

The Standing Ovation Problem

John H. Miller and Scott E. Page

July 27, 2006

1 Introduction

Over the last decade, research topics such as learning, heterogeneity, networks, diffusion, and externalities, have moved from the fringe to the frontier in the social sciences. In large part this new research agenda has been driven by key tools and ideas emerging from the study of complex adaptive systems. Research is often inspired by simple models that provide a rich domain from which to explore the world. Indeed, in complex systems, Bak's (1996) sand pile, Arthur's (1994) El Farol bar, and Kauffman's (1989) *NK* system have provided such inspirations. Here we introduce another model we believe to offer food for thought—the Standing Ovation Problem (SOP). This introduction is especially appropriate given the focus of this special issue on complex adaptive social systems. The SOP has much to offer (1) it is easily explained and part of everyone's common experience; (2) it simultaneously emphasizes some of the key themes that arise in social systems, such as learning, heterogeneity, incentives, and networks; and (3) it is amenable to research efforts across a variety of fields. These features make it an ideal platform from which to explore the power, promise, and pitfalls of complexity modeling in social science.

Inspired by the seminal work of Schelling¹ (1978), the SOP possesses sufficient structure to generate nontrivial dynamics without imposing too

¹Schelling's work represents one of the first examples of a style of theorizing about social systems that naturally arises using computational models—notwithstanding the notable absence of a computer. His work focused on the macro-behavior that emerges from micro-motives, and often relied on models that emphasized agents driven by simple behavioral algorithms placed in interesting spatial contexts.

many *a priori* modeling constraints. The basic SOP can be stated as: A brilliant economics lecture ends and the audience begins to applaud. The applause builds and tentatively, a few audience members may or may not decide to stand. Does a standing ovation ensue or does the enthusiasm fizzle?

Though ostensibly simple, the social dynamics responsible for a standing ovation are complex. As the performance ends, each audience member must decide whether or not to stand. Of course, if the decision to stand is simply a personal choice based on the individual's own assessment of the worth of the performance, the problem becomes trivial. However, people do not stand solely based upon their own impressions of the performance. A seated audience member surrounded by people standing might be enticed to stand, even if he hated the performance. This behavioral mimicry could be strategic (the agent wants to send the right signal to the lecturer). It could be informational (maybe the lecture was better than he thought), or it could be conformal (the other people are standing so he stands as well so as to not feel awkward). Regardless of their source, the introduction of these peer effects sets the stage (so to speak) for interesting dynamic behavior.

The SOP admits a variety of modeling strategies. Over a four year period, we confronted graduate students entering the Santa Fe Institute's summer workshop in computational modeling and complexity with the SOP and found the problem to be rich in possibilities and insights. Fundamental to the success of the SOP is its ability to force modelers to confront the core methodological issue in complex adaptive social systems, namely, how does one model a system of thoughtful, interacting agents in time and space. Moreover, the SOP forces modelers to take the details seriously. These details include: How do agents influence one another? How sophisticated are agents? How does information spread among agents? In what order do events occur? At what time scales do events occur?

Methodological perspectives can be deeply ingrained. Prior to presenting the SOP to graduate students in economics, we tested it on Cal Tech undergraduates. Though Cal Tech undergraduates are hardly a random sample, we did find that their modeling efforts differed in fundamental ways from those of graduate students in economics. The undergraduates assumed that individuals sat next to close friends (or, even went to the lecture with dates). In contrast, very few economic graduate students included the possibility of friends in their models. This difference might be a reflection of the social life of budding economists, but we remind you that the comparison group here

is Cal Tech undergraduates. We suspect that the divergence in assumptions is much more due to the emphasis on individual choice that pervades most of modern economic theory, rather than social differences between the two groups of students.

The SOP is an apt metaphor for social situations in which agents make binary decisions and interact spatially. It applies to a wide ranging set of phenomena such as whether to send children to public or private school, to commit crimes (Glaser, *et al.*, 1996), to violate the law (Picker, 1997), to riot (Granovetter, 1978), to search for jobs (Menczer and Tassier, 2001), to retire (Axtell and Epstein, 1999), to vote for a particular party (Mayer and Brown, 1998), to experiment with drugs, to engage in unprotected or premarital sex (Durlauf, 1997), to pay your electric bill, or even whether to decorate your house with strands of multi-colored bulbs during the holiday season. These various phenomena all share elements of the SOP: people are socially influenced, they have varying degrees of sophistication, and information flows over a network.

The SOP can be used to explore some intriguing policy questions. We often pose the following question to our students: suppose you can place some skills in the audience, where would you place them, and how should they act in order to maximize (or minimize) the probability of a standing ovation? Other policy questions can also be addressed, for example, consider the architecture of the performance hall. Does the presence of a balcony alter the probability of an ovation? Of course, whether the Phantom of the Opera receives a standing ovation (or not) is of little or no global concern, but if we interpret standing as taking drugs, committing crimes, abstaining from dangerous sexual practices, attending school, then we can attach much more normative significance to our ability to prevent and create ovations.

Along with policy prescriptions, these models often provide other insights. For example, suppose we lower the overall level of satisfaction with the performance while increasing its variance—intuitively, this type of change would seem to lessen the probability of an ovation. However, in many models, the opposite occurs. On its own, this is a cute insight, but given the tendency of social scientists to rely on means it suggests that we may easily miss some key drivers of social systems—when social influences are present, the tail (of the distribution) may wag the dog.

2 The Methodology of Agent-Based Models

A computational, agent-based model includes interacting agents who rely on (possibly adaptive) computer algorithms to determine their behavior.² Agents can interact in both space and time, creating dynamic patterns, and potentially perpetually novel behavior. Computational models need not be bound by the limits of analytic tractability, though the usual modeling desiderata of elegance and parsimony still apply. The models also permit heterogeneity in not only agent preferences but also their behavior. The behavior of agents in these models might be “tunable” in several dimensions, ranging from hyper-rational and hyper-informed to simple and naive.

The computational approach to theorizing discussed here has been highly successful in the physical sciences; Notwithstanding this success, the inherent elements of social systems potentially combine to create a much more difficult modeling task. Models of interacting carbon atoms can ignore some of the most perplexing issues that arise in systems of even minimally intelligent social beings—carbon atoms do not (as far as we know) form expectations about their world or strategize about their behavior. Even in simple social environments like the standing ovation, issues of expectations and strategy can quickly begin to dominate the analysis. Of course, there is always the possibility that a limited description (which is the essence of good modeling) might still be able to capture these key elements of social interaction.

Note that the way we limit the descriptions of our models is closely tied to our tools. Mathematics requires a different set of refinements of the world than computation. Whether constructing a mathematical or a computational model, the tools at our disposal partially determine the simplifications that we, as scientists, must impose on the world. Even the most advanced mathematical models of social phenomena, such as those used in general equilibrium theory, exhibit the residue of this trade off (for example, in market models the absence of a compelling story of how prices form).

In modeling the SOP, one must explicitly account for many aspects of social interaction. Here, we shall discuss just three: the spread of information, the timing of events, and the behavior of the agents.

Computational models allow for a variety of assumptions about informa-

²Rather than give a full treatment of computational, agent-based models here, we refer interested readers to papers by Holland and Miller (1991), Judd (1997), Tesfatsion (1997), and Page (1999).

tion transmission. Information can be assumed to emanate from a single source and flow to the agents according to a distribution as in many mathematical models. Alternatively, information can be given a more explicit micro-structure. For example, agents can have friends and can get information from their friends, from their friends' friends, and so on. This enhanced micro-structure can lead to interesting dynamics, especially in the case of standing ovations where the agents that can communicate their information best—those in the front row—have the worst information about what others are doing.

Timing often plays a critical role in computational models (and, one suspects, in the world in general). Economists have long recognized the importance of timing, for example, consider the differences arising between the Stackelberg and the Cournot outcomes in an oligopoly model or the strategic issues that emerge in extensive versus normal form games. Yet, much of current modeling ignores timing issues by concentrating on equilibria and asymptotic behavior. Such a focus does indeed make timing irrelevant—to the model—but it does not eliminate such issues from the real world. Seemingly irrelevant timing constructions can lead to drastically different outcomes. What happens, when, matters. Specifically, whether agents make decisions and take actions synchronously, asynchronously, or endogenously can lead to important differences in model behavior (Glance and Huberman, 1993; and Page, 1997). The notion of process is largely ignored (or, at the very least, assumed to be of a particular form) in a lot of current modeling. Early computational results suggest that process may be a much more important determinant of behavior than previously assumed.

Another area for investigation is into the behavioral assumptions of our agents. In computational models, agents tend not to be fully rational. Of course, there is no *a priori* reason why fully rational agents cannot be used in such models, and the fact that these models do not use such agents is more likely a reflection of the flexibility of the tools and a desire to investigate non-optimizing agents. If agents rely on non-rational rules, then the question arises of how best to be human. While there may be only one way to be smart in the world, there are many ways to be less than perfect. If each modeler implements a different way to be less than perfect, *and* if each of these ways of being less than perfect leads to a different model behavior, then the basis for a science of adaptive agents is lost. However, strong evidence has emerged from computational models that there may be large equivalence

classes of adaptive behavior, and thus we may indeed be able to form a science of adaptation. This equivalence appears to hold in the SOP. The SOP solutions formed to date might be metaphorically thought of as cars in a show room: they differ in their trim packages, but they all respond roughly in the same way to similar inputs.

Operationally, we advocate including several types of agent behavior and emphasizing results that appear invariant to such choices. For example, in Kollman, Miller, and Page (1992) we find that political parties tend to converge to the center of the platform space whether they use random search, hill-climbing, or a genetic algorithm to locate new platforms. We also experimented with other behavioral rules, and found that they all generated similar convergence. Many others have commented on the need for robustness in computational models. Holland (1988) makes the important point that we need to avoid “brittleness.” Though at one level the notion of avoiding brittleness in models seems obvious, a lot of our current theories in the social sciences seem more concerned with modeling unicorns than horses.

To some extent, robustness rests in the eye of the beholder. Computational modeling allows the researcher to easily alter the assumed behaviors and parameters in an effort to identify the key factors driving the results. One way to search systematically the space of alternatives, proposed by Miller (1998), is called Active Nonlinear Tests (ANTs). ANTs uses automated, nonlinear search algorithms to search over a model’s parameters and assumptions in an effort to break the model’s conclusions. The fact that models (computational or not) can be broken should not be a surprise—interesting models must be responsive to outside forces. Nonetheless, the inherent flexibility of computational modeling necessitates a careful exploration of brittleness.

3 Mathematical Theories

Using the standing ovation problem as a backdrop, we can comment on several distinct mathematical research agendas that concern diffusion, information aggregation, conformity, information cascades, and growth. We cannot overstate the importance of continually trying to link the various insights gained from different modeling tools to each other. Advances on a particular theoretical topic (like, understanding standing ovations) may require coordinated efforts that involve many different tools (say, pure mathematics and

computation) that can exploit each tool’s comparative advantages and the various “insight” externalities among the methods.

The peer effects in a standing ovation might exist for many reasons: audience members may gain utility by matching their behavior to that of surrounding agents (Durlauf, 1995; Bell, 2002; and Bernheim, 1994), agents may interpret standing as a discrete signal of quality and want to send the right signal themselves (Banerjee, 1992), or they may want to collectively send the correct signal. Each of these types of peer effects has been studied in depth mathematically. Taken as a group these models create a convincing analysis and demonstrate the ability of mathematical theory to provide novel and powerful insights. They also, however, leave one wanting, as what they produce does not have the “look and feel” of standing ovations, as they do little to uncover the relevant spatio-temporal dynamics. In short, they leave the auditorium door ajar for computational approaches.

3.1 Information Cascades

In the simplest approach to the SOP, audience members can send one of two signals: stand or sit. The inability to express intensity of preference can lead to inefficient herding, often referred to as information cascades. In an information cascade (Banerjee, 1992; and Bikhchandani, *et al.*, 1992), agents receive information sequentially and make binary choices, for example, buy or don’t buy. Agents know both their own signal and the choices of previous agents. The simplest models include a two-state world, *GOOD* or *BAD*, and agents who receive one of two signals, *HIGH* or *LOW*. If an agent receives the signal *HIGH* the probability that the true state of the world is *HIGH* equals p , where $p > 0.5$. Similarly, if an agent receives the signal *LOW*, then the probability that the true state of the world is *BAD* is p . Signals satisfy independence and agents receive them sequentially. Upon receiving her signal, an agent must choose either *GOOD* or *BAD*. If she is correct, then she obtains a utility of one, otherwise she gets nothing.

Rational agents would not only follow their signals since the choices of previous agents contain information. If an agent’s signal would lead him to buy, and he sees that the eight agents in front of him all chose not to buy, then he may rationally conclude not to buy. An *information cascade* occurs when at some point all agents ignore their signal and simply choose according to the actions of the previous agents. It can be shown that cascades are

inevitable. If the first few people all make the same choice, the others follow like lemmings. Rational lemmings, we might add, but lemmings nonetheless (Bikhchandani, *et al.*, 1992).

This simple model can be generalized to allow for different mathematical structures for the information, for public information, and for heterogeneity in the quality of information that the agents receive. These generalizations yield several counterintuitive insights: public information can be *ex ante* utility decreasing for some agents and the probability of an incorrect cascade can increase if the first agent decides to become better informed.

The standing ovation problem differs from information cascade models in three respects. First, decisions need not be made sequentially. Second, agents can change their decisions. An agent who initially decides not to stand may opt to get on her feet on the basis of later information. And third, agents get an initial signal of the reaction of others within their sightlines. This, as we shall show, can prevent cascades.

Taken together, these differences imply that an information cascade should be less likely to happen in a standing ovation than in a sequential, binary-decision herding model. Consider a stylized example with six audience members. We will show how sequential updating can lead to a cascade, and that this cascade would not occur in a standing ovation model. Suppose that the signals (*HIGH* and *LOW*) and states of the world (*GOOD* and *BAD*) are as before. Assume that agents 1 through 3 are in the front row, agents 4 and 5 are in the second row, and agent 6 is in the back row.

Agent 1 Agent 2 Agent 3
 Agent 4 Agent 5
 Agent 6

Assume that agents in the front row can look to their left and right, and that agents in rows two and three can see all agents in front of them, but cannot look to their left or right. Suppose that the signals are as given below:

GOOD BAD GOOD
BAD BAD
BAD

To see that sequential decisions by rows would lead to a cascade, let row one update until reaching an equilibrium. Initially, agents 1 and 3 stand, followed by agent 2.³ Imagine that agents in row two have been blindfolded until the first row reaches an equilibrium. They know that with equal probability either two or three agents were standing initially; therefore, their optimal decision is to stand. Agent 6's decision relies upon identical logic. If all five agents are standing, then she intuits that either two or three of the agents in the first row were standing initially. She can make no inferences about the signals of the second row, and her optimal decision is also to stand.

In a standing ovation, agents neither make decisions or receive information in a predefined sequence. Agents 4 and 5 have information that two of four signals were *HIGH*. Assume that they randomize in this environment. One quarter of the time, they would both stand. Even if both agents 4 and 5 stand, agent 6 will not. Her information, which in this example is complete, is that four of the six signals were *LOW*. The fact that information improves with row number stems the cascade.

3.2 Democratic Ovations

In democracies, agents care about outcomes and not their individual votes. This concern with outcomes prevents inefficient cascades in noiseless environments.

A number of recent papers construct models where agents receive private signals about candidate quality. Each agent has a prior probability over which decision (*GOOD* or *BAD*) generates higher utility, receives a signal, and updates her prior accordingly. The models rely on either Nash or Bayesian-Nash Equilibrium as the solution concept (Austen-Smith and Banks, 1996; and Feddersen and Pesendorfer, 1998), and generate provocative findings. First, in equilibrium an agent may vote insincerely, that is the opposite of how she would have voted if she were the only voter (Austen-Smith and Banks, 1996). Second, with a large number of agents and the possibility of abstentions, elections aggregate information effectively, nearly always obtaining the correct outcome (Feddersen and Pesendorfer, 1998). Third, information cascades do not occur if agents only care about the outcome of the election and not about

³We assume that if an agent is indifferent between standing and sitting after acquiring new information, then given inertia she remains in her current state.

their individual vote in a noiseless environment (Fey, forthcoming).

These models would seem to say that standing ovations occur only if they should occur, but this conclusion must be tempered. First, even though agents are rational, they are optimizing only in a limited sense. Here they only make one decision at one time, and since these decisions are a Nash Equilibrium, no one wants to change her decision. Nevertheless, in the context of the problems they are addressing, this limitation stifles much of the interesting behavior. For example, a jury makes deliberative decisions that allow for multi-period signaling (Coughlin, 1998) similar to the SOP. Jurors may at first vote not guilty and then, based upon the signals of others, switch and decide to vote guilty. The overlapping, interacting signals cannot be captured easily by a static Nash analysis. The ability to delay the switching of a vote to relate intensity of preference, a common feature of both juries and standing ovations, plays an important role in information aggregation that is absent from these models.

In addition to constraining signaling, these models do not include communication networks. Within a jury, all agents receive all signals, so the criticism doesn't apply. However, for elections with a large number of voters, assuming that agents' decisions are not correlated with those of the agents with whom they are in communication, ignores how information spreads. Information does not emanate from a giant sun; Instead, it passes through a vast information network consisting of newspapers, television, radio, friends, books, and public figures. Some of these sources may resemble the audience members in the front row, as they may have great visibility but little to say,⁴ while other sources may have much to say, but no visibility.⁵ Regardless of the specific topology of influence, information and behavior should be strongly correlated with respect to these connections.

Standard game theoretic models make an aggregative assumption about information and ignore the details for the sake of tractability. The art of modeling hinges on when to focus the microscope and when to misplace one's spectacles. The omission of communication and information networks, though often accepted within the social sciences, biases the investigation of social phenomena such as information aggregation. Part of the reason for ignoring communication networks has been a lack of analytical tools. Recent

⁴Celebrities come to mind.

⁵Academics come to mind.

advances in computational techniques, make such investigations a realistic possibility.

3.3 Pure Conformity

The standing ovation can also be interpreted as pure conformity and not as a strategic attempt to signal. In many situations, people modify their behavior to match their neighbors, peers, or friends (Bernheim, 1994). Evidence suggests that many other species also imitate the behavior of their neighbors (Galef, 1976). Standing ovations might be caused by this preference for conformity. Efficiency becomes irrelevant under pure conformity. If everyone prefers to act identically to everyone else, then either all standing or all sitting are efficient outcomes.

Creating a standing ovation model with preferences for pure conformity requires little work, and, not too surprisingly, eventually either agents all stand or all sit. However, the interesting part of these models is not the asymptotic behavior that is so well illuminated by the mathematical models, but rather the dynamics that take us to that state. As in life, the journey is often more interesting than the destination.

3.4 Growth and Coordination

Durlauf (1995) constructs models where agents can take one of two possible actions. These agents reside on a giant lattice and the actions of their neighbors influence their own choices. The distinction between these models and conformity models is subtle. In these models, the actions of others determine an aggregate variable that in turn influences the costs and benefits of the two actions. So, instead of choosing A because her neighbors chose A , an agent chooses A because her neighbors' actions make A less costly than B .

Using random and mean field theory, Durlauf (1995) shows that there can be up to three equilibria, of which two can be stable. Unlike the pure conformity case, the two stable equilibria can be Pareto ranked. Suppose we have two actions A and B , and attach greater personal satisfaction with A (for example, A might represent staying in school, choosing not to do drugs, or acquiring new technology, while B might denote dropping out, committing crimes (Glaser, *et al.*, 1994), or sticking to existing technology). As in conformity models, neighbors on the lattice influence payoffs. The

two equilibria can be interpreted as good or bad neighborhoods, schools, or growth rates. This class of models suggests that decentralized interaction need not lead to the preferred outcome as the agents can all coordinate on the wrong action.

3.5 Diffusion Models

A standing ovation might also represent the spread of a new product or new technology. Two assumptions can be made in the SOP so that it accords well with models of diffusion (see Lave and March, 1975). First, the decision to stand must be irreversible. Second, agents must become more likely to stand as more other agents stand. This second assumption makes sense if either the costs fall or the benefits rise as more people take the action. If there are network externalities, such as with the diffusion of the telephone, then it is the latter. If there are returns to scale in production, such as may be the case with washing machines, then it is the former. In either case, the more that do it, the more likely that others do it. This second assumption also makes sense if the action is diffused through interpersonal contact. If people can only get the flu from other people, then the more people that have the flu, the more likely that someone who doesn't have it will.

Let's assume that people stand to conform with others. In this case, in the SOP the benefits of standing increase with the number of people who stand. It is also true that standing spreads largely through interpersonal contact, though unlike in diffusion models of disease there is not random mixing. In the SOP, people either stay in or in front of their seats.

A main result from the diffusion models is that the number of people who take the action fits an S-shaped curve. At first only a few people take the action, but as more do, this creates incentives or causes more others to do so as well. Eventually, almost everyone is taking the action and the rate at which new people take the action falls.

4 Reactions to the SOP

We intentionally frame the SOP very loosely. When students first discuss the problem, they typically emphasize several features that they would like their models to embody. First, in deciding whether to stand at the end of

a performance, audience members should balance their desire to provide an honest signal of their enjoyment level of the performance against the pressure to conform to others. Such conformity may be strategic, informational, or purely preference based. The natural focus on agents modifying their behavior based on the actions of their neighbors makes the interdependence of the agents quite salient to the students, and forces them to directly confront one of the most fundamental issues in modeling social agents.

Students also tend to be very cognizant of the underlying complexity of the SOP. This complexity arises in information, expectations, and actions, and may make applying traditional solution concepts difficult at best. Students often struggle with employing either rational (RATs) or rule of thumb (ROTs) agents. Under ROTs, the models need not assume that all agents rely on the same behavioral rules. For example, some agents may care only about the two people on either side of them, while others may try to calculate (given their field of vision) the percentage of audience members standing. Students often endow agents with limited and diverse information. While most students assume that people sitting in the front get less information than people in the back, some students create models in which a portion of the audience chooses to turn around and scan the theater before deciding what to do. In either case, audience members obtain different information as a result of micro-level assumptions on the informational flow. Such micro-level flows may not be well approximated by the more typical assumptions of information emanating from a central source with an exogenous parametric distribution.

Finally, students struggle with the appropriate level of elaboration in the model. Students usually start from a skeletal model and over time ratchet up the level of realism. These later models may vary the types of people sitting in the various sections of the theater (for example, people who sit in the front might be predisposed toward enjoying the performance), rely on continuous signals like the decibel level and length of applause, or allow people that get up to leave or grab their coats to cancel out an emerging ovation.

SOP belongs to a rich class of problems—decentralized dynamical systems consisting of spatially distributed agents who respond to local information. Such models force students to contemplate the interplay between the micro-level rules of agents and the macro-level behavior of the system. On the one hand, subtle changes in agents' behavior can have unanticipated and large effects on final outcomes. On the other hand, some macro-level events prove

to be invariant to a large class of micro-level rules. Even in those cases where the end states do not differ (for example, if everyone in the first three rows stands then many models lead to the entire audience standing), the dynamics leading to those states often does.

5 Some SOP Results

We now formally describe a particular approach to modeling the SOP. This description resembles several models constructed by students. Agents are seated in a rectangular auditorium with R rows and C seats per row. At the conclusion of a performance, whether academic, musical, or comical (or all three), each agent makes an evaluation of the performance's quality. Let $q_{ij} \in Q_{ij} = [0, 1]$ represent the quality signal received by the audience member seated in the i th row and j th seat. Higher values of q_{ij} represent greater perceived quality. For the moment, values may be thought of as either private or common with idiosyncratic noise, or some convex combination of the two. Each audience member possesses an exogenous threshold level in addition to his or her quality evaluation. The threshold level for the agent in the i th row and j th seat, T_{ij} , equals the minimal quality required for that agent to stand immediately. Thus, if an audience member's private value weakly exceeds the threshold, that is if $q_{ij} \geq T_{ij}$, she stands immediately. If not, she remains seated.

Let $s_{ij}^t \in \{0, 1\}$ denote whether or not the audience member is standing ($s_{ij}^t = 1$) or sitting ($s_{ij}^t = 0$), t time periods after the completion of the performance, and let S^t be the total number of audience members standing at time t . Therefore,

$$S^0 = \sum_{j=1}^C \sum_{i=1}^R s_{ij}^0$$

equals the number of agents standing immediately. In interesting cases, only a fraction ($0 < S^0 < R \cdot C$) of the audience stands immediately. With a portion of the audience standing, those who remain seated must decide whether or not to stand, and those audience members standing must decide whether to remain standing or to sit. Both decisions rely on local information—possibly the number of neighbors standing or the percentage of audience members within sight who are standing—as well as the initial quality appraisal of the

individual. For example, someone seated surrounded entirely by people who are standing most likely will stand, unless she abhorred the performance. Similarly, unless she felt that the performance was stupendous, an isolated standing person may decide to sit if her neighbors do not join the standing ovation quickly.

The behavior of an audience member at a particular point in time can be represented by a heuristic which maps her information and quality appraisal into an action, either sit or stand. Periods are considered as discrete units. The continuous time case shall be addressed later in the discussion of random asynchronous and endogenous asynchronous updating. Let K_{ij} be the seat assignments visible to the audience member in the i th row and j th seat. Define a *behavioral rule* at time $t > 0$ by $F^t : K^t \times Q_{ij} \rightarrow \{0, 1\}$ for $t \geq 1$. Recall that in period 0, the decision whether to stand depends only on the agent's threshold and q_{ij} .

As formulated, the behavioral rules may depend on the time period. An audience member may require fewer neighbors to stand initially than in later periods in order to be induced to stand herself. Suppose, for example, that audience members can observe the actions of everyone seated in the row in front of them but are incapable of observing any other agents. Suppose that everyone in the front row stands immediately and that no other members of the audience do. Consider two scenarios. In the first, the behavioral rule of each agent is to stand only if all visible agents are standing. It follows that in the first period, the second row stands, in the second period the third row stands, and so on until the $R - 1$ period when the entire audience is standing. In the second scenario, suppose that an audience member stands only if all visible agents are standing prior to the end of the third period. With this rule, the first through fourth rows will stand, but then the standing will cease. These scenarios demonstrate the interplay between the micro-level rules of the audience members and the resulting macro-level phenomena which forms the core of this inquiry.

In the previous example, people in the front rows have more signaling power than people in the rear. While people in the front can be seen by nearly everyone, people in the rear cannot. If the entire front row of audience members were to stand at the conclusion of a performance, they make their preferences known to everyone in the audience. In contrast, if the people in the back row were to stand, their preferences might only be known by people in the one or two rows adjacent to theirs. The large influence of the front

rows becomes especially apparent when considering the seeding of standing ovations.

6 A Computational Model of Standing Ovations

We now construct a computational model that approximates the formal model described in the previous section. We then vary the utility functions of the audience members as well as the heuristics that they use to decide whether or not to stand. We begin with threshold-based behavior, and then add features as we proceed. In the first scenario, each audience member uses a majority rule heuristic—if a majority of the people that she sees are standing, she stands, if not she sits. Previously, we discussed a variety of issues that we can address within the SOP. Here we consider two of them: the timing of updating and the information structure.

Computational modeling allows great flexibility in the implementation of timing and process. To demonstrate how timing can be implemented, and, more importantly, how it can make a difference to the predictions of the model, we consider three possible procedures for updating: synchronous, asynchronous random and asynchronous incentive based. Under synchronous updating, all agents update simultaneously; Under random asynchronous updating, agents update one at a time based on a random order (capturing the spirit of continuous time updating) (Glance and Huberman, 1993); Under incentive based asynchronous updating the order is not random but depends upon incentives (Page 1997). Those people surrounded by people taking the opposite action are the first to update.

Synchronous Updating: *At the start of each discrete time period, all agents update in unison.*

Asynchronous Updating—Random Order: *Within each discrete time period, the agents are permuted into a random order and updated in that order. The $k + 1$ st agent to update can observe the behavior of the first k agents before updating.*

Asynchronous Updating—IncentiveBased: *Within each discrete time period, the agents update one at a time based on an explicit ordering rule that captures incentives.*

Such rules could depend on a variety of factors, for example, geography (those in the front move first), incentives (those who will gain the most utility move first), and so on. Here, we use an incentive-based rule—agents who have the most to gain by updating move first. In the SOP, this means agents who are least like the people that surround them react first. A person engulfed by a standing ovation stands faster than someone for whom the threshold has just been met.

The second issue concerns the information structure imposed by the neighborhoods. At one extreme, audience members might see the behavior of the entire audience. This corresponds to the assumption of global information. At the other extreme, a person may only see a limited distance, say, the three people immediately ahead and the two on either side. We shall consider two neighborhood structures. In each, X denotes the agent and each F denotes a visible neighbor.

Five Neighbors: *Agents look at the two neighbors on either side and the three agents directly ahead of their current location.*

$$\begin{array}{c} F \ F \ F \\ F \ X \ F \end{array}$$

Cones: *Agents look at the two neighbors on either side of them, the three agents in the row directly ahead, the five agents two rows ahead, and so on.*

$$\begin{array}{c} F \ F \ F \ F \ F \ F \ F \\ F \ F \ F \ F \ F \\ F \ F \ F \\ F \ X \ F \end{array}$$

The model proceeds as follows. We assume a square auditorium with 400 seats. Initially, audience members make their decisions based solely on perceived quality. We let audience members' initial quality assessment lie in the interval $[0, 1]$ and give each of them the same standing threshold of 0.5. Thus, an audience member initially stands if and only if her perceived quality exceeds 0.5. After the initial ovation, each agent decides what to do entirely on the basis of what other audience members are doing. An agent stands if and only if a majority of her neighbors are standing.⁶ Admittedly, we

⁶In the case of a tie, we assume that she sits or stands with equal probability.

assume an extreme transition in behavior between the initial and subsequent periods—it is likely that someone who stands initially, won’t immediately sit if barely less than a majority of her neighbors are standing. Nevertheless, this assumption is sufficient to generate some interesting results, and the symmetry induced by the use of identical rules for sitting and standing greatly simplifies the analysis.

We introduce three measures to compare outcomes under the different scenarios. *Number of Iterations* (NI) denotes the number of periods until a steady state is achieved. For incentive based updating, we let one period equal 400 agent decisions. *Stick in the Muds* (SM) equals the percentage of people that do the opposite of the majority in the steady state, and *Informational Efficiency* (IE) equals the percentage of the time that the majority of agents in the steady state takes the same action as the majority did initially. The higher IE the better the information aggregated. The fewer number of SM’s the better as well.

In the five neighbor scenario we find:

Five Neighbors

	Number of Iterations		Stick in the Muds		Informational Efficiency	
Async—Random	10.3	(0.62)	34.9	(1.01)	72.0	(4.51)
Sync	20.2	(1.45)	25.0	(1.91)	57.0	(4.98)
Async—Incentive	2.3	(0.05)	27.5	(1.57)	53.0	(5.02)

Notice that random, asynchronous updating leads to a higher IE and more SM’s, than either of the the other two updating rules. This suggests a tradeoff between the two, that aggregating information efficiently requires some stick in the muds. However, this is not universally true. Synchronous updating dominates Asynchronous Incentive-Based updating on both counts. It has a higher IE and fewer SM’s. It is disappointing that incentive-based asynchronous updating, probably the most realistic timing assumption, performs worst on both measures. One good thing about Asynchronous Incentive-Based updating is that it converges really quickly. The system evolves fast relative to the adaptation rate of the agents. This is likely the result of the growth of regions of SMs. By comparison Synchronous updating takes a very long time to settle down into a steady state. This is because members of the crowd can stand and sit many times while trying to coordinate.

Under cone neighborhoods we observe:

Cones

	Number of Iterations		Stick in the Muds		Informational Efficiency	
Async—Random	8.0	(0.92)	22.9	(0.94)	60.0	(4.92)
Sync	18.8	(1.12)	13.6	(1.61)	51.0	(5.02)
Async—Incentive	2.0	(0.00)	16.8	(1.45)	51.0	(5.02)

We find similar patterns to those seen under the five-neighbor scenario. Note, however, that even though the agents see more agents with cone neighborhoods, and thus they should have better information, the IE is lower. This occurs because the agents in the front have enormous influence. Almost everyone cues off of the behavior of the front row agents, and we find a phenomenon that is similar to an information cascade. Note as well that the number of SM drops considerably.

If we compare these computational findings with the mathematical literature, we see many of the same features (1) The system often converges to the “wrong” equilibrium, that is, most people can be standing even though most did not like the play. (2) Greater pressures to conform - as captured by the cones - leads to less efficient aggregation of information (3) A plot of the number of people standing over time tends to be roughly S-shaped as predicted by diffusion models and (4) People in the front can have a large effect. by this model.

That said, the mathematical results that we described tend to suggest starker outcomes than we see in the models. The mathematical models imply (via added noise) that all agents eventually take the same action. This rarely happens in the computations. Witness the number of stick in the muds. Also, with the exception of the cascade models which make a strick assumption, the mathematical models typically ignore the sequencing of updating. Yet we see how important that is. The mathematical models also obscure many of the interesting dynamics. Though we do see something that is roughly like an S-shaped curve, it is relatively easy to upset that finding so that the ovation is more gradual. And, even if the S-shaped curve is correct, it only characterizes the number of people standing, not their spatial locations. In sum, the mathematical insights do not reveal the full story, but certainly help us frame and understand the results emerging from the computational model. Similarly, the computational insights can now begin to inspire new directions in the mathematics.

7 Conclusions

The SOP offers a platform for considering worlds with social learning, diffusion, networks, and heterogeneity. It is also a platform for exploring important methodological issues related to computational modeling. We have found that our students find modeling the SOP both a fun and inspiring enterprise—one that is different enough from the usual problems they face to shake them from paradigmatic complacency, yet grounded enough to key social concerns to avoid apathy. By simultaneously being normatively benign yet closely related to many important social issues of the day (the effectiveness of democracies, the spread of crime, divergences in growth, and financial inefficiencies), the SOP pleases both the puzzle solvers and those concerned with the “real world.”

The SOP can be extended in many directions. As we mentioned one can add balconies, introduce variables that allow for continuous signaling, or even include people leaving the auditorium. Many of these variations have real-world analogs. In a recent workshop, we put a new spin on the SOP. We encouraged our students to consider the entrance dynamics that determine the initial seating decisions of agents. In such a world, local groups begin to suddenly play a big role. Ultimately, the two problems can be linked. The entrance dynamics may imply a very different pattern of seating than our usual assumptions may imply: people predisposed to like the presentation may sit up front, groups with similar tastes may cluster, and so on. Such initialization patterns may have a big impact on the ultimate dynamics of the ovation, or on the complexity of the agent’s strategies and information processing. Exactly what will happen, we cannot say at this point. We will have to wait until the performance ends.

References

- Arthur, W. Brian (1994), "Inductive Reasoning and Bounded Rationality: The El Farol Problem," *American Economic Review (Papers and Proceedings)* 84:406–11.
- Austen-Smith, David, and Jeffrey Banks (1996), "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *American Political Science Review* 90:34–46.
- Axtell, Robert, and Joshua Epstein (1999), "Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement," Brookings CSED Working Paper No. 1.
- Bak, Per (1996), *How Nature Works: The Science of Self Organized Criticality*, Springer Verlag: New York.
- Banerjee, Abhijit (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107:797–817.
- Bell, Ann (2002), "Dynamically Interdependent Preferences in a General Equilibrium Environment," *Journal of Economic Behavior and Organization* 47:309–33.
- Bernheim, B. Douglas (1994), "A Theory of Conformity," *Journal of Political Economy* 102:841–77.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A Theory of Fads, Fashions, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy* 100:992–1026.
- Coughlin, Peter (2000), "In Defense of Unanimous Jury Verdicts," *American Political Science Review* 94:375–94.
- Durlauf, Steven (1995), "Statistical Mechanics Approaches to Socioeconomic Behavior," Santa Fe Institute working paper 96-08-069.
- Durlauf, Steven (1997), "The Memberships Theory of Inequality: Ideas and Implications," Santa Fe Institute Working paper 87-05-047.
- Feddersen, Timothy, and Wolfgang Pesendorfer (1998), "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts," *The American Political Science Review* 92:23–35.
- Fey, Mark (forthcoming), "A Note on the Condorcet Jury Theorem with Supermajority Voting Rules," *Social Choice and Welfare*.
- Galef Jr., Bennett (1976), "Social Transmission of Acquired Behavior: A Discussion of Tradition and Social Learning in Vertebrates," *Advances in the Study of Behavior* 6:77–100.

- Glance, Natalie and Bernardo Huberman (1993), "Evolutionary Games and Computer Simulations," *Proc. Natl. Acad. Sciences (USA)* 90:7716–18.
- Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman (1996), "Crime, and Social Interactions," *Quarterly Journal of Economics* CXI:507–48.
- Granovetter, Mark (1978), "Threshold Models of Collective Behavior," *American Journal of Sociology* 83:1420–43.
- Holland, John (1986), "Escaping Brittleness: The Possibility of General-Purpose Learning Algorithms Applied to Rule-Based Systems," in *Machine Learning: An Artificial Intelligence Approach*, Volume II, (R. Michalski, J. Carbonell, and T. Mitchell, eds), pp. 593–623, San Mateo, CA: Morgan Kaufmann.
- Holland, John, and John H. Miller (1991), "Artificial Adaptive Agents in Economic Theory," *American Economic Review (Papers and Proceedings)* 81:365–70.
- Judd, Ken (1997), "Computational Economics and Economic Theory: Complements or Substitutes?," *Journal of Economic Dynamics and Control* 21:907–42.
- Kauffman, Stuart (1989), "Adaptation on Rugged Fitness Landscapes," in *Lectures in the Sciences of Complexity* (D. Stein, ed), pp. 527–618, Redwood City: Addison-Wesley.
- Kandori, Michihiro, George Mailath, and Rafael Rob (1992), "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* 61:29–56.
- Kollman, Ken, John H. Miller, and Scott E. Page (1992), "Adaptive Parties in Spatial Elections," *American Political Science Review* 86:929–37.
- Lave, Charles, and James March (1975), *An Introduction to Models in the Social Sciences*, New York: Harper and Row.
- Mayer, David, and Thad Brown (1998), "The Statistical Mechanics of Voting," *Physical Review Letters* 81:1718–21.
- Menczer, Filippo, and Troy Tassier (2001), "Emerging Small-World Referral Networks in Evolutionary Labor Markets," *IEEE Transactions on Evolutionary Computation (Special Issue on Computational Economics)* 5:482–92.
- Miller, John H. (1998), "Active Nonlinear Tests (ANTs) of Complex Simulations Models," *Management Science* 44:820–30.
- Page, Scott E (1997) "On Incentives and Updating in Agent Based Models." *Computational Economics*. 10 pp 67–87.

- Page, Scott E. (1999), "Computational Models from A to Z," *Complexity* 5:35–41.
- Picker, Randal (1997), "Simple Games in a Complex World: A Generative Approach to the Adoption of Norms," *University of Chicago Law Review* 64:1225–88.
- Schelling, Thomas (1978), *Micromotives and Macrobehavior*, New York: W.W. Norton.
- Tesfatsion, Leigh (1997), "How Economists Can Get A-Life," in *The Economy as a Complex Evolving System II* (W. Brian Arthur, Steven Durlauf, and David Lane eds), pp. 533–65, Reading, MA: Addison Wesley.