Epidemiology or Marketing?
The Paradigm-Busting Use of Complexity and Ethnography

[*] M.H. Agar, Friends Social Research Center, Baltimore MD
S. Guerin, Redfish Group, Santa Fe NM
R. Holmes, Redfish Group, Santa Fe NM
D. Kunkle, Redfish Group, Santa Fe NM

ABSTRACT

Based on ethnographic work with youth in Baltimore County, an agent-based model was developed to test the finding that circulation of narratives explained the rise and fall of heroin epidemic incidence curves. This paper features the use of Design of Experiments to evaluate the parameters in that model. Results of the analysis suggest the drug epidemics can be better understood as diffusion of a commodity rather than as infection by a disease, the view of the medically dominated substance abuse field. Policy implications of this change in views are sketched in the conclusion.

Keywords: Design of Experiment, Substance Abuse, Ethnography, Agent-Based Model

INTRODUCTION

As we listened to interviewees telling us about illicit drug epidemics in which they had participated, we noticed that they often offered a "folk-explanation" of how and why a particular illicit drug took off. Those folk-explanations, however partial, typically resembled stories of diffusion of consumer products more than infectious disease. People described how early experimenters generated stories about experiences they had with a new drug and, if stories were positive, they would circulate through social networks and encourage further experimentation as time went on, But, with truly dangerous drugs, the effects of continual use would eventually become apparent, so negative stories would increase and experimentation would diminish.

Simple as it sounds, this shift from "epidemiology" to "consumer diffusion" is a major and fundamental change in how illicit drug use is viewed, a change with implications for drug policy and intervention. It is a change in paradigm, in the classic Kuhnian sense of the term.

In this article, we describe an illicit drug case to support this paradigm shift—what we tongue-in-cheek call "paradigm-busting"—but we also use the case to exemplify a more general argument. That argument lays out a research strategy for "paradigm-busting," for setting out to change a framework for viewing a human health problem by

[*] *Corresponding Author Address:* Michael Agar, c/o Redfish Group, 624 Agua Fria St., Santa Fe NM 87501; email: magar@anth.umd.edu. Support by NIDA grant DA-10736 is gratefully acknowledged.

first exploring the phenomenological experience of those close to it. In a second paradigm-busting step, one takes the alternative framework for viewing the health problem that ethnography with those nearest to it always generates, and then one explores it using recent computer modeling techniques from complexity theory, specifically techniques of agent-based modeling. In a third step, one returns to the world of the health problem with fresh eyes, looking at things through a different paradigm. If the exercise has no real world consequences for action, if it doesn't clarify actual cases in new ways, then the exercise was pointless.

Ethnography

Now to the first part of our paradigm-busting strategy: "Ethnographic" research is often thought of as a matter of "data collection" to learn about and document some social world. Most people, including many of its practitioners, ignore its "paradigm-busting" function. But here the "What if" question that can bust a paradigm comes from conceptual systems and social practices learned in the field rather than from the inspiration of genius. A "local" way of thinking and acting suggests a "What if" alternative to an established paradigm for describing and explaining a particular group.

Ethnography, by investigating local ways of making sense of things, can show that an official outsider expert framework and local concepts are paradigms apart. Ethnographic results in fact usually offer candidate "What if…" paradigm-busting questions.

Agent-Based Models

What in the world does this have to do with agent-based models (ABM)? Such models derive from Complex Systems research, that field that busies itself with the study of nonlinear dynamic systems, systems with multiple interactions moving through time that can produce surprising results. In the human realm, ABMs allow us to model emergent results of social dynamics, if we can strip these dynamics down to a few features whose interaction we believe to be critical based on our ethnographic work.

What links paradigms, ethnography and agent-based models is the paradigm-busting question mentioned earlier—"What if…." Axelrod (1997), for instance, describes ABMs as a cognitive laboratory, a way to try out ideas, the computational version of the "thought experiment." And Epstein and Axtell (1996), in their pioneering book on artificial societies, note that their approach is neither "deductive" nor "inductive," but rather what they call "generative." Both of these foundational views support use of the "What if…" question.

Given a particular phenomenon, the question arises, "What if we modeled it a different way?" In fact, that is exactly the exercise we will present in this article. In the case to be presented, a consumer product model generates illicit drug epidemics as well or better than a biomedical model, so the question "What if drug epidemics are more like marketing than disease" gains more credibility. It also gains some rather interesting applications, but more on that in the conclusion.

The Drug Field

Now let us shift to the "drug field," that collection of researchers, clinicians, law enforcement and policy makers—and users—who focus on the use of illicit drugs. Within that field, two competing paradigms have coexisted since the early 20th century—law enforcement and the "medical model." The competition between the two goes back to Supreme Court interpretations of the Harrison Narcotics Act of 1914. Subsequent court disputes centered on whether or not a physician could legally "treat" an opiate addict. The conflict between "legal" and "medical" remains a centerpiece of drug policy today. In this article, though, we focus only on the medical paradigm, the one that prevention and treatment rely on, as well as research that explains drug use by appeal to social and psychological deficiencies.

The medical paradigm subsumes drug use under the category of "disease," as terms like "prevention" and "treatment" already imply. Physicians provide the relevant expertise and biomedical research provides the appropriate way to develop knowledge. Root causes are, in the end, a matter of universal human biology. Diffusion of a drug among its users is a matter for epidemiology, with the infected transmitting the disease to susceptibles. The relevant empirical unit is the "case," the infected individual who has been classified according to "diagnostic criteria," the case whose "cure" is the goal of the field.

Even though "medical" has of course expanded in general during the previous century to include non-biological personal and social factors, staffing patterns and funding priorities at such centers of the drug field as the National Institute on Drug Abuse show that the medical paradigm remains dominant in all the ways implied above. A glance at the drug field in most other countries would show a similar hegemony in the medical arena.

The problem is, many aspects of drug use don't fit a medical paradigm in any straightforward way. What kind of disease is it that some people want to catch while others decide not to catch it? What kind of disease is it where portals of entry and exit and vectors are group-specific and symbolic rather than biological? What kind of disease consists of positive effects? What kind of disease is it that is encouraged if the social and political position of those who manufacture and distribute the psychoactive drug is mainstream? Think of Valium, Prozac, Ritalin, Oxycontin.

What if… What if illicit drugs were more like consumer products than they are like a disease?

The Ethnographic Background

Let us now describe our specific paradigm-busting case. The first step for research in human worlds, as outlined at the beginning of this article, is ethnographic work with the purpose of generating "what if" questions based on local concepts and practices. It is actually embarrassing how easy it is to do this. The "official" frameworks among experts are typically distorted by distance and expert interests. It is amazing that any policy works at all. In fact, when policies do work, it is probably more often

testimony to the adaptive ingenuity of the represented population rather than to any accuracy in the representation.

The ethnographic part of the paradigm-busting process is given short shrift in this article, because much of the ethnographic and historical research we did on illicit drug epidemic cases has been published elsewhere (see, for example, Agar and Reisinger, 2000, 2001). For now, we will just say that our research on illicit drug epidemics, not to mention numerous other studies, suggests that circulation of stories through social networks drives increases and decreases in illicit drug use. Specifically, in one study of white suburban youth involved in heroin experimentation in the Baltimore suburbs in the late 1990s, youth described the "buzz" around heroin, stories that changed with time as early experiments by risk-takers evolved into widespread experimentation and then turned more complicated as negative stories about physical dependence worked against those early positive accounts.

This sounds more like a product evaluation by consumers than a disease being transmitted from infected to susceptible. While the biological basis of a drug experience is relevant, the critical issue for the youth was the phenomenology of the experience—good or bad or both. Those experiences—whether one's own, or witnessed, or simply heard about—were then in turn conveyed to others. The dynamic that explained an epidemic of use, in other words, was driven by interaction among good and bad stories, with good stories appearing initially and bad increasing with time.

It looked like dynamic circulation of narratives among agents would generate an epidemic incidence curve, the classic S curve, all by itself. In fact, most illicit drug epidemics show a flattening incidence curve well before any policy reaction takes place. Perhaps use increases and decreases "naturally," due to inter-agent dynamics rather than due to externally imposed sanctions. "What if…" the experts hadn't made much difference? What if they were taking credit for something that was already over and done with?

Based on these "What If…" ethnographic conclusions, we moved on to step two, work with agent-based models. A model based on circulating narratives does indeed generate incidence curves like those observed in epidemiological graphs based on such sources as arrests and treatment admissions. It is this part of the paradigm-busting research process that we want to feature in this article. We will show the importance of three model parameters, learned from the youth as important, that will force the old medical paradigm to break. Using Design of Experiment (DOE) approaches, we will show that the DOE analysis does in fact support the need to re-think medical paradigms for understanding, explaining and intervening in illicit drug use epidemics.

The Agent-Based Model

First we need to describe the agent-based model. The model, called DrugTalk, has gone through several incarnations. Common to all of them is a simple idea born of what the youth taught us in our Baltimore research. The model gives each agent a risk and an attitude. "Risk" is the willingness to try something new and unknown. "Attitude" is the degree of aversion to illicit drug use. Risk is fixed, but attitude can change, depending on

what an agent experiences and "hears" from other agents. Whether or not an agent uses a drug depends on whether or not risk is greater than attitude.

At first heroin is made available at one location on a torus, on which agents move randomly. If an agent does use the drug, it evaluates the experience, communicates with its primary social network, and offers them the drug as well. And all agents, all the time, check the attitude value of the agents that surround them as they move about.

Let us use part of an earlier article in this section (Agar, 2004) to better describe the details of the program. Each agent must have a risk and an attitude value at the beginning, since the comparison between the two determines whether or not an agent will try a drug. All agents are assigned the same attitude value initially. This represents a general orientation to use on the part of a particular population, a "norm," if you will. Attitude of individual agents will change during a simulation run, sometimes dramatically.

Risk, on the other hand will be different for each agent. And, in contrast to attitude, an agent's risk will not change during the simulation. The assumption is that risk is a fairly stable and pervasive characteristic of an agent, that there are risk-takers and risk-avoiders, and that those proclivities hold up across different situations. Rogers' Diffusion of Innovation (Rogers, 1995) reports numerous studies across many different domains. A robust result of those studies is that people's willingness to take a chance on an innovation is normally distributed. With this body of work as background, then, risk values are assigned to agents using a random-normal distribution.

Barabassi (2002) argues that social networks show an inverse power law distribution. That is, a few agents will have a lot of social links and many agents will have just a few links. After simple trial and error, an exponent of 1.5 produced a reasonable-looking distribution for a 500-agent world. Notice that there is no restriction on which agents might be assigned to a network. That is, the same agent might be selected at random more than once, or even several times, or perhaps never. And it doesn't matter what the selected agent's network looks like, either. The resulting network, if graphed with number of agents on the Y axis and size of network on the X axis, will look like an inverse power law distribution. But the overall network, expressed as a digraph, will look very different from time to time.

Each agent moves randomly. Then the first thing it does is to check-the-buzz. If it has become an addict, it doesn't bother to check, because it doesn't matter what other people are "saying" about heroin any more. Check-the-buzz corresponds to what youth often told us, that you pick up on stories about drugs wherever you go, from people other than those in your personal network. A party, a club, an event, school, a part-time job--drug stories are often "tellable" in these settings, since they can be dramatic, surprising, something out of the ordinary.

How does the buzz get checked? Each agent keeps a record of how many positive and negative experiences it has had with the drug. To check the buzz, an agent just adds up the total number of positive and negative experiences among the agents on its own patch or within a radius of two patches. The agent who is doing the checking then adjusts its attitude by these numbers, subtracting the positive total from its attitude to make it more likely to use and adding the negative total to make it less likely.

Here we introduce a bias based on Tversky and Kahnemann's prospect theory (Tversky, Kahneman and Slovic, 1982). People, say the hundreds of studies that have now been done, want to minimize loss more than they want to maximize gain. Therefore, an agent will put more emphasis on the negative total than it will on the positive total. So the negative total is multiplied by two to represent this effect.

The overall effect of checking the buzz is low when compared to procedures to come. This is as it should be, since hearing things from strangers you just happen to run across has less effect than a story from a trusted and long term friend. However, one buzz-checking experience can have a major effect on attitude. If an addicted agent is also in buzz range, the agent who is checking will raise their attitude by 20. (An addict is defined by a certain number of uses—a parameter—set to five in this case). Twenty is a substantial change since the range is zero to 100. (By the way, the range is always kept between zero and one hundred. It can go no higher or lower.)

There is some justification for this number—not this exact number, but rather for a number that represents a "big" difference. For one thing, youth reported such reactions—"I was experimenting, or thinking of trying it out, and then I ran into so-and-so who'd turned into a junkie, and it really turned me off." Other evidence comes from Musto's concept of "generational forgetting" (Musto, 1999)—after an illicit drug epidemic impacts one generation, the next generation tends not to use, since they've seen use go from pleasant early on to devastating for addicts and communities down the road. Recent observations in Baltimore and other cities suggest that African-American youth, having witnessed the crack epidemic, will have nothing to do with use of the drug, though a few will sell it as a lucrative niche in the underground economy. For many such youth now, a "drug-related problem" means dealing, not using.

Next comes the moment of truth. If an agent is on a red patch, meaning heroin is available there, it compares its risk to its attitude, and if risk is higher, it uses the drug.

Right after it uses, an agent evaluates the experience with a function called how-was-it, unless the agent is already an addict, in which case it doesn't matter any more. To understand how this function works, we first have to look at two more parameters that will play a major role in the analysis to come, the first called goodstuff?, the second, badstuff? Each one can vary between zero and 100. Broadly speaking, this number means to be a kind of quality evaluation. For the moment, we ignore problems of individual variation and context and assume there's some kind of average that makes sense. Overall, does the drug produce a pretty good or a pretty bad experience? And notice that both things can be true—in other words a user might have an experience that he/she would describe as both good and bad.

The rest of how-was-it is simple. After an agent has used, it generates a random number between zero and one hundred, and if goodstuff? is larger than that number, the agent records a positive experience. Then it changes its attitude in a favorable direction—i.e. it decreases it—by an amount equal to ((1 / positive) * 20). Notice how the effect of the evaluation diminishes with increased use. The first positive experience reduces attitude by 20; the second, by 10; the third, by 6.67; and so on.

And then, independently of how the goodstuff? evaluation went, the agent does the same evaluation using badstuff? The difference here, of course, is that if badstuff? is

larger than a random number between zero and 100, the value of the agent's attitude increases to make it less likely to use. And another difference, corresponding to the prospect theory principle that people are risk-aversive, as described earlier: This time the value changes by 40 instead of by 20. The impact changes with experience, just as it did with "goodstuff?"—40 the first time, 20 the second time, 13.33 the third time, and so on.

The justification for the diminishing impact lies in intuitions about "habit," that the first experience of anything is the most significant, with subsequent experiences showing a "I'm getting used to it" effect. There is a literature that supports this assumption going back to old-fashioned behaviorist psychology that we take for granted here.

The next thing that agents do, right after they use, is let their network know with tell-the-network. Recall that the model was set up with an inverse power law social network distribution, that is, a few agents will have large networks and a lot of agents will have small networks. An agent who has just used checks its network members. If a network member is already an addict, the agent who used has no influence on its attitude. But if the network member isn't an addict, a couple of things might happen.

First of all, if the agent who has just used is itself an addict, it will "turn off" the members of its network by adding 20 to their attitude. Recall that the same thing happened if an agent found an addicted agent nearby when it checked the buzz around it.

If the agent who just used is not an addict, then something different happens. For each agent in its network, the agent who just used "pulls" them in the direction of its attitude, whatever it might be. It does this by the simple mechanism of assigning the agent in its network the average of its own and that agent's attitude.

The agent who used will have an attitude that reflects its history of positive and negative experiences from checking the buzz and evaluating its own use. Tell-the-network will move the agent in its network towards its current attitude that reflects those experiences. The assumption is that if the agent who used is becoming more positive, it will make its network more positive. If it is becoming more negative, then it will make its network more negative. Since all agents begin with the same attitude value, the attitude carries the cumulative positive and negative history of an agent with the drug, so we want it to pull its network members in the direction of how that history has changed after drug use.

Whatever the outcome of all this influencing, or lack thereof, the agent who just used always offers heroin to all the agents in its network, no matter what. If the agents in the network have a risk greater than their attitude, they use the heroin and evaluate the experience, as the original agent did, with the same procedure, how-was-it. But at that point the network member stops. In other words, the agent in the network does not, in turn, offer heroin to other agents in its own network. Perhaps it should, not immediately, in that particular tick of the program, but with some time lag.

That's pretty much the interesting part of the program. Comparison with actual cases together with observations of how DrugTalk generally behaves show that we're on the right track. The paradigm-busting argument learned from the youth—that a new drug is a consumer product—works as well when we model it as it did when we heard it in

interviews. But—the classic problem with ABM's—there are many parameters that can vary in the model. One strategy here is simply to set up multiple runs to explore the space of possible outcomes by sweeping parameter values with regular intervals, something explored in a preliminary way in an earlier article. We decided to try something different, the Design of Experiment (DOE) approach.

Design of Experiment

DOE was developed to analyze real-world experiments where there is a practical limit to the number of experiments that can be performed, either because they are expensive or because they are slow. With an agent-based models like DrugTalk, as we saw in the previous sections, there are many parameters to the model—*attitude*, *goodstuff?*, and so on. Ideally, to explore the model, we would want to run it under all possible combinations of parameter values. The number of runs would be enormous. The motive for using DOE is that traditional 'parameter sweeps' suffer from combinatorial explosion. Sweeping 10 parameters (as in this study) through 10 values each, requires 10^{10} = 10 billion experimental runs. If each simulation took 1 minute to run, this would represent a significant investment in time, close to twenty thousand years, in fact.

The sorts of questions that DOE can answer are slightly different in the context of simulations than in real-world experimentation. The latter tend to focus on optimization and prediction, trying to closely specify the input values that produce some desired effect in the output. In contrast, when DOE is applied to simulations, it answers broader questions:

- searching for insights or developing basic understanding of a simulation or system
- finding robust configurations, decisions or policies
- comparing configurations, decisions or policies (Kleijnen et al, 2004).

Typical DOE designs involve specifying a series of experiments that take place at the maxima and minima of the various parameters (in a 2-level design) or that take place at the maxima, minima and midpoint (in a 3-level design). The maxima and minima are chosen by the experimenter and represent the range over which he/she wants to study the simulation. However, instead of running all possible experiments with these two or three values per variable (because again, 2^n and 3^n explode with n), a limited subset of experiments is carried out. This subset is balanced so that the values of any variable are equally represented; e.g. in the three-level experiment used in this study, one-third of the experiments were run with *goodstuff?* at its minimum value, one-third at its midpoint and one-third at its maximum value. Depending on how restricted the subset, combinations of variables show the same balance (e.g. the nine possible combinations of *goodstuff?* and *badstuff?* appear the same number of times in the experiments). It is this balance that gives DOE its validity.

The current study uses a 3-level design of 81 experiments that evaluates the 'main effects' of 10 variables. (In terms of the 'balance' concept above, this means that the three values of each of the 10 simulation inputs are equally represented in the

experiments.) Main effects are a simple way of finding out what the important variables are: for a particular input variable, the 81 experimental results are divided by high, average and low values into three groups of 27 and the mean of each of these three groups is calculated. This is repeated for each input variable; the three means for each input are then plotted in graphs. These graphs show which input variables have a large impact on the simulation output (i.e. there is a big difference between the smallest and largest values in that variable's plot).

DOE Results

The DOE analysis contains one small and one large surprise, both of which suggest parts of the medical paradigm that need busting for the drug field. Examine Figures 1 through 3. Each figure shows the results for a different key outcome variable from the model. Figure 1 shows the effect of the ten parameters on the total number of users. Figure 2 shows effects for the total number of addicts. And Figure 3 shows effects for the way agents become more or less at risk.

Each figure, in turn, contains ten charts, one for each of the parameters we looked at. Each chart shows an angular line. The more sharply vertical the line, the stronger is the effect of that parameter on that outcome. Limitations of space prohibit a full discussion of all the parameters tested, though we will be happy to provide additional information on request. But most of them should be familiar from a reading of the model details in a previous section.

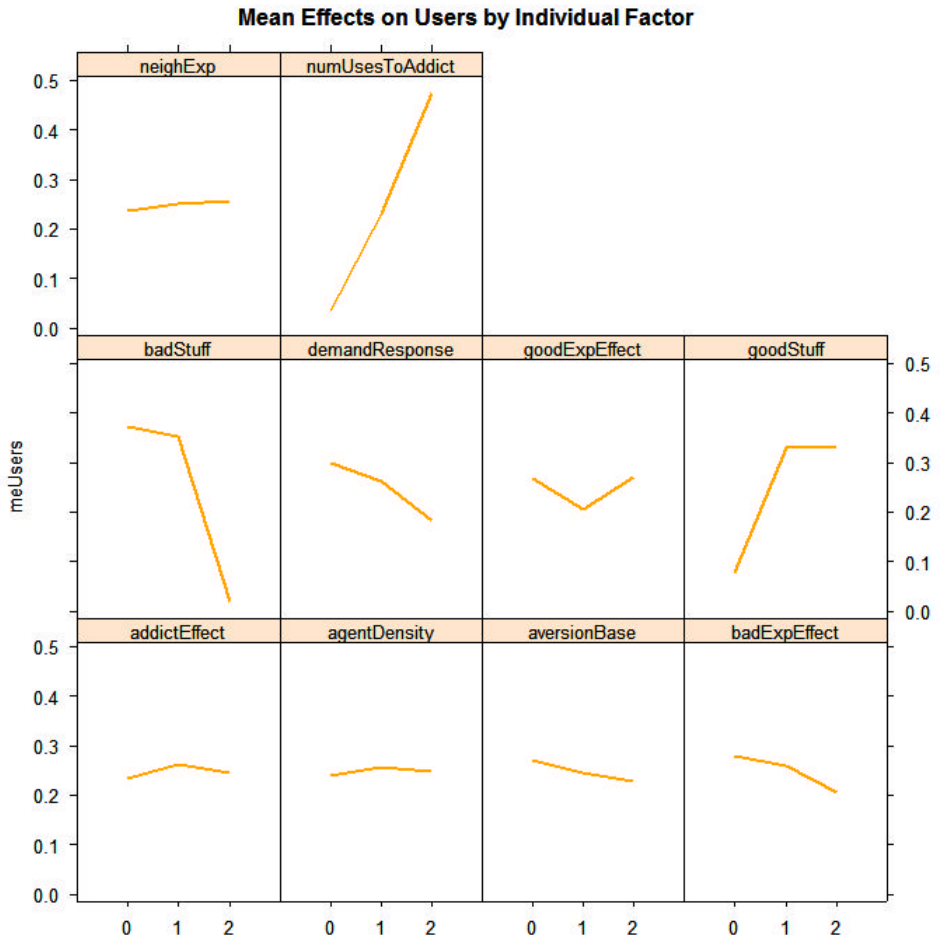

FIGURE 1 Parameter Effects on Number of Agents that Used the Drug at Least Once

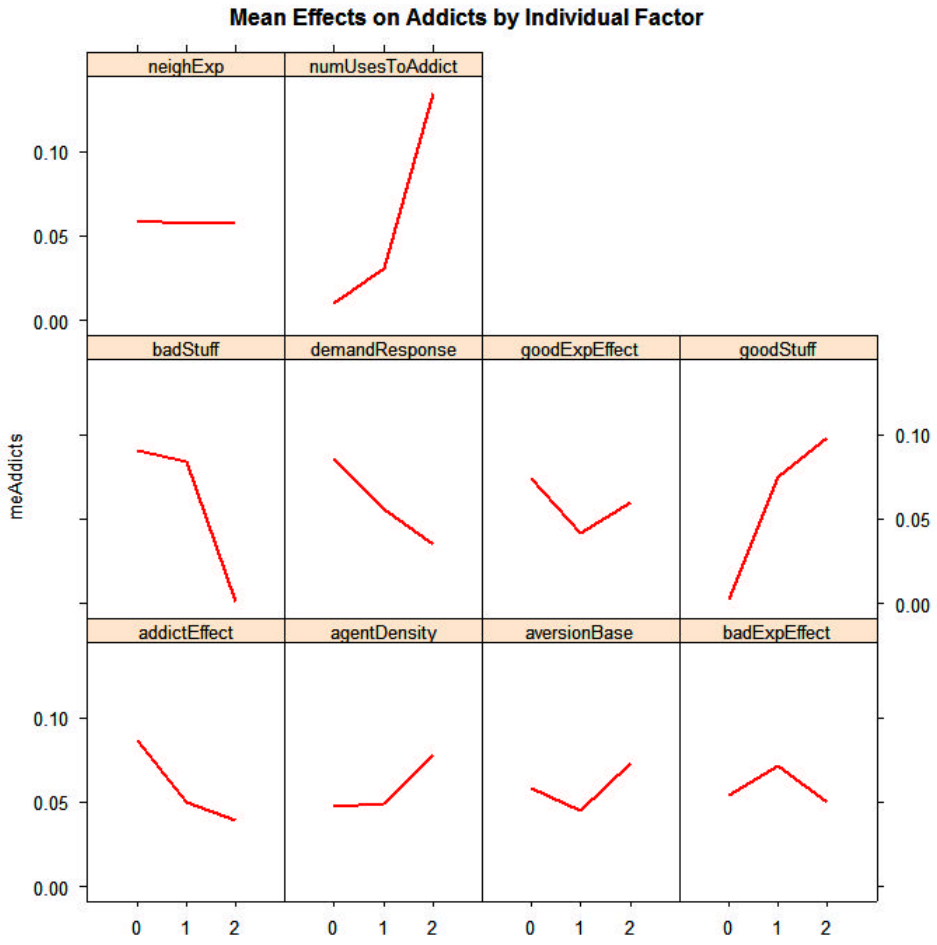

FIGURE 2 Parameter Effects on Number of Addicts

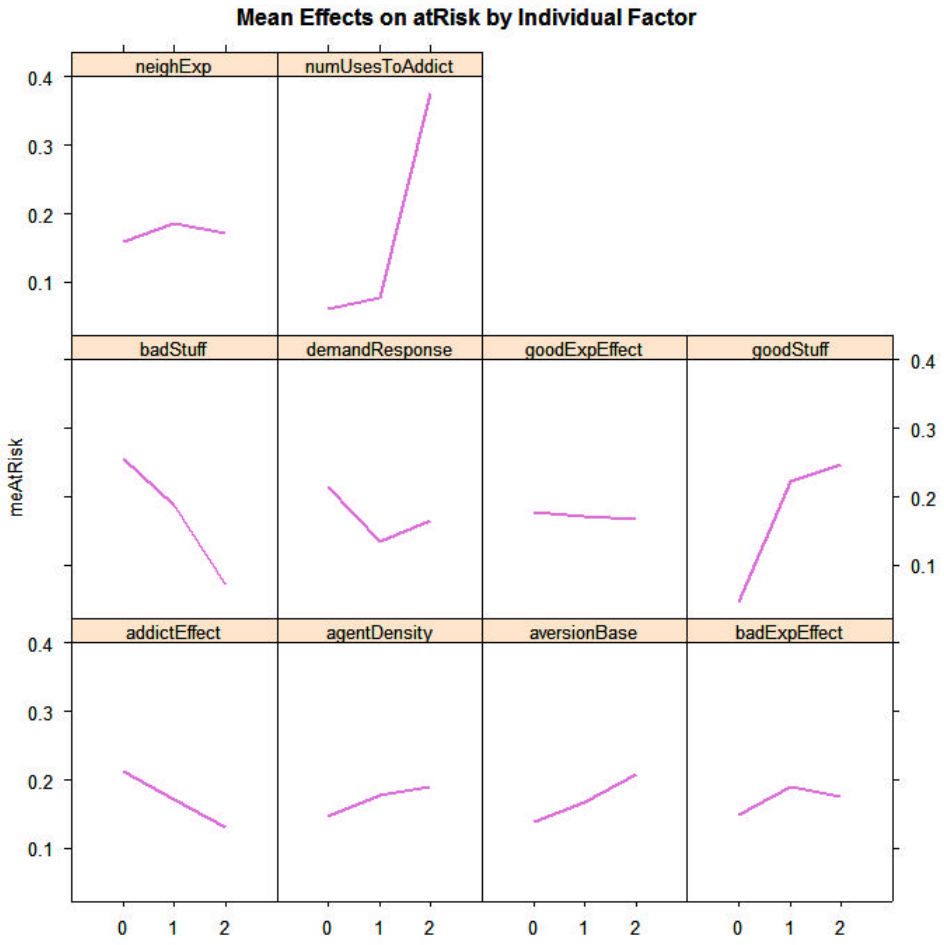

FIGURE 3 Parameter Effects on Number of Users "At Risk"

Here are the seven parameters we tested that are not already discussed above. The abbreviations refer to the titles in Figures 1 through 3:

neighExp - The exponent that defines the initial network distributions

demandRepsonse - The speed with which additional heroin patches are created

goodExpEffect - The strength of a particular good experience

addictEffect - The effect of having an addict in the neighborhood

agentDensity - The density of agents in the model

aversionBase - The initial setting of attitude

badExpEffect - The strength of a particular bad experience

Some of the charts we neglect here do show vertical angularity. But if the reader looks across the Figures and scans for repeated extreme angularity for the same parameter, three stand out in striking fashion.

First the small surprise, actually two parameters, but they are two sides of the same coin. They are “goodstuff” and “badstuff,” informal labels that echo ordinary conversation. They represent the quality of the drug as the user experiences it. Recall that “badstuff” has a stronger impact to reflect the findings of Prospect Theory. And recall that the effect of “goodstuff” and “badstuff” on agent attitude declines with number of uses. Beginnings are most important.

“Goodstuff” and “Badstuff” collapse many things into a single number. Drug effects on a particular person at a particular time can change with biochemistry, with set and setting, and with the particular biographical and historical situation. These are complicated parameters. In the end, though, they’re appropriate at a phenomenological and social-interactional level because, in the end, a new user is a person who tries something and tells stories to other people about how good or bad it was.

The DOE analysis foregrounds the importance of these parameters. It makes the hidden fact explicit, because the fact was invisible to the medical paradigm. For an experimenter, an illicit drug is a commodity to be evaluated, not a disease to be caught or avoided. Ironically enough, when an earlier version of DrugTalk was presented at the UCLA conference on agent-based modeling in the social sciences, the organizers put it in a session called “marketing.” Untainted as they were by a medical paradigm, they saw the model a different way, right away.

The idea that a drug is a commodity that behaves like other commodities is not an alien concept in the drug field, though it is not a frequent one, either. The idea just doesn’t fit the medical tradition. In that tradition, any use is to be discouraged, so any use must be negative. At times it seems like an implicit guideline operates: There can be no reason why anyone would want to use an illicit drug—i.e. to catch a disease—so use must be caused by pathology in the biological or psychological or sociological system.

This change in how illicit drugs should be understood—commodity, not disease—surely counts as a paradigm shift, a problematic one given mainstream U.S. war on drugs policy. Would the concept allow more effective intervention? Before considering that question, let us describe the second surprise, the large one.

One parameter in DrugTalk is, how many uses have to occur before “addiction” sets in, “numUses ToAddict” in the Figures. “Addiction” is a loaded and ambiguous term, since the original meaning was the actual physical addiction produced by opiates. Now the drug field uses the more general DSM-IV concept of “dependence.” That concept defines “dependence” in psychological and sociological terms, that is, in terms of how an individual’s behavior changes.

The critical changes basically show a shift from personal control over use of the chemical to chemical control over much, even most, of what a person does. Smoking a joint on Saturday night is one thing: needing a joint six times a day is another. When most of what you do with your time is get the chemical and use the chemical and figure out how to get money to buy the chemical and think/talk about the chemical, etc., you

obviously are dependent on that chemical. That is, in fact, a problem for you, for your non-dependent family and friends, for your studies or your work, and for your community.

Dependency should certainly make a difference in outcomes, so the fact that it did in the DOE analysis is not a surprise. A product that makes you dependent should literally capture market share. As Warren Buffet explained, speaking about a legal drug, "I'll tell you why I like the cigarette business. It costs a penny to make. Sell it for a dollar. It's addictive. And there's fantastic brand loyalty." The DOE, in fact, showed that dependence is the most important parameter of all.

The large surprise, though, is this: Intuitively one would think that the faster a drug produces dependence—the fewer uses it takes—the more addicts it will produce in the end. Get them quick and you'll get a lot of them. But as it turned out, the DOE supported the opposite conclusion. The longer it took, the more uses it took before an agent became dependent, the more addicts were produced in the end.

How could this be? What in the model explains this peculiar result? Once again the model makes a hidden fact clear. Recall that it builds an event into an epidemic, something youth often talked about when we interviewed them. Once dependent persons appear in an agent's friendship network or in its neighborhood, those "visible" examples of what that particular drug can do to a person have a negative effect. In fact, such events produce the strongest increase in attitude that ever occurs in the model. So that is probably the explanation: If no dependent agents show up for a long time, attitude will increase more slowly and less dramatically than if dependent agents do quickly appear.

Colleagues in the drug field sometimes joke after a presentation of this model. The best thing to do for a new wave of heroin experimentation would be to fly in dozens of addicts and distribute them throughout the social world of the group that is experimenting. The DOE analysis explains the joke. It's not funny.

So another bit of the medical paradigm encounters difficulties. First of all, DOE results suggest that a new drug that makes a big splash is to be less feared than a stealth drug that can be used for substantial periods of time before "dependency" appears publicly and deters use.

Second, if a drug seldom causes dependency in a way that will publicly deter use, it won't go away once it gets going absent draconian punishment. Consider marijuana as the classic case.

Third, and most devastating for the medical paradigm: The old notion of "addiction" as a matter of biological dependence is clearly inadequate. This is not news, as already noted, since the field now talks in terms of DSM-IV. But the notion of "dependence" as primarily an individual problem with intrapsychic causes is inadequate as well. It might be important for clinical work, but it will not explain the shape of an incidence curve. The negative effect comes from social impact. This is a robust theme of ABMs in general—individual level properties won't explain system level phenomena.

Critical for the power of the "addict" parameter in DrugTalk are what agents "see" around them as the social consequences of use become public. A biologically addicted psychopath who behaved himself in public would not have an impact on other

agents' attitudes. Without being aware of it, we told the model that what counts under "dependence" is the observation by other agents that continual use of an illicit drug can have a negative impact on their social world. That's the theme of many of the stories that the youth told us—"And then I saw so-and-so, he was a junkie, and what a mess."

Policy Implications

As a result of ethnographic research, the agent-based model, and the DOE analysis, we see that drugs can be viewed as a commodity like any other. And we test the idea that the major deterrent to dependence is personal experience and/or stories from networks that dependence is a socially destructive condition. The implications of this paradigm change are massive and beyond the scope of this presentation. Let us just outline a few:

1. In a social world that is open to illicit drug experimentation, any drug that is high on goodstuff and low on badstuff will be tried if the market can provide it. A wave of experimentation will occur. Trying to prevent this wave is futile.
2. Credible drug education must recognize the positive quality of the product, something they seldom do, as far as we are aware.
3. Many drugs can produce traumatic results on first use, and these should be a topic in prevention. But they must be presented so that they correspond with actual experiences with which the population will already be familiar. And they must not be presented as the inevitable, or perhaps even likely, outcomes of experimentation, nor must they be overemphasized by way of comparison with positive effects.
4. The most critical part of prevention is to prevent dependence. Educational materials should feature what life is like on the other side of dependence, realistically, with examples.
5. Group sessions can serve as "story amplification" devices. Assuming dependence has already occurred, likely given the time lag between epidemic and policy response, potential and actual users will already be familiar with its effects. Program time should be dedicated to participants telling stories about themselves and people they know. Material for prevention is in fact available in the worlds of experimenters and non-users. This only amplifies what naturally happens anyway, as reflected in ethnographic interviews and in DrugTalk. Group sponsors must accept that some of the stories that will be told about experimentation will be positive.
6. The most important programs will deal with early intervention, something which is now rare. By this is meant that if an experimenter shifts to a user shifts to a frequent user, he/she is "at risk" for dependence. Early intervention is an effort to intervene with a serious user on the edge of dependence and pull him/her back. Identifying such serious users typically occurs among the friends, family, organizations or communities where they spend their time. Early intervention referral may be a productive use of such law enforcement mechanisms as the drug courts.

7. Dependence will occur, and such dependent users will require treatment. Part of their treatment could be community service where they can serve as speakers, assuming they are peers of the experimenting population, to tell stories of the line between use and dependence, how they crossed it, and what the personal consequences were. Often when former addict speakers are brought in, in our experience, a forty-something ex-heroin addict addresses a high school group. This is less credible than a session with a peer.

There are other implications to spell out, and the list above is already controversial when measured against traditional war on drugs practices. As far as we know, suggestions such as those in the list above have not been tried extensively or consistently. They of course might not work, but they should be tried. We are at a juncture where it is widely recognized that the war on drugs has failed. New alternatives are in order. The problem with the medical paradigm, not to mention the legal paradigm that we have not dealt with here, is that they have not generated any new ideas.

While our primary purpose here has been to demonstrate an ethnography/ABM collaboration as a paradigm-busting device, we also want to emphasize that both the old and the new paradigms that define a particular application may well have massive social and political consequences. They certainly do for DrugTalk. The opportunity for real change in social practices is enormous, though implementation raises political issues that go well beyond the research framework suggested here.

REFERENCES

Agar, M.H., 2004, "Agents in Living Color: Towards Emic Agent-Based Models," *Journal of Artificial Societies and Social Simulation*, in press.

Agar, M. H. and H.S. Reisinger, 2000, "Explaining Drug Use Trends: Heroin Use in Baltimore County," in *Illicit Drugs: Patterns of Use-Patterns of Response,* A. Springer and A. Uhl, eds., pp. 143-165, Insbruck Austria: STUDIENverlag.

Agar, M.H. and H.S. Reisinger, 2001, "Trend Theory: Explaining Heroin Use Trends," *Journal of Psychoactive Drugs* 33(3): 203-212.

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

Barabasi, A., 2002, *Linked: The New Science of Networks*, Cambridge MA: Perseus Books.

Epstein, J.M. and R. Axell (SP?), 1996, *Growing Artificial Societies: Social Science from the Bottom Up,* Cambridge MA: MIT Press.

Kleijnen, J.P.C., Sanchez, S. M., Lucas, T. W. and Cioppa, T. M., 2004, "A User's Guide to the Brave New World of Designing Simulation Experiments," *Working Paper, Tilburg University, Tilburg, The Netherlands.*

Musto, D.F., 1999, *The American Disease: Origins of Narcotic Control*, New York: Oxford University Press.

Rogers, E.M., 1995, *Diffusion of Innovations*, New York: The Free Press.

Tversky, A., Kahneman, D., and Paul Slovic, 1982, *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge UK: Cambridge University Press.