Learning To Be Thoughtless:
Social Norms And Individual Computation

Joshua M. Epstein

Center on Social and Economic Dynamics
Working Paper No. 6
September 1999

ABSTRACT

This paper extends the literature on the evolution of norms with an agent-based model capturing a phenomenon that has been essentially ignored, namely that individual thought--or computing--is often inversely related to the strength of a social norm. In this model, agents learn how to behave (what norm to adopt), but--under a strategy I term Best Reply to Adaptive Sample Evidence--they also learn how much to think about how to behave. How much they're thinking affects how they behave, which--given how others behave--affects how much they think. In short, there is feedback between the social (inter-agent) and internal (intra-agent) dynamics. In addition, we generate the stylized facts regarding the spatio-temporal evolution of norms: local conformity, global diversity, and punctuated equilibria.

Joshua M. Epstein is a Senior Fellow in Economic Studies at The Brookings Institution and a member of the External Faculty of The Santa Fe Institute. For valuable discussions he thanks Peyton Young, Robert Axtell, John Steinbruner, Clifford Gaddy, and Miles Parker. For production assistance, he thanks Trisha Brandon.

Two Features of Norms

When I'd had my coffee this morning and went upstairs to get dressed for work, I never considered being a nudist for the day. When I got in my car to drive to work, it never crossed my mind to drive on the left. And when I joined my colleagues at lunch, I did not consider eating my salad bare-handed; without a thought, I used a fork.

The point here is that many social conventions have *two* features of interest. First, they are self-enforcing behavioral regularities (Axelrod 1984 and 1986, Young 1995, Lewis, 1969). But second, once entrenched, we conform *without thinking about it*. Indeed, this is one reason why social norms are useful; they obviate the need for a lot of individual computing. After all, if we had to go out and sample people on the street to see if nudism or dress were the norm, and then had to sample other drivers to see if left or right were the norm, and so on, we'd spent most of the day figuring out how to operate, and we wouldn't get much accomplished. Thoughtless conformity, while useful in such contexts, is frightening in others--as when norms of discrimination become entrenched. It seems to me that the literature on the evolution of norms and conventions has focused almost exclusively on the first feature of norms--that they are self-enforcing behavioral regularities, often represented elegantly as equilibria of n-person coordination games possessing multiple pure-strategy Nash equilibria (Young 1993 and 1995, Kandori, Mailith, and Rob 1991).

Goals

My aim here is to extend this literature with a simple agent-based model capturing the second feature noted above, that *individual thought--or computing--is inversely related to the strength of a social norm.* In this model, then, agents learn how to behave (what norm to adopt), but they also learn *how much to think about* how to behave. How much they're thinking affects how they behave, which--given how others behave--affects how much they think. In short, there is *feedback* between the social (inter-agent) and internal (intra-agent) dynamics. In addition, we are looking for the stylized facts regarding the spatio-temporal evolution of norms: local conformity, global diversity, and punctuated equilibria (Young, 1998).

An Agent-Based Model

This model posits a ring of interacting agents. Each agent occupies a fixed position on the ring and is an object characterized by two attributes. One attribute is the agent's "norm," which in this model is binary. We may think of these as "drive on the right (R) vs. drive on the left (L)." Initially, agents are assigned norms. Then, of course, agents update their norms based on observation of agents within some sampling radius. This radius is the second attribute and is heterogeneous across agents. An agent with a sampling radius of 5 takes data on the five agents to his left and the five agents to his right. Agents update, or "adapt," their sampling radii incrementally according to the following simple rule:

Radius Update Rule

Imagine being an agent with current sampling radius of r. First, survey all r agents to the left and all r agents to the right. Some have L (drive on the left) as their norm and some have R (drive on the right). Compute the relative frequency of Rs at radius r; call the result F(r). Now, make the same computation for radius r+1. If F(r+1) *does not* equal F(r), then increase your search radius to r+1[1]. Otherwise, compute F(r-1). If F(r-1) *does* equal F(r), then reduce your search radius to r-1. If neither condition obtains (i.e., if $F(r+1)=F(r)\neq F(r-1)$), leave your search radius unchanged at r.

Agents are "lazy statisticians," if you will. If they're getting a different result at a higher radius ($F(r+1)\neq F(r)$), they increase the radius--since, as statisticians, they know larger samples to be more reliable than smaller ones. But they are also lazy. Hence, if there's no difference at the higher radius, they check a lower one. If there's no difference between that and their current radius ($F(r-1)=F(r)$), they reduce. This is the agent's *radius update rule.* Having updated her radius, the agent then executes the Norm Update Rule.

Norm Update Rule

This is extremely simple: *match the majority within your radius*. If, at the updated radius, Ls outnumber Rs, then adopt the L norm. In summary, the rule is: *When in Rome, do as the (majority of) Romans do, with the (adaptive) radius determining the "city limits."* This rule is equivalent to Best Reply to sample evidence with a symmetric payoff matrix such as:

	L	R
L	(1,1)	(0,0)
R	(0,0)	(1,1)

Following Young (1995), we imagine a coachman's decision to drive on the left or the right. "Among the encounters he knows about, suppose that more than half the carriages attempted to take the right side of the road. Our coachman then predicts that, when he next meets a carriage on the road, the probability is better than 50-50 that it will go right.

[1] When we say "not equal", we mean the difference lies outside some tolerance, T. That is, $|F(r+1)-F(r)|>T$ for inequality, and $|F(r+1)-F(r)|\leq T$ for equality. For our runs, T=0.05.

Given this expectation, it is best for him to go right also (assuming that the payoffs are symmetric between left and right)." The coachman "calculates the observed frequency distribution of left and right, and uses this to predict the probability that the next carriage he meets will go left or right. He then chooses a best reply," which Young terms "best reply to recent sample evidence." Best reply maximizes the expected utility (sum of payoffs) in playing the agent's sample population[2].

The departure introduced here is that each individual's sample size is itself *adaptive*[3]. In particular, as suggested earlier, once a norm of driving on the left is established (firmly entrenched) real coachmen don't calculate anything--they (thoughtlessly and efficiently) drive on the left. So, we want a model in which "thinking"--individual computing--declines as a norm gains force, and effectively stops once the norm is entrenched. Of course, we want our coachmen to start worrying again if suddenly the norm begins to break down. Of the many adaptive individual rules one might posit, we will explore the radius update rule set forth above.

Overall, the individual's combined (norm and search radius) updating procedure might appropriately be dubbed *Best Reply to Adaptive Sample Evidence.*

Noise

Finally, there is generally some probability that an agent will adopt a random norm, a random L or R. We think of this as a "noise" level in society.

Graphics

With this set-up, there are two things to keep track of: the evolution of social norm patterns on the agent ring, and the evolution of individual search radii. In the runs shown below, there are 190 agents. They are drawn at random and updated asynchronously. Clearly, each agent's probability of being drawn k-times per cycle (190 draws with replacement) has the Binomial distribution B(n,1/n), with n=190. Agents who are not drawn keep their previous norm. After 190 draws--one cycle--the new ring is redrawn below the old one (as a horizontal series of small contiguous blue and yellow dots), so time is progressing down the page. There are two Panels. The left Panel

---

[2] For arbitrary payoff matrices, Best Reply is *not* equivalent to the following rule: Play the strategy that is optimal against the most likely type of opponent (i.e., the strategy type most likely to be drawn in a single random draw from your sample). For our particular set-up, these are both equivalent to our "match the majority" update rule. These three rules part company if payoffs are not symmetric.

[3] In best reply models, the sample size is fixed for each agent, and is equal across agents. See Young (1996).

shows the evolution of norms, with L-agents colored blue and R-agents colored yellow. Each entire Panel displays 275 cycles (each cycle, again, being a sequence of 190 random calls.) The right window shows the evolution of search radii, with "hotter" colors for higher radii: Yellow if r>=8; Red if r=6 or 7; Green if r=4 or 5; Blue if r=2 or 3, and Black if r=1.

<u>Runs of the Model.</u>

We present six runs of this model[4]. Once more, we are looking for the stylized facts regarding the evolution of norms: Local conformity, global diversity, and punctuated equilibria (Young, 1998). But we wish also to reflect the rise and fall of individual computing as social norms dissolve and become locked in.

[4] A single random seed is used in all runs.

Run 1. Monolithic Social Norm, Individual Computing Dies Out

For this first run, we set all agents to the L norm (coloring them blue) initially and set noise to zero. We give each agent a random initial search radius between 1 and 50 (artificially high to show the strength of the result in the monolithic case). There is no noise in the decision-making. The uppermost line (the initial population state) of the right graph (190 agents across) is multicolored, reflecting the random initial radii. Let us now apply the radial update rule to an arbitrary agent with radius r. First look out further. We find that F(r+1)=F(r), since all agents are blue. Hence, try a smaller radius. Since F(r-1)=F(r), the agent reduces r to r-1. Now, apply the norm update rule. At this new radius, match the majority. Clearly, this is Blue, so stay Blue. This is the same logic for all agents. Hence, on the left panel of figure 1, the Blue social norm remains entrenched, and, as shown in the right panel, individual "thinking" dies out--radii all shrink to the minimum of 1 (colored black).

Figure 1. Monolithic Norms Induce Radial Contraction

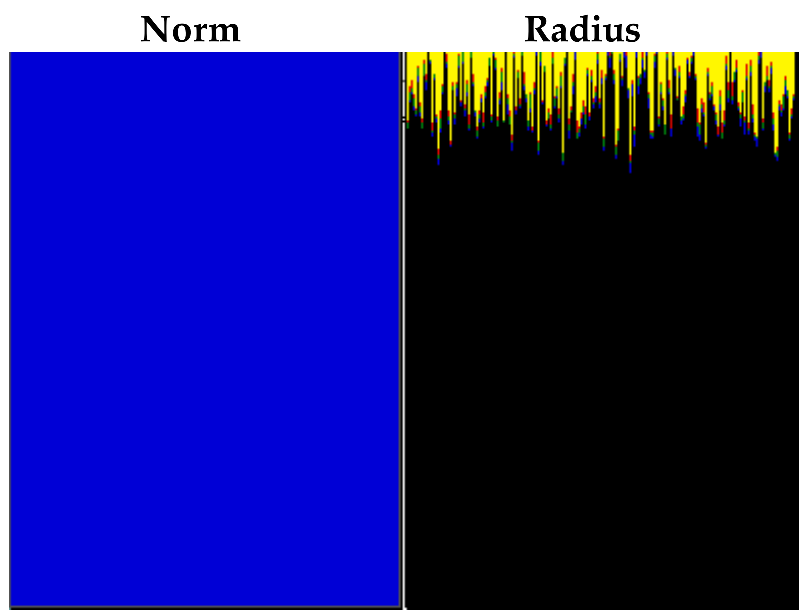

Run 2. Random Initial Norms, Individual Computing At Norm Boundaries

With noise still at zero, we now alter the initial conditions slightly. Rather than set all agents in the L norm initially, we give them random norms. In figure 2, we see the results.

Figure 2. Local Conformity, Global Diversity and Thought at Boundaries

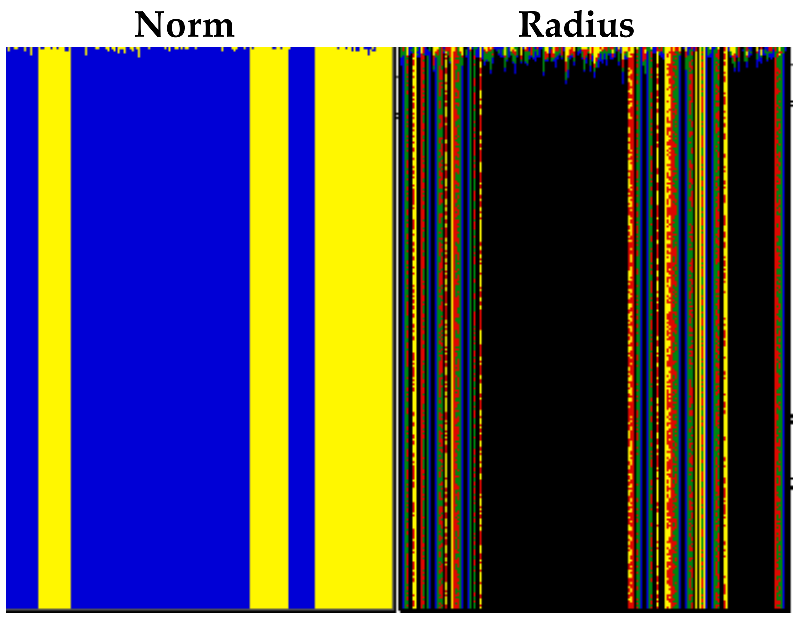

In the left panel, there is rapid lock-in to a global pattern of alternating local norms on the ring. In the right panel, we see that deep in each local norm, agents are colored black: there is no individual computing, no "thinking," as it were. By contrast, agents at the boundary of two norms must worry about how to behave, and so are bright colored.

Run 3. Complacency in New Norms

In the 1960's, people smoked in airplanes, restaurants, and workplaces, and no one gave it much thought. Today, it is equally entrenched that smoking is prohibited in these circumstances. The same point applies to other social norms (e.g., revolutions in styles of dress) and to far more momentous political ones (e.g., voting rights, segregation of water fountains, lunch counters, and seats on the bus). After the "revolution" entirely new norms prevail, but once entrenched, people become inured to them; they are observed every bit as thoughtlessly (in our sense) as before. I often feel that the same point applies to popular beliefs about the physical world; these represent a procession of conventions rather than any real advance in the average person's grasp of science. For example, if you had asked the average 14th Century European if the earth were round or flat, he'd have said "flat." If, today, you ask the average American the same question, you'll certainly get a different response: "round." But I doubt that the typical American could furnish more compelling reasons for his correct belief than our 14th Century counterpart could have provided for his erroneous one. Indeed, on this test, the "modern" person will likely fare worse: at least the 14th Century "norm" accorded with intuition. Maybe we're going backward! In any event, there was no "thinking" in the old norm, and there's little or no thinking in the new one. Again, the point is that after the "revolution," new conventions prevail, but once entrenched, they are conformed to as thoughtlessly as their predecessors. Does our simple model capture that basic phenomenon?

In Run 3, we begin as before, with randomly distributed initial strategies and zero noise. We let the system "equilibrate," locking into neighborhood norms (as before, these appear as vertical stripes over time). Then, at t=130, we shock the system, boosting the level of noise to 1.0, and holding it there for ten periods. Then we turn the noise off and watch the system re-equilibrate. Figure 3 chronicles the experiment.

Figure 3. Re-Equilibration After Shock

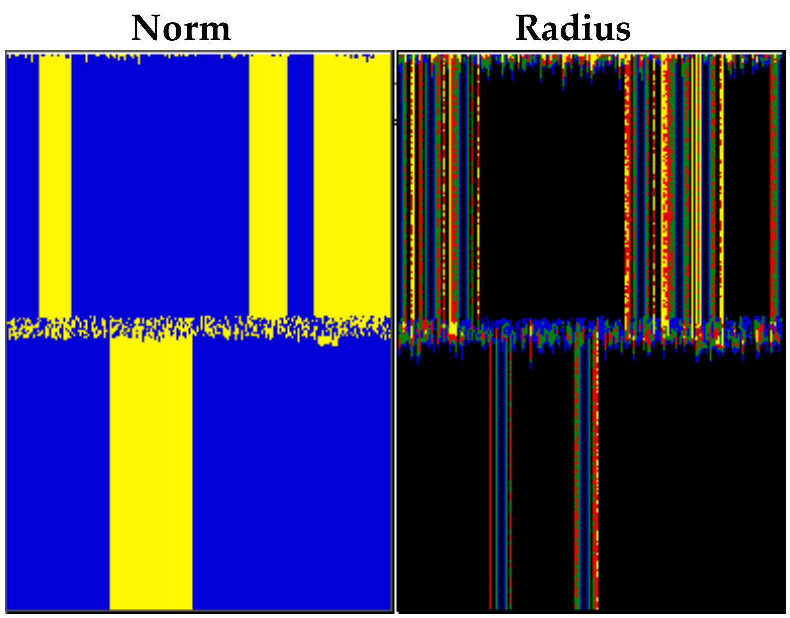

After the shock, an entirely new pattern of norms is evident on the left-hand page. But, looking at the right-hand radius page, we see that agents who were thoughtlessly in the L norm (Yellow) before the shock are *equally thoughtlessly* in the R norm (Blue) after, and *vice versa*.

Run 4. Modest Noise Level and Endogenous Neighborhood Norms

Now, noise levels of zero and one are not especially plausible. What norm patterns, if any, emerge endogenously when initially random agents play our game, but with a modest level of noise (probability of adopting a random norm)? This run and the next use the same initial conditions as Run 2, but add increasing levels of noise. With noise set at 0.15, we obtain runs of the sort recorded in Figure 4.

Figure 4. Moderate Noise and Endogenous Norms

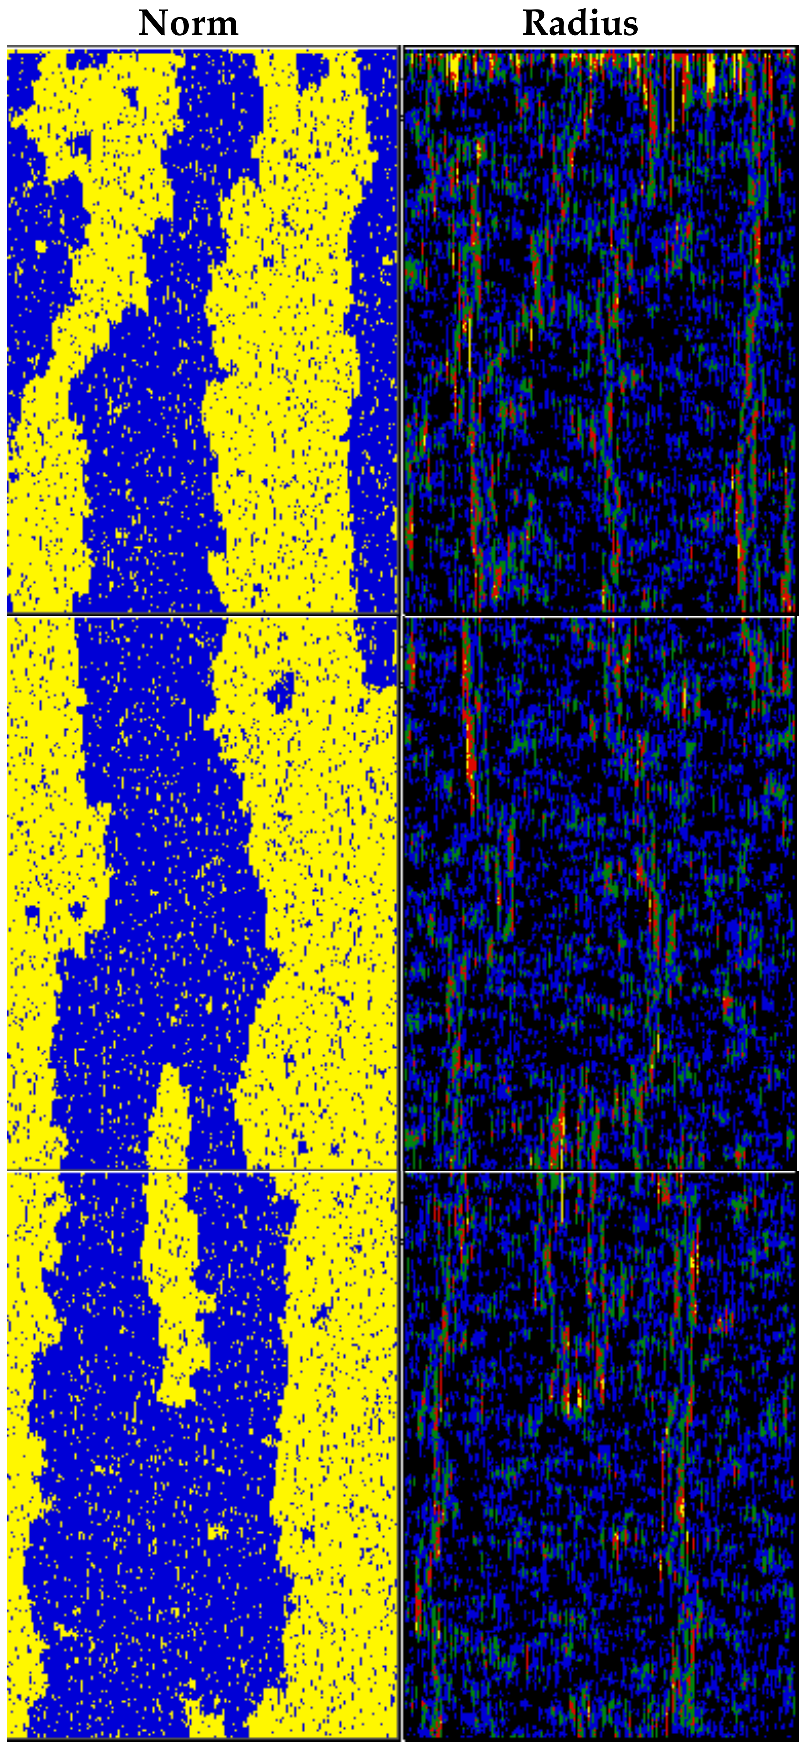

Again, we see that individual computing is most intense at the norm borders--regions outlining the norms. We also see the emergence and disappearance of norms, the most prominent of which is the yellow island that comes into being at around t=500, and then disappears after some 150 periods. One can think of islands as *punctuated equilibria*. We increase the noise level to 0.30 in Run 5.

Run 5. Higher Noise and Endogenous Neighborhood Norms

The result is a more elaborate spatial patterning than in the previous run. Again, however, we see regions of local conformity amidst a globally diverse pattern.

Figure 5. Higher Noise and Endogenous Norms

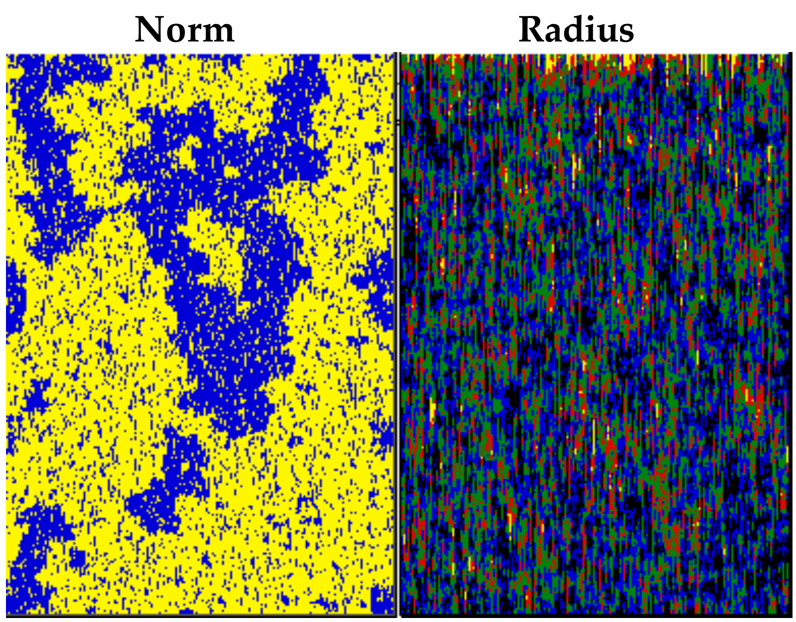

In this run, we see the emergence of a yellow island and later a blue one, punctuated equilibria once more.

Run 6. Maximum Noise Does Not Induce Maximum Search

Finally, we fix the noise level at its maximum value of 1.0, meaning that agents are adopting the Left and Right convention totally at random. One might assume that, in this world of maximum randomness, agents would continue to expand their search radii until the entire right graph of radii was yellow. But this is not what happens, as evident in figure 6.

Figure 6. Maximum Randomness in Norm Does Not Induce Maximum Search

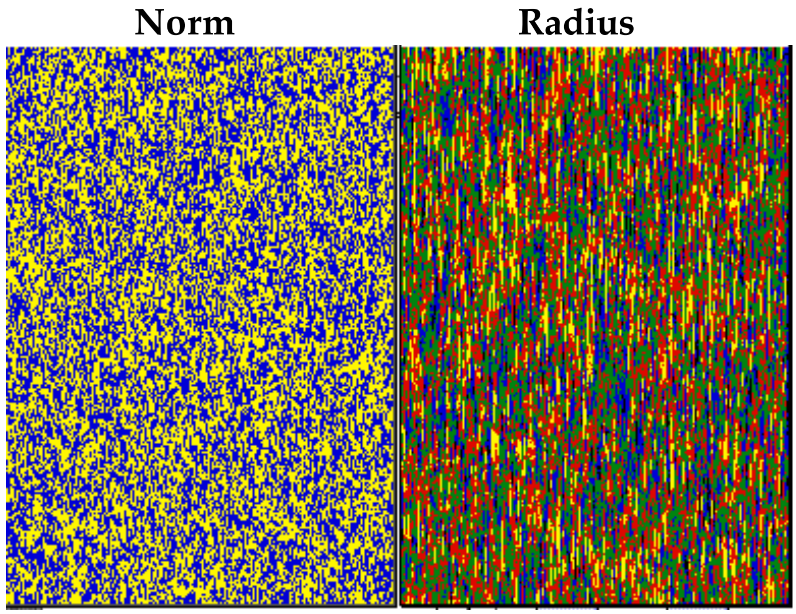

Thinking--individual computing--is minimized in the monolithic world of Run 1. But, it is not maximized in the totally random world of this run. A natural concluding question, then, is at what level of randomness--or other measure of macroscopic complexity--is it maximized?

Summary

My aim has been to extend the literature on the evolution of social norms with a simple agent-based model that generates the stylized facts regarding the evolution of norms--local conformity, global diversity, and punctuated equilibria--while capturing a feature of norms that has been essentially ignored: that individual computing is often inversely related to the strength of a social norm. Obviously, many refinements, sensitivity analyses, and further extensions are possible. But the present exposition meets these immediate and limited objectives.

References

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

______. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

Bicchieri, Christina, Richard Jeffrey, and Brian Skyrms (eds.). 1997. *The Dynamics of Norms*. New York: Cambridge University Press.

Canning, David. 1994. "Learning and Social Equilibrium in Large Populations." In *Learning and Rationality in Economics*, edited by Alan Kirman and Mark Salmon. Oxford: Blackwell.

Feldman, David P, and James P. Crutchfield. 1997. "Measures of Statistical Complexity: Why?". Santa Fe Institute Working Paper 97-07-064.

Gell-Mann, M., and S. Lloyd. 1996. *Complexity* 2(1):44—52

Glaeser, Edward L., Bruce Sacerdote, and José Scheinkman. 1996. "Crime and Social Interactions." *Quarterly Journal of Economics* 111:507—48.

Kandori, Michihiro, George Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long-Run Equilibria in Games." *Econometrica* 61:29—56.

Kirman, Alan and Mark Salmon (eds.). 1994. *Learning and Rationality in Economics*. Oxford: Blackwell.

Lewis, David. 1969. Convention: A Philosophical Study. Cambridge, MA: Harvard University Press.

Ullman-Margalit, Edna. 1977. *The Emergence of Norms*. Oxford: Oxford University Press.

Vega-Redondo, Fernando. 1996. *Evolution, Games, and Economic Behaviour*. New York: Oxford University Press.

Young, H. Peyton. 1993a. "The Evolution of Conventions." *Econometrica* 61:57—84.

______. 1993b. "An Evolutionary Model of Bargaining." *Journal of Economic Theory* 59:145—68.

______. 1995. "The Economics of Convention." *Journal of Economic Perspectives* 10:105—22.

______. 1998. *Individual Strategy and Social Structure*. Princeton: Princeton University Press.