and helpfulness towards strangers, and the industriousness and entrepreneurship of workers. Numerous studies of organizations have indicated how the success of work groups depends on "corporate culture" or "social capital" (e.g. Nahapiet and Ghosal (1998), Adler and Kwon (2002), Hermalin (1999), Knack and Keefer (1997)). Otherwise similar groups are capable of very different levels of performance based on a shared expectation of individual contributions to the group. Inter-organization relationships are another possible source of regional uniformity, as in Krugman(1991), Uzzi (1997), and Owen-Smith and Powell (2004). While uniform behavior within groups has been the subject of numerous inquiries stressing social influences on behavior, relatively little theoretical research in economics (we cite some examples below) has addressed the puzzle of geographical variation: why and when do different groups adopt different norms?[2]

In this paper we examine a variety of regional variation phenomena using a model that incorporates local interaction within networks and social influence on choices. We build upon much previous work, but especially the recent literature on evolutionary game theory (e.g. Young, 1998; Ellison, 1993). Rather than deducing equilibrium behavior from game-theoretic solution concepts predicated on strong rationality assumptions, these models describe the aggregate behavioral patterns that emerge when individuals adopt relatively simple, *boundedly rational* decision rules. Predictions focus on the sta-

[2]While numerous works (e.g. Kremer 1993, Selod and Zenou 2001) have pointed to the presence of multiple equilibria to explain geographic variation, the latent existence of multiple equilibria does not explain the simultaneous selection of different equilibria at different, possibly contiguous, locations.

ble, long-run behavior of the dynamics, which depend not only on the behavioral rules but also on the topology of the social interactions. Within this literature we follow most closely those papers that address the prospects for behavioral uniformity (and diversity) under various social maps and payoff structures, such as Morris (2000), Goyal (1996), and Sugden (1996). The most direct influence is the Young and Burke (2001) model of the evolution of sharecropping norms. Unlike the latter paper, however, the current framework does not rely on a specific parametrization of payoffs, and admits a broader range of interpretations. In this more general setting analytical results are not forthcoming, and we rely on a computational approach.[3] Fortunately, the results are remarkably sharp.

Our model has a network of agents with a defined neighborhood relation.[4] A central assumption is that choices of neighbors exert a direct social influence on an agent's decision. A particular choice yields a greater payoff, and so becomes more likely, if more neighbors have recently made the same choice. Decisions are myopic and error-prone. Agents take the current choices of neighbors as given and, in each period, try to choose an optimal response. However, with small probability they make the wrong choice. As this stochastic dynamic process evolves in time, we are able to witness the emergence of the characteristic pattern of locally uniform and globally diverse choices. As parameters of the model are varied, we also observe the presence of critical points that, when crossed, lead to a sudden qualitative change in the behavior of the system. As a result, norms within a region tend not to change gradually and, instead, respond suddenly as important thresholds are crossed. In the absence of errors the network could

[3] A computational approach to problems such as these is taken by Epstein and Axtell (1996) and Epstein (2001).

[4] The nature of social interactions in this model is admittedly simple, but nonetheless enables sharp and robust results. For example we assume a fixed exogenous neighborhood structure. We discuss possible extensions for future research, such as endogenous network formation, in the conclusion.

get locked into a number of possible norms. Errors, even when they are small, allow us to refine our predictions considerably.

The local uniformity is clearly a consequence of the local social influence (or network externality) assumption. While network externality models can generate behavioral uniformity, they tend to do so on a global scale. We want to get diversity without resorting to the untenable assumption that groups are isolated. Global diversity arises from the assumed heterogeneity in the environment, where heterogeneity occurs within as well as across regions. Global diversity can arise also in models with homogeneous agents, as in Morris (2000), Sugden (1996), and Goyal (1996), among others. The extension to the heterogeneous case is non-trivial, and we are motivated by the richness of the applications that are afforded in such settings. For example, in the Young and Burke (2001) model, regions differ in soil quality. One region may have more fertile soil on average, although all regions have both high and low quality plots. Some contracts are more ideally suited to high quality plots, others to low quality plots. In the presence of local social influence, the contract chosen for a low quality plot is likely to depend upon the average soil quality in the region. In any region the landlord and tenant of the low quality plot will be drawn towards the contract others choose, and in fertile regions this is likely to be the contract appropriate for high quality plots. A similar idea appears in the BFP model of medical procedure choice. The ideal procedure for a patient depends upon individual characteristics, such as age, which vary within the population. Physicians are also influenced by the choices of neighboring physicians. Locally uniform treatment will result, where the local norm depends on the typical patient characteristics in the region. For instance, in regions where older people are relatively more numerous the norm that emerges is for the use of procedures better suited for older patients, *even on younger patients.* In each case, an identical transaction (between identical landlords and tenants on identical plots, or identical physicians and patients) will result in very different outcomes at different locations. The predicted relationship between average regional characteristics and local norms holds up rather well in data sets on sharecropping in Illinois (Young and Burke, 2001) and cardiac care in Florida (BFP).

The environment we present here extends BFP's finding of robust geographical variation and generalizes it to other applications. In BFP, stable long-run variation requires an infinite geography in which physicians are assumed to be located on the integers $\mathbb{Z}$, and their neighbors are the adjacent physicians. Here we obtain regional variation for a finite number of agents and for a larger set of spatial arrangements. Moreover, errors in decision-making are shown to be crucial for producing regional variation in the finite case. When regional variation is a stable outcome of the noiseless process the noisy dynamic process approximates exactly this equilibrium (leading us to believe that regional variation is stable in the sense of Kandori, Mailath and Rob (1993) and Young (1993)). Remarkably, even when regional variation cannot be a long-run outcome we find that it arises as a robust phenomenon for considerable lengths of time.

The rest of the paper is organized as follows. We describe the general model in section 2, and present the results from our simulations in section 3. Finally, we draw together our conclusions in section 4, emphasizing what we learn about the emergence of norms, and discussing the key phenomena observed—regional variation, criticality, and lock-in.

2 Model

There are M agents or decision makers. Each agent x has a set of neighbors, $\mathcal{N}(x)$. We consider two arrangements of the agents (called geographies): (i) a circle, and (ii) a torus. The circle has the virtue of simplicity, while still capturing the notion of local, overlapping interactions. While any individual is influenced only by the actions of two neighboring agents, the behavior of others much farther away may exert an indirect influence. At the same time we are interested in more complex interaction structures because results on the circle may not be robust. The torus model gives each agent 4 rather than 2 neighbors, and affects the behavior at the boundaries between regions, as described below. We find that regional variation may arise and persist in both types of graphs, but the differences between the cases indicate that network size is important. In the circle, we index the M agents by the numbers $1, 2, \ldots, M$ and define the neighborhoods by

$$\mathcal{N}(i) = \begin{cases} \{2, M\} & \text{if } i = 1 \\ \{1, M-1\} & \text{if } i = M \\ \{i-1, i+1\} & \text{if } M > i > 1 \end{cases}$$

In effect, agents are located along a single dimension, with the neighbors being those at adjacent locations. However, we define 1 and M to be neighbors. In the torus geography, we do the same with a two dimensional arrangement. Suppose that each location is identified by two coordinates (i, j). We assume that $1 \leq i \leq K$ and $1 \leq j \leq K$, together with $M = K^2$. Now the neighbors of (i, j) are $\{(i+1, j), (i-1, j), (i, j+1), (i, j-1)\}$, with the obvious modification at the boundaries.

At each date, each agent observes a private signal $\sigma_i \in \{\alpha, \beta\}$ and chooses a decision $d_i \in \{A, B\}$. Payoffs depend upon (d_i, σ_i), as well as on the decisions of neighbors.[5] Let n be the number of neighbors of agent i who chose action A *in the previous period*; we denote payoffs by $\pi(d_i, \sigma_i, n)$. Define the payoff difference $\Delta^\sigma(n) = \pi(A, \sigma, n) - \pi(B, \sigma, n)$. Our key *social influence* or *network externality* assumption is as follows:

Assumption 1: $\Delta^\sigma(n)$ is an increasing function of n for all σ.

In other words, A becomes relatively more attractive if more neighbors decided to play A in the previous period. The previous definitions lead to two possibilities: (i) $\Delta^\sigma(n) \leq 0$

[5] Mathematically, the structure here is an *interacting particle system*. These are discussed by Liggett(1999) and Schinazi (1999). Such systems were introduced fruitfully into game theory by Blume (1993).

or $\Delta^\sigma(n) \geq 0$ for all n, or (ii) $\Delta^\sigma(n)$ changes sign for some value of n. In the former case the choices of neighbors do not affect the ranking of actions, whereas in the latter case they do. We focus on case (ii)—for both signals the relative ranking of the actions changes with n.[6] Of particular interest is how rankings change with the signal. Let N be the total number of neighbors (which is an even number for both the circle and the torus). When $n = N/2$, the neighbors are equally split between playing the two actions, so the effect of the two is, in some sense, neutralized. We assume:

> *Assumption 2:* At $n = N/2$, the sign of $\Delta^\sigma(n)$ changes with the signal—specifically, $\Delta^\alpha(N/2) > 0$ and $\Delta^\beta(N/2) < 0$.

This assumption captures the idea that A is the better choice when signal α is observed, whereas B is the better choice if signal β is observed (once social influences are neutralized). Together with the assumption 1, and the focus on case (ii), this implies that an agent should choose A if all N neighbors do so, and B if all neighbors choose B. In other words, "lock-in" to either action is a possibility. We assume error in decision-making.

> *Assumption 3:* Given (σ, n), with probability $(1 - \varepsilon)$ an agent chooses the action that maximizes payoffs, and with probability ε chooses the inferior action.

Our final assumption relates to properties of signals. A region is a fixed, contiguous, set of locations. Signals arrive with fixed probability within a region, but the probability can differ across regions. To simplify, we consider two regions, East and West. In the circle, East is defined as the set $\{i | i \leq M/2\}$, in the torus by $\{(i, j) | j \leq M/2\}$. An important feature of the definition of regions here is that they are contiguous, and not isolated. There are other ways to model contact between regions. For instance, a model

[6] For case (i), our arguments imply that the emergent behavior would be to play the action superior for all n.

with very similar properties to our circle model is the following—we have *two* circles and, occasionally, an agent in one circle imitates a randomly selected individual from the other circle.[7]

> *Assumption 4:* The probability that a location $x \in$ East receives signal α is p, and the probability location $y \in$ West receives signal α is q. In general $p \neq q$.

We now illustrate the reach of this model with three applications.

We begin with a finite version of the *medical procedure choice* problem discussed in BFP. The decision-makers are physicians who, in each period, get a new patient with a specific condition (say, coronary atherosclerosis). Signals are now to be interpreted as patient characteristics, e.g. age (α is 'old' and β is 'young'). After observing patient characteristics a physician must choose between two procedures (drugs (A) or surgery (B)). Physicians are influenced by the choices made recently by their neighbors—either because they talk to, and learn from, neighbors, or from a desire to conform with local practice. In the manner of assumption 2, procedure A is better for α (old) patients while B is better for β (young) patients. Patient characteristics (age distributions) can differ across regions ($p \neq q$). We want to know whether stable patterns of procedure use evolve, and whether patients in different regions are likely to receive different treatments.

Our second example is of technology choice with a network externality. The decision-makers are problem solvers who belong to one of two professional groups. There are two available technologies (A and B) from which an individual must choose. An agent's neighbors are people he interacts with, typically from the same profession. We index people in such a manner that a region comprises all the people in the same profession (the people with ties across professions are placed on the boundary). For concreteness, the East comprises graphic designers, and West writers. At each date, each individual

[7] We thank Rinaldo Schinazi for this observation.

gets a task which may be intensive in the use of images (α) or of text (β). Image-intensive tasks are best solved using technology A, while technology B is best for text-intensive tasks. There is also a network externality present. You are more likely to use a technology if your neighbors use the same technology. People in both professions get both types of tasks, but graphic designers are more likely to get image-intensive tasks ($p > q$). The questions now concern whether stable patterns of technology adoption arise, and whether technology use differs across professional groups.

Our final example concerns corruption of government employees, and we assume the geography of a circle. The agents are now officials, who choose to either solicit a bribe (A) or not (B). The payoffs imply that officials are more likely to solicit bribes if their neighbors do so (perhaps because this lowers the risk of getting caught, or reduces the stigma associated with corruption). A signal is now the arrival of a private individual with some business transaction. This person has some observable characteristic that can take one of two values, α or β, such as rich or poor, doctor or lawyer, member of one ethnic group or another. A given individual's corruptibility, that is his willingness to accept the bribe rather than report the official, cannot be observed in advance, but it is correlated with the observable characteristic. For example, suppose that 70% of α's are corruptible (accept all bribes), but only 30% of β types are. Define payoffs conditional on corruptibility with the notation $\hat{\pi}(d_i, b_k, n)$, where $b_k \in \{c, \neg c\}$ indicates corruptibility. Let $\hat{\pi}(A, c, n) = n$, $\hat{\pi}(A, \neg c, n) = n - 2$, and let $\hat{\pi}(B, b_k, n) = 0$ for either value of b_k. Given these payoffs and the conditional probabilities of corruptibility, the expected payoff for an official that solicits a bribe from an α type is thus

$$\pi(A, \alpha, n) = 0.7n + 0.3(n - 2) = n - 0.6.$$

Similarly, the expected payoff from attempting to bribe a β type is $\pi(A, \beta, n) = n - 1.4$ Assuming $N = 2$, it can readily be confirmed that assumptions 1 and 2 hold. We assume regions differ in their proportions of observable types, and so the rate of corruptibility will also differ across the regions. Again we want to know whether the emergent patterns of corruption display regional variation, as well as which circumstances lead to a non-corrupt governance norm. Another interpretation along these lines might involve fraudulent activity on the part of the employees of a firm, where such activity requires cooperation between agents, and the signal indicates the expected profitability of the fraud opportunity.

3 Results

We begin by considering the case where the error probability is zero (the "zero noise" case). In the torus model, it is easy to confirm that possible long run outcomes (absorbing states of the Markov chain) include the two uniform states, A at each location or B at each location, the regional variation state, in which A is played in one region and B in the other, as well as the "blinking cycle" depicted in Figure 1.[8] "Blinking" involves a cycle between two states—in each state, all four neighbors of a location are occupied by people making the opposite choice; and each individual alternates between their two choices. There are, in addition, many other possibilities—for instance a combination of regional variation and a strip that is in the blinking pattern (e.g. as in Figure 3).

In Figure 2, we illustrate convergence, starting from a random initial condition, for the following parameter values: $M = 40$, $\varepsilon = 0$, $p = 0.3$, $q = 0.7$. In this instance, we happen to get convergence to the regional variation state. The iteration at which we take a snapshot appears at the top of the frame. Figure 2 isn't the only possibility. Another simulation, starting again from a random initial condition, arrived at iteration

[8] All figures were generated using Matlab. Each location is represented as a diamond, which will be either red or blue depending on the choice. The background is light green. However, when adjacent cells are of the same color, Matlab fills in the background using that same color. The Matlab program is available at www.garyfournier.com.

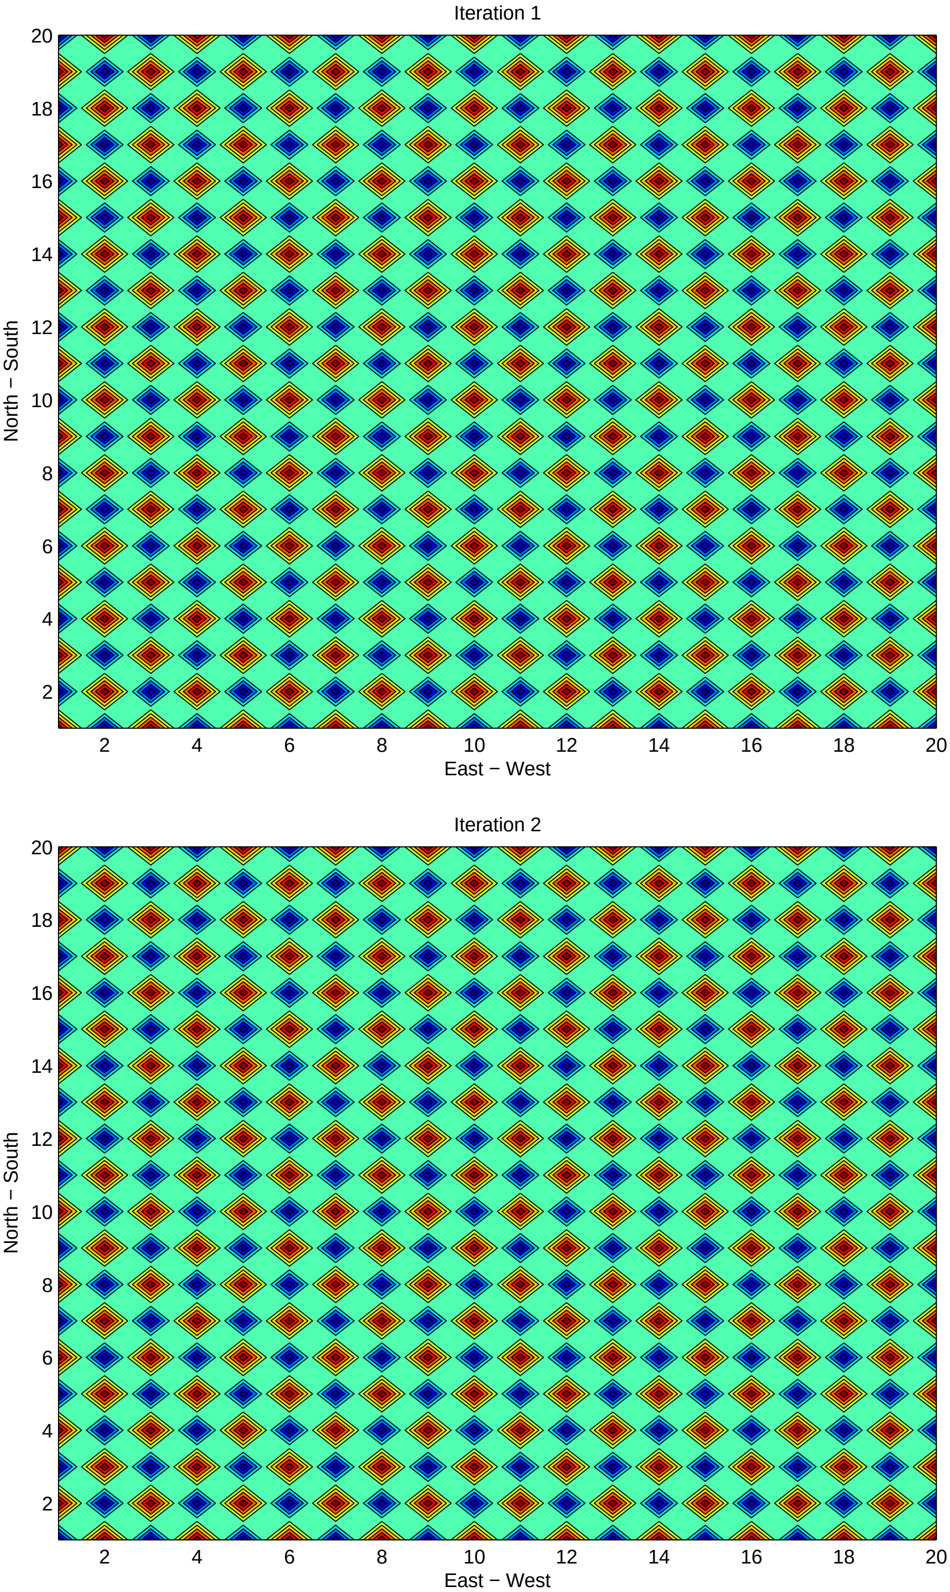

Figure 1: The “blinking” pattern.

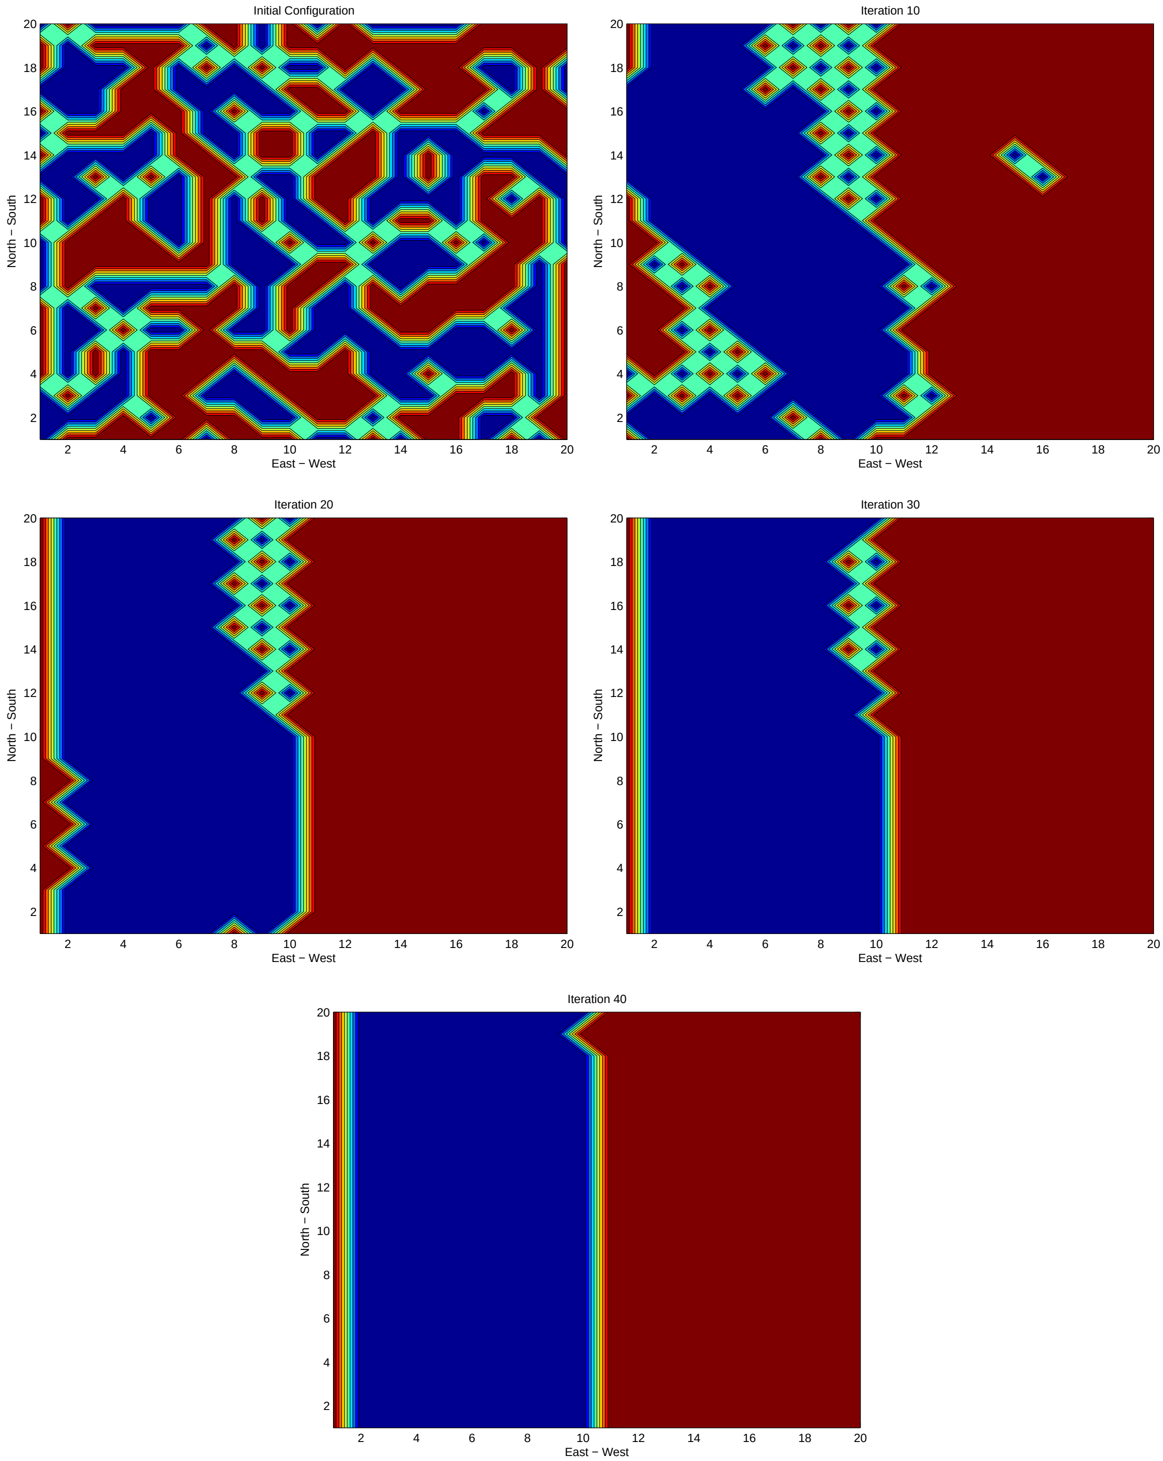

Figure 2: Convergence in the zero noise case from a random initial configuration.

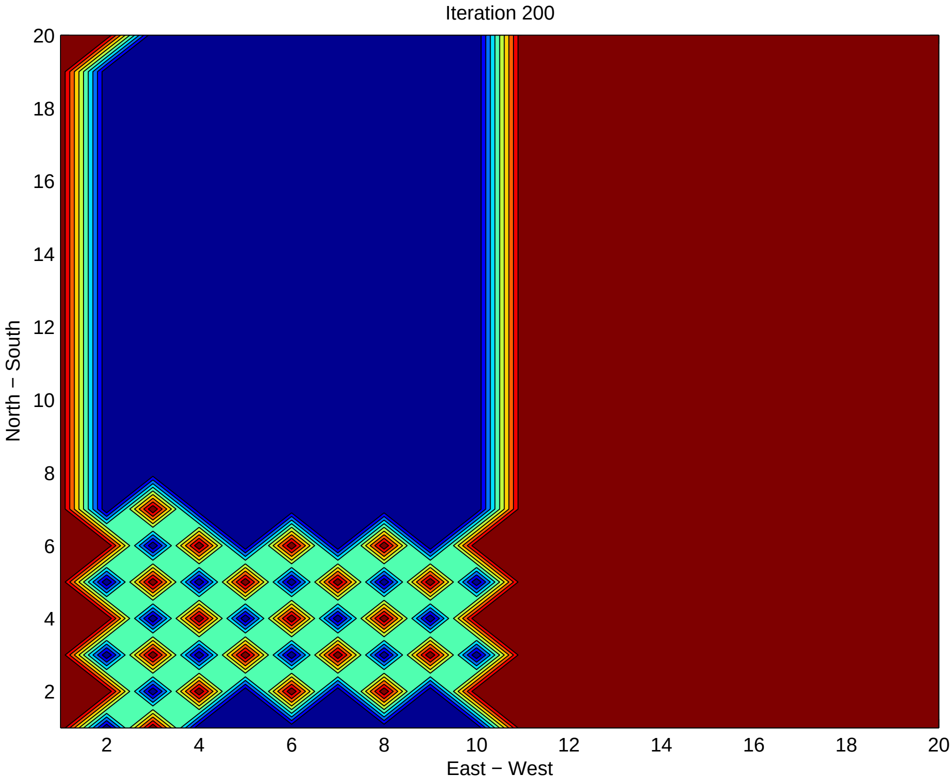

Figure 3: Alternative simulation result for the zero noise case.

100 in the state in Figure 3. In fact, in an experiment with 27 simulations (of 100 iterations), the regional variation state arose 6 times, a uniform state never arose, and the remaining outcomes were mixtures as in Figure 3. Next we introduce noise. Figure 4 demonstrates convergence, starting from a random initial state, to a state of regional variation, letting the error probability be $\varepsilon = 0.05$. Again, the time period at which each snapshot was taken is labelled in the frame. For the same parameter values, in 27 consecutive simulations, the regional variation state emerged *every time.*

In the next simulation we reduce ε to 0.01, and depict the results in Figure 5. Convergence to regional variation still occurs, but it is likely to require more iterations than it would with a higher error probability.

In Figure 6 we show that the other possible steady states of the noiseless process (uni-

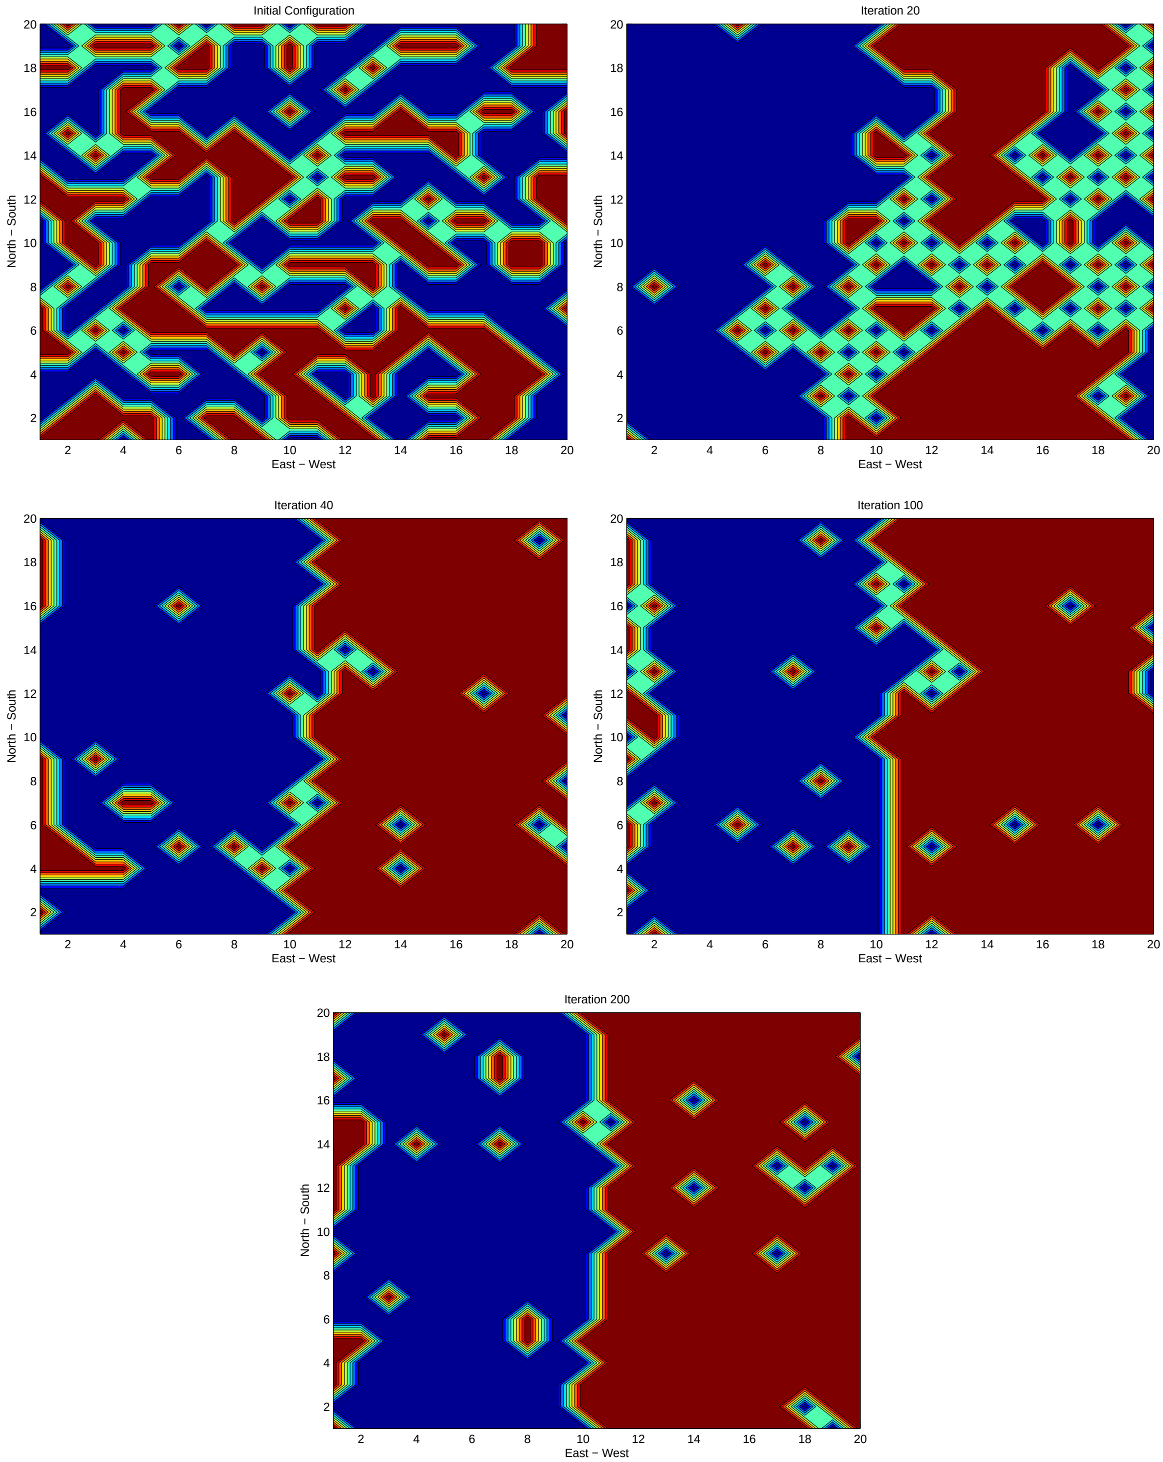

Figure 4: Convergence in the noisy case from a random initial configuration ($\varepsilon = 0.05$).

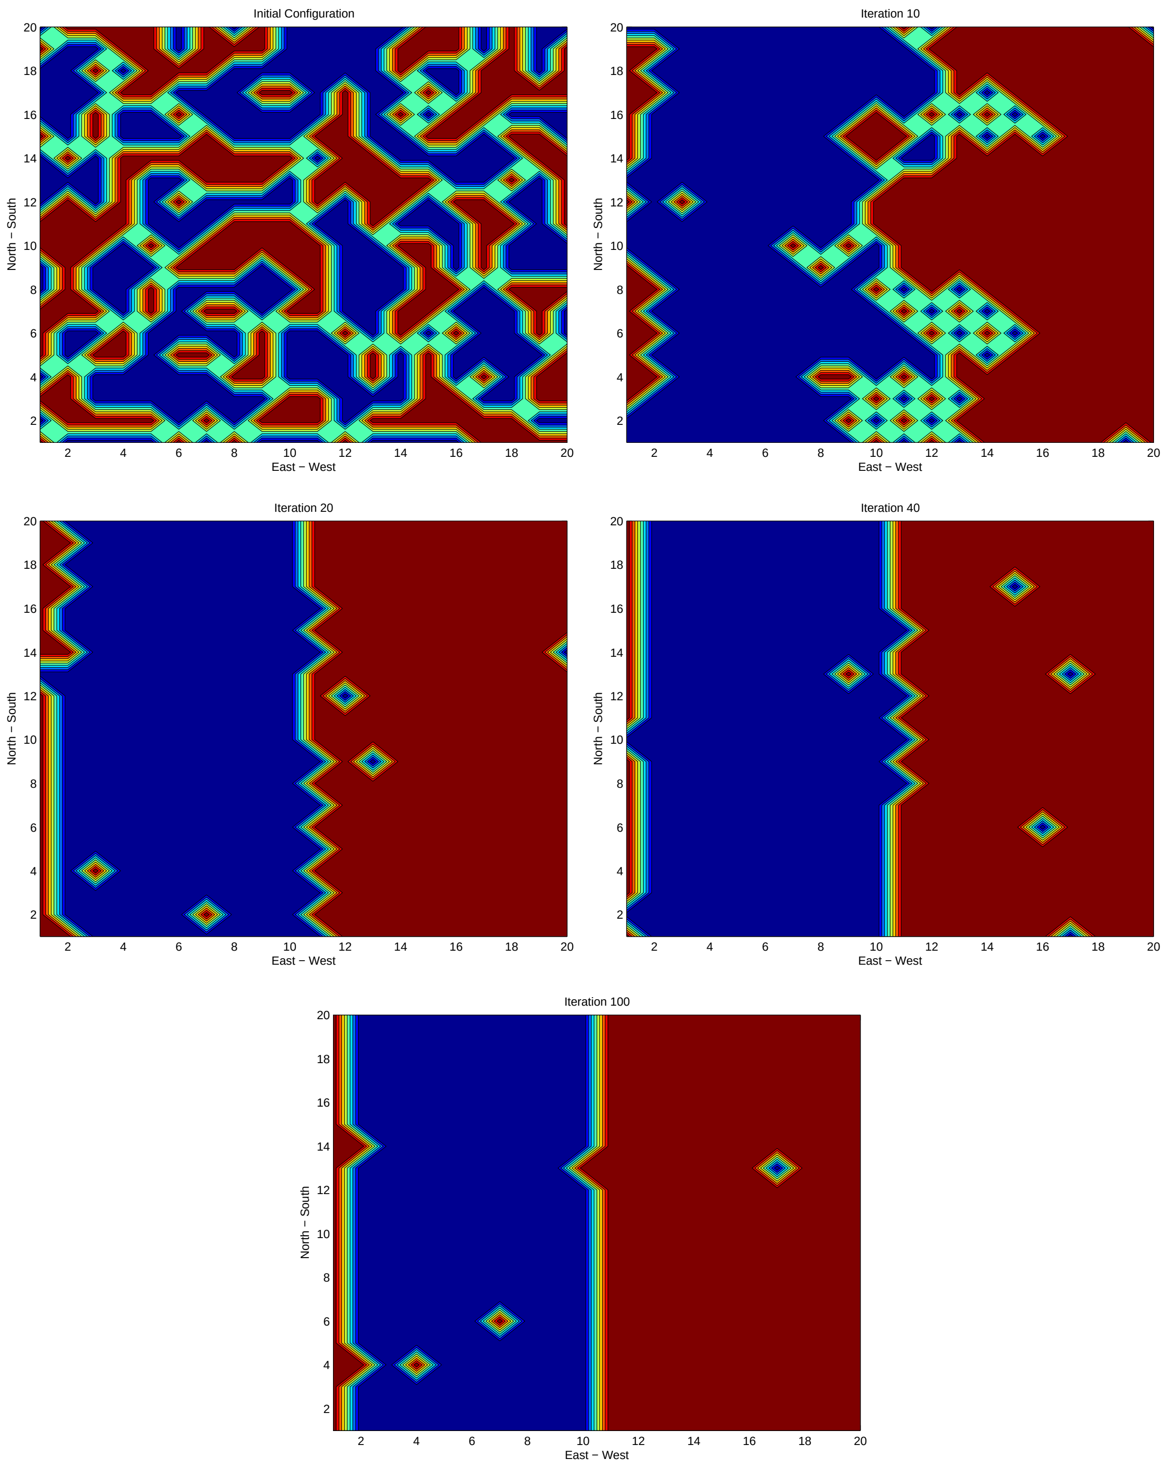

Figure 5: Convergence in the noisy case from a random initial configuration ($\varepsilon = 0.01$).

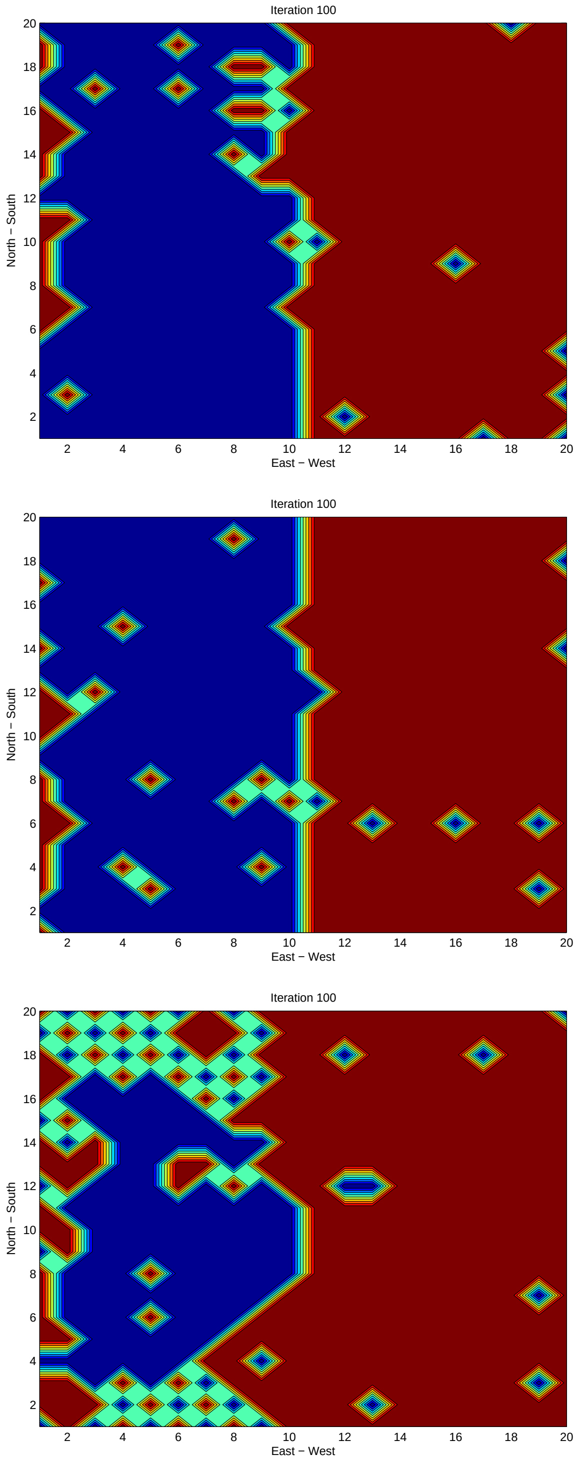

Figure 6: Convergence in the noisy case starting from, respectively top to bottom, the two uniform states and the blinking state. ($\varepsilon = 0.05$).

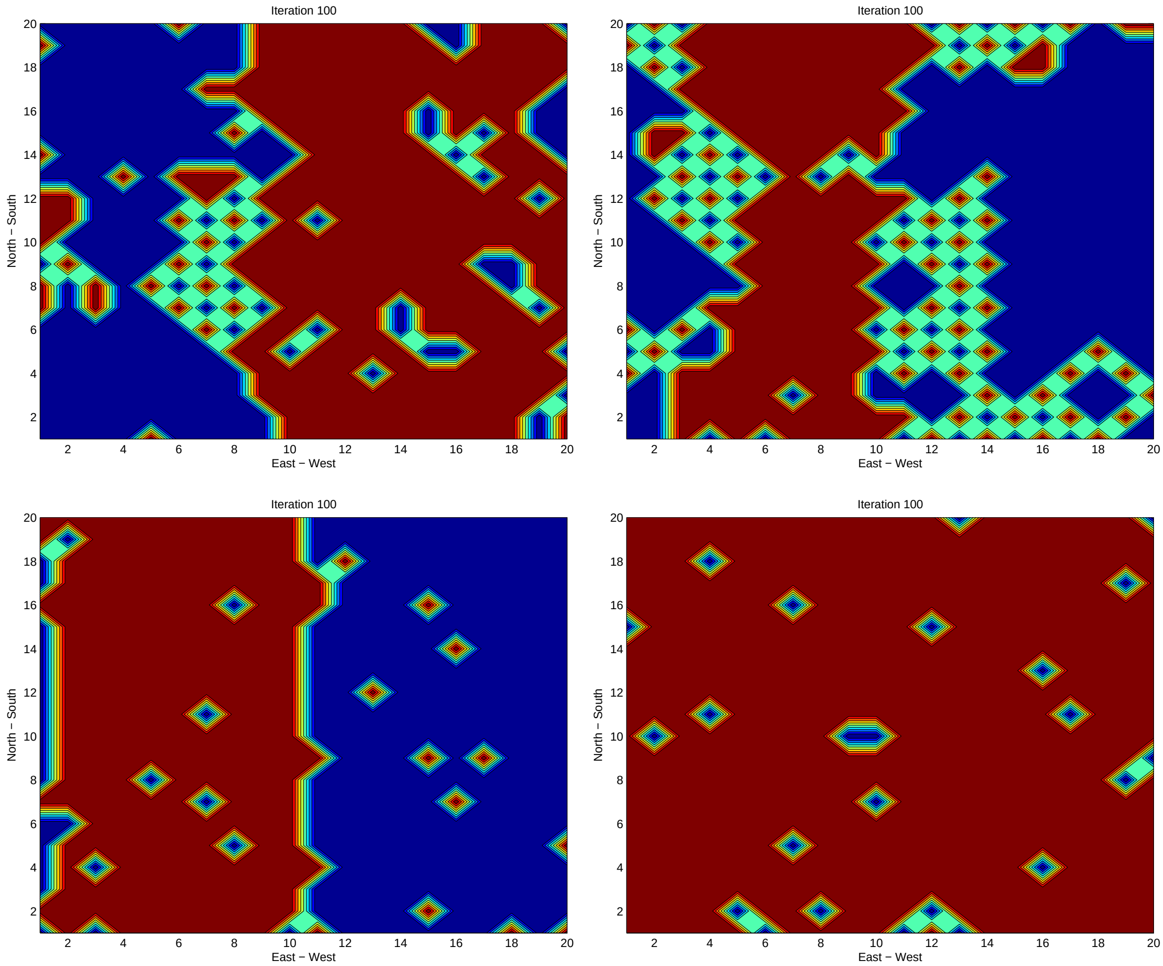

Figure 7: Criticality: $p = 1/2$ and $q = 1/2$ are the thresholds.

form and blinking) are unstable, even when the error probability is very small. The simulations begin, respectively, in (1) the uniform state with everyone playing A, (2) everyone playing B, and (3) one of the blinking states. With error, the stable local norm can always get a toe-hold, and then spreads by contagion. In each case we get fairly quick convergence to the regional variation state. Of the three, the blinking state took longest to disappear.

Figure 7 illustrates the "criticality" phenomenon: the existence of critical parameter thresholds that, when crossed, effect precipitous change in the behavior of the system. However, on either side of the critical point variation in fundamentals is irrelevant. The parameters that exhibit threshold effects in this model are p and q. In the east, the prevailing norm depends upon whether p is greater than, or less than 1/2. We start, in the top left hand frame of Figure 7, with the long-run outcome when $p = 0.55$ and $q = 0.45$. Increases in p above 0.5, or decreases in q below 0.5, make virtually no difference. In the top right hand frame, we reverse the parameter values (to 0.45 and 0.55 respectively). In both these graphs regional differences are not as sharp as they were in Figures 4–6, but the pattern is unmistakeable. In the bottom left hand frame, we have $p = 0.3$ and $q = 0.7$ (that is we reverse the values from the Figure 4). In the final frame, $p = q = 0.6$. It is clear that the critical value determining the norm for a region is the probability of the α signal.

Finally we consider the circle model, highlighting just the points of contrast with the torus. One key difference is that regional variation is no longer a long run steady state of the noiseless process. The reader can readily verify that the variation state will always eventually be supplanted by a globally uniform state, never to return. The probability of getting to a uniform state from the regional variation state, however small, is bounded away from zero. And once either uniform state arises it will persist forever. Thus the long run outcome may be global uniformity, but it could also be the blinking

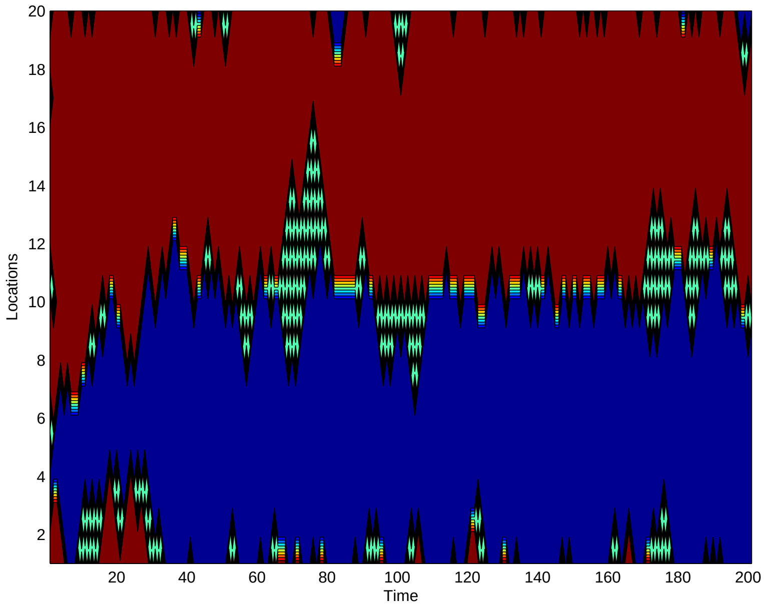

Figure 8: Slow convergence: regional variation is a robust phenomenon ($\varepsilon = 0$).

state depending on early random events. Figure 8 shows something interesting, however. In this model we show the full evolution in one graph, and the X-axis denotes time. We consider first the noiseless case. Despite the fact that regional variation is not a steady state here, we see after 100 iterations that it could arise and persist for multiple time periods. Convergence to one of the steady states is *very slow,* and regional variation is a robust phenomenon. Another robust outcome is some combination of a regional norm with the blinking pattern. The same is true for the torus model.

Finally, we start from a uniform state and introduce small noise ($\varepsilon = 0.05$) to get convergence to regional variation. It should be pointed out that small regions of the blinking pattern can arise and persist for multiple periods. Therefore convergence is not as sharp as in the simulations on the torus. The blinking pattern is, to a large extent, an artifact of the simultaneity of decisions. We have simulated a model in which decisions are asynchronous (at each date, only one person gets to make a choice). In that model the blinking pattern (as well as combinations of regional variations with this pattern)

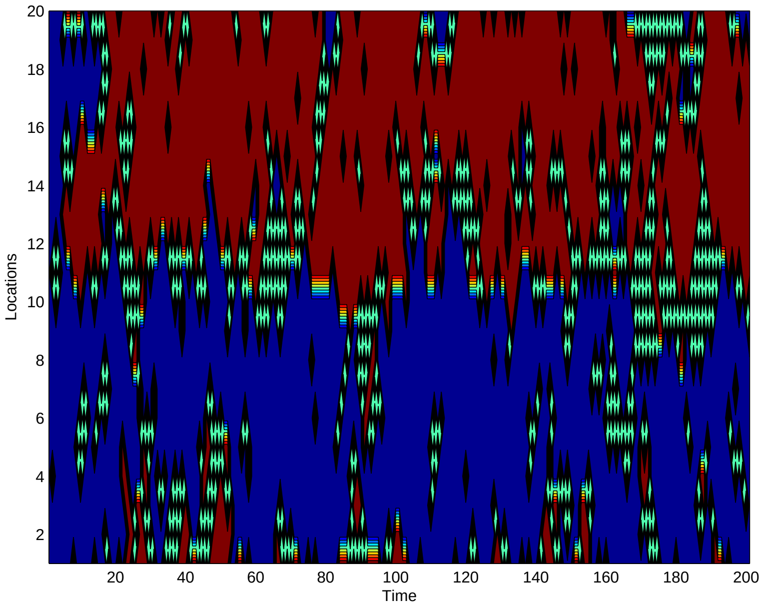

Figure 9: Convergence to long-run outcome ($\varepsilon = 0.05$).

does not arise.

4 Conclusion

The model tells us that alternative norms may coexist in close proximity to each other for indefinite periods of time, despite a tendency toward local conformity. We focus on situations involving a network externality, which tends to promote uniformity among interacting agents. Conformity within a given region will be close to complete, despite heterogeneity in payoffs. The distribution of signals determines what the local norm will be. Consider the medical practice example, where the signal is patient *type* (e.g. old or young). The majority type within a region will dictate the content of the norm, and thus norms can differ across regions depending on the region's type mix. When the dominant type differs across regions, local norms respond accordingly, and we witness global diversity as the stable phenomenon. The concept of "lock-in" has been very influential in the study of institutions and organizations, and our results in the noiseless case illustrate how it could arise. But small noise, as in the theory of Kandori, Mailath, and Rob (1993), Young (1993) and Ellison (1993), allows us to refine predictions considerably. New patterns of behavior, although they arise by error, can spread contagiously until they become locally prevalent. Lock-in does not have quite as tenacious a hold as in the noiseless case.

The locality of interactions is crucial for our results. In such settings, global majorities need not dictate global norms. Alternative norms will survive as long as there exist regions (or subgroups) within which the globally dominant signal or type forms a minority, even though interactions straddle regional boundaries. As informal evidence, we observe that minority languages are sustained by the presence of ethnic residential enclaves. We also observe ethnically and regionally specific slang and dress codes, as well as pockets of dedicated Mac users in a Windows-dominated world. Furthermore, as our criticality results indicate, stable coexistence does not require extreme differences in the composition of the population across regions or groups. The criticality of the 50% threshold also means that norms can shift rapidly within a region with even a small change in demographics around the threshold. The results caution against making inferences about preferences, both within and across regions, based on observed behavior. A single norm can accommodate a diversity of types, just as relatively small changes in group characteristics may cause a discrete shift in the dominant behavior.

The condition of regional variation, involving local uniformity and global diversity, embeds certain social tensions. We have shown that within any region or group there may be a large number of agents (i.e. the minority type) who would be better off (a) living in a different region, in which their preferred norm prevails, or (b) living in the same region but being a member of the majority type. In the signal interpretation, any given individual will face inferior payoffs whenever she receives the locally less common signal type. We have treated location and characteristics as exogenous. However, if geographic location or social network were made endogenous, we would expect self-selection into locations or networks by type, i.e. spontaneous physical or social segregation. Alternatively if locations remain fixed but type could be chosen, people could simply adapt their preferences to their surroundings. Even racial identity might be viewed as endogenous, as in Akerlof and Kranton (2000) and Bodenhorn and Ruebeck (2003).

Depending on the application under consideration, the degree to which location or characteristics are in fact endogenous will vary, as will the welfare implications. While the prospect of spontaneous segregation echoes Schelling (1971), in the context of our model segregation by characteristics, i.e. the signal, may be socially preferred to an outcome in which regions or groups have a mix of types. Luckily, the implications need not be politically unsavory. For example, ignoring transport costs, it is desirable that medical patients be transferred to the treatment location that specializes in the treatment that best suits her type. In the case of corruption or corporate malfeasance, however, endogeneity may have negative consequences. We might expect, for example, that initially honest types, witnessing rampant corruption or receiving frequent invitations to embezzle, might eventually suffer moral decay. If so, replacing dishonest workers with honest workers on a piecemeal basis would be futile. Segregation and assimilation are not inevitable in every application, however. For example, graphics and text processing may be complements within a firm. If so the firm faces a tradeoff between compatibility across workers or tasks and supplying the best tool for each task. In our model the benefits of compatibility produce local uniformity, but an alternative outcome might involve innovations which render the opposing technologies more compatible.

References

[1] Adler, P.S. and S-W Kwon (2002), "Social Capital: Prospects for a New Concept," *Academy of Management Review*, 27: 17-40

[2] Akerlof, G.A. and R.E. Kranton (2000), "Economics and Identity," *Quarterly Journal of Economics*, 115(3):715-753.

[3] Arthur, W.B. (1989), "Competing Technologies, Increasing Returns, and Lock-In by Historical Events," *Economic Journal* 99(4): 116-131.

[4] Bodenhorn, H. and C.S. Ruebeck, "The Economics of Identity and the Endogeneity of Race," NBER Working Paper 9962. URL: http://www.nber.org/papers/w9962.

[5] Burke, M.A., G.M. Fournier, and K. Prasad (2004), "Physician Social Networks and Geographical Variation in Medical Care," Florida State University working paper. URL: http://mailer.fsu.edu/˜ mburke/

[6] Burke, M.A., and H.P. Young (2003), "On the Distributional Consequences of Contractual Norms: The Case of Cropshare Contracts," Florida State University working paper. URL: http://mailer.fsu.edu/˜ mburke/

[7] Blume, L. E. (1993), "The statistical mechanics of strategic interaction," *Games and Economic Behavior,* 4:387-424.

[8] David, P. (1985), "Clio and the Economics of QWERTY," *American Economic Review*, 75(2):332-337.

[9] Epstein, J. M. and R. Axtell (1996), *Growing Artificial Societies: Social Science from the Bottom Up,* MIT Press.

[10] Epstein, J. M. (2001) "Learning to be Thoughtless: Social Norms and Individual Computation," *Computational Economics,* 18, 9-24.

[11] Ellison, G. (1993), "Learning, Local Interaction and Coordination," *Econometrica*, 61(5): 1047-1073.

[12] Glover, J.A. (1938), "The Incidence of Tonsillectomy in School Children," *Proceedings of the Royal Society of Medicine* 31: 95-112.

[13] Goyal, S. (1996), "Interaction Structure and Social Change," *Journal of Institutional and Theoretical Economics*, 152(3):472-94.

[14] Hermalin, B. (1999), "Economics & Corporate Culture," SSRN Working Paper. URL: http://ssrn.com/abstract=162549.

[15] Kandori, M., G. Mailath, and R. Rob (1993), "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* 61: 29-56.

[16] Knack, S. and P. Keefer (1997), "Does Social Capital Have an Economic Payoff? A Cross-Country Investigation," *Quarterly Journal of Economics*, 112(4):1251-88

[17] Kremer, M. (1993), "The O-Ring Theory of Development," *The Quarterly Journal of Economics*, 108(3):551-575.

[18] Krugman, P. (1991), "Increasing Returns and Economic Geography," *Journal of Political Economy,* 99(3): 483-99.

[19] Liggett, Thomas M. (1999), *Stochastic Interacting Particle Systems: Contact, Voter and Exclusion Processes*. Heidelberg: Springer-Verlag.

[20] Morris, S. (2000), "Contagion," *Review of Economic Studies* 67(1): 57-78.

[21] Nahapiet, J. and S. Ghoshal (1998), "Social Capital, Intellectual Capital, and the Organizational Advantage," *Academy of Management Review*, 23(2): 242-266.

[22] Owen-Smith, J. and W.W. Powell, "Knowledge Networks as Channels and Conduits: The Effects of Spillovers in the Boston Biotechnology Community," *Organization Science*, 15(1): 5-21.

[23] Schelling, T.C. (1971), "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, 1:143-186.

[24] Schinazi, R. (1999), *Classical and Spatial Stochastic Processes.* Birkhauser.

[25] Selod, H., and Y. Zenou (2001), "Social Interactions, Ethnic Minorities and Urban Unemployment," *Annales d'Economie et de Statistique*, 63-64:183-214.

[26] Sugden, R. (1996), "The Coexistence of Conventions," *Journal of Economic Behavior and Organization*, 28(2): 241-256.

[27] Uzzi, B. (1997), "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness," *Administrative Science Quarterly*, 42: 35-67.

[28] Wennberg, J., and M. Cooper, eds. (1999), *The Dartmouth Atlas of Health Care 1999*, Chicago: Dartmouth Medical School and American Hospital Association.

[29] Wennberg, J., and A. Gittelsohn (1973), "Small Area Variations in Health Care Delivery," *Science* 182: 1102-1108.

[30] Young, H. P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton, NJ: Princeton University Press.

[31] Young, H.P. (1993), "The Evolution of Conventions," *Econometrica* 61: 57-84.

[32] Young, H.P., and M.A. Burke (2001), "Competition and Custom in Economic Contracts: a Case Study of Illinois Agriculture," *American Economic Review* 91(3): 559-573.