

Coordination Failure as a Source of Congestion in Information Networks

Ann M. Bell, William A. Sethares, *Member, IEEE*, and James A. Bucklew, *Senior Member, IEEE*

Abstract—Coordination failure, or agents' uncertainty about the action of other agents, may be an important source of congestion in large decentralized systems. The *El Farol* problem provides a simple paradigm for congestion and coordination problems that may arise with over utilization of the Internet. This paper reviews the *El Farol* problem and surveys previous approaches, which typically involve complex deterministic learning algorithms that exhibit chaotic-like trajectories. This paper recasts the problem in a stochastic framework and derives a simple adaptive strategy that has intriguing optimization properties; a large collection of decentralized decision makers, each acting in their own best interests and with limited knowledge, converge to a solution that (optimally) solves a complex congestion and social coordination problem. A variation in which agents are allowed access to full information is not nearly as successful. The algorithm, which can be viewed as a kind of habit formation, is analyzed using a weak convergence approach, and simulations illustrate the major results.

Index Terms—Adaptation, decentralized networks, *El Farol*, habit formation, learning, minority game, weak convergence.

I. INTRODUCTION

STANDARD models of congested public resources focus on the costs that an individual user imposes on other potential users. For example, each person who travels on a congested highway or visits a popular web site increases the waiting time of subsequent users. Congestion arises because individuals do not consider the effects of their actions on other users. Explicitly charging users for these unobserved costs can eliminate the socially inefficient congestion of a scarce, shared resource. However, this approach often utilizes equilibrium solutions in which all agents are fully informed about the structure of the problem and the behavior of other agents. Consequently, the relationship between agents' behavior and the congestion they experience is easily discerned. This reliance on information-intensive equilibrium solutions limits the usefulness of these models in solving resource allocation problems in large-scale systems such as the Internet. In contrast, this paper focuses on imperfect information and coordination failure across agents as a source of congestion in large decentralized systems. We utilize a coordination problem or simple congestion game, framed by Arthur¹ [3], as a simplified

model simplified model of a large class of congestion and coordination problems that arise in modern engineering and economic systems. *El Farol* is a bar in Santa Fe, NM.² The bar is popular but becomes overcrowded when more than 60 people attend on any given evening. Everyone enjoys themselves when fewer than 60 people go, but no one has a good time when the bar is overcrowded.

How should an agent decide whether or not to go out to the bar, given that the actions of other agents are unknown? The problem set up emphasizes the difficulty of coordinating the actions of independent agents without a centralized mechanism. The analogy between the *El Farol* problem and decentralized resource allocation is noted by Greenwald *et al.* [10], as well as in our previous work [4], [22]. Glance and Huberman [9] and Huberman and Lukose [13] also consider the dynamics of congestion on the Internet when externalities similar to those found with public goods prevail. Unlike the standard public good framework, in the *El Farol* scenario, fully informed optimizing agents will not increase consumption of a publicly available resource until it experiences an inefficient level of congestion: If agents could predict the behavior of other agents perfectly, the bar would never be crowded, and all patrons would have a good time.³ The only source of congestion, at least in a deterministic framework, is the inability of agents to coordinate their actions.

Arthur originally posed the *El Farol* problem to illustrate the aggregate dynamics of a system composed of bounded rational agents who rely on inductive learning. Agents attempt to predict the aggregate behavior of other agents, which simultaneously depends on all agents' predictions. Consequently, the interaction between individual learning strategies and the environment that agents face plays a key role in determining the dynamics of the system. Using *El Farol* to model the Internet environment emphasizes that congestion can arise from coordination failure across agents, as well as from absolute constraints on bandwidth. Furthermore, in contrast to many game theoretic treatments of learning and coordination, the level of congestion at *El Farol* depends on the actions of a relatively large number of individual agents. These features make it an especially useful tool for analyzing information technology systems characterized by decentralized decision making and rapid endogenous changes in the operating environment.

A. Overview

In our previous treatments [4], [22], we proposed a deterministic adaptive algorithm based on habit formation that enabled

²Arthur's *El Farol* scenario is also known as the Santa Fe bar problem.

³The stochastic or mixed-strategy framework may suffer from socially inefficient congestion as discussed below.

Manuscript received December 5, 2000; revised October 28, 2002. The associate editor coordinating the review of this paper and approving it for publication was Dr. Gonzalo Arce.

A. M. Bell is with the NASA Ames Research Center, Moffett Field, CA 94035-1000 USA (e-mail: abell@mail.arc.nasa.gov).

W. A. Sethares and J. A. Bucklew are with the Department of Electrical and Computer Engineering, University of Wisconsin-Madison, Madison, WI 53706 USA.

Digital Object Identifier 10.1109/TSP.2002.808139

¹The market entry game analyzed by [18] and [23] has a similar structure as does the "minority game" analyzed by Challet and Zhang [7] and Savit *et al.* [20].

agents to solve the *El Farol* coordination problem in a decentralized environment while avoiding the seemingly random fluctuations in aggregate attendance that Arthur's simulations demonstrated. Here, we consider the *El Farol* problem in a stochastic setting. We analyze a stochastic adaptive algorithm analogous to the one introduced in [4] and [22] and consider the dynamic and equilibrium characteristics of the system in relation to the mixed and pure strategy equilibria of the corresponding game. We employ some novel convergence results [5] that approximate the dynamics of the (stochastic) system using (deterministic) ordinary differential equations (ODEs) and allow a concrete description of its convergence and stability properties. In addition, we demonstrate that the information structure plays a crucial role in determining the behavior of the system.

B. How do Agents Decide to Attend *El Farol*?

In Arthur's simulations, agents attempt to predict how many others will attend *El Farol* each time using a simple kind of deterministic inductive reasoning. If they predict attendance will be less than 60, then they go to the bar; if they predict attendance will be greater than 60, then they stay at home. Each agent uses a number of "rules of thumb" such as simple averages, moving averages, and linear or nonlinear filters to formulate predictions and then acts on the prediction that was correct most frequently in the recent past. When Arthur simulated a bar-going society of 100 inductively rational agents, he found that the population at the bar tended to hover near 60 though attendance varied greatly, often exceeding 70 or dropping below 50. The time series of aggregate attendance appeared random, despite the deterministic rules of the underlying agents.

Here, we consider the *El Farol* problem in the stochastic setting. We briefly discuss the characteristics of pure and mixed strategy equilibria of the corresponding congestion game and then frame our adaptive learning rule in terms of a mixed strategy profile. There are several advantages to considering the stochastic version of the adaptive learning rule: a clearer problem statement, a simpler algorithm that is amenable to detailed analysis, and more general results. The analysis demonstrates that the type and characteristics of the equilibria actually observed depends crucially on the nature of the information available to agents. In particular, we show that limiting the information available to agents leads them to successfully coordinate on a Pareto efficient equilibrium while providing more information leads to an inefficient outcome. Our results emphasize the critical role that information exchange plays in alleviating congestion that arise from coordination failure.

A somewhat unusual feature of the *El Farol* problem statement is the discontinuous transition from uncrowded to crowded that occurs when the 61st patron arrives. While this may seem like an unrealistic assumption for a bar, discontinuities and extreme nonlinearities are prevalent in information technology applications. For example, when a network server divides resources equally among users, the performance of the entire system can dramatically decrease with the addition of a single user. Many routers handling data packets have fixed queue lengths: Additional packets are dropped. When data from two sources arrive simultaneously, exceeding queue capacity, packets from both users may be dropped, leading to

long delays for both messages. The quality of audio and video data streams degrades rapidly when packets are dropped. In general, systems that experience congestion at a bottleneck will respond nonlinearly when traffic increases even slightly above the capacity of the bottleneck. The preference structure of the *El Farol* problem mimics the discontinuous and nonlinear responses to increases in traffic observed in information systems.

In addition, the discontinuity in agents' response to attendance levels helps distinguish between congestion arising from overuse of a public good and congestion arising from coordination failure. When the value of attendance declines slowly in response to larger turnouts, agents will continue to attend until the value of attendance for all bar goers has been reduced to the value of staying at home. Congestion in this case may be optimal for the individual but nonetheless inefficient for society: Everyone could be made better off by a compensation scheme that induces some agents to stay home. The discontinuous preferences utilized in the *El Farol* framework help minimize the importance of individually optimal but socially inefficient congestion.

C. Other Approaches to *El Farol*

The *El Farol* problem has received a fair amount of attention from computer scientists and physicists as well as from researchers in the area of complex systems. Casti uses the *El Farol* problem to frame his definition of a complex adaptive system as one with "a medium-sized number of intelligent, adaptive agents interacting on the basis of local information" [6, p. 10]. The dynamics of Arthur's system are entirely deterministic (only the initial values of agents parameters are chosen randomly) the resulting pattern of attendance appears random. The uncertainty or apparent randomness in the system is entirely endogenous, created by the interaction between the number of agents attending the bar and the set of prediction rules active at any given time.

Johnson *et al.* [14] consider how the variance in the *El Farol* problem changes in response to the number of predictors available in the entire system and the number of predictors that each agent selects. Zambrano [24] applies results from Bayesian game theory to show that a system composed of Bayesian learners will converge to the set of Nash equilibria. Greenwald *et al.* [10] examine whether or not boundedly rational agents can learn their way to a mixed strategy equilibrium. Note that agents in their model are not able to distinguish the effects of their own actions on aggregate attendance, which we demonstrate is a critical factor in determining system behavior. Challet and Zhang [7] simplify the *El Farol* problem even further by considering a "minority game" in which agents choose one of two groups to join and receive positive payoffs when they choose the smaller group. The information available to agents is limited even further: They only observe which group was the minority and not the number of agents who chose that group.

II. *EL FAROL* AS A GAME

The *El Farol* problem is a type of congestion game, first characterized by Rosenthal [19]. In congestion games, each agent chooses a resource to utilize. The agent's utility depends on the

number of other agents who choose to utilize the same resource. Finding a Nash equilibrium of a congestion game is equivalent to a constrained minimization problem.

We consider the *El Farol* problem as a one-shot simultaneous move game. Let agents have identical payoffs: b is the payoff an agent receives for attending a crowded bar, and g is the payoff an agent receives for attending an uncrowded bar. Without loss of generality, let h , which is the payoff received for staying home, be zero. Let M be the total number of agents and \mathcal{N} be the maximum capacity of an uncrowded bar.

In a deterministic setting where agents utilize only pure (deterministic) strategies, a Nash equilibrium occurs when exactly sixty agents choose to attend. There are $\binom{100}{60}$ such equilibria. There are no symmetric pure strategy Nash equilibria. Pure strategy Nash equilibria are Pareto efficient.

Arthur's approach sidesteps the usual game theoretic considerations by focusing on the process of prediction in an endogenously evolving environment rather than on the binary choice between the strategies of attending and staying home. The only information available to agents is attendance in each time period. It is often reasonable to assume that agents do not and need not remove themselves from the aggregate statistics before reacting to them. However, because the *El Farol* problem contains a knife-edge response to increased attendance, the analysis of equilibria depends crucially on how the agent accounts for his or her own behavior.

Suppose that agents use predictive rules like those suggested by Arthur and that attendance at *El Farol* for the last ten periods has been exactly 60. How should an individual agent decide whether or not to attend in this case? Common sense suggests that agents who have attended the bar every period should continue to attend every period. On the other hand, agents who have not attended at all in the last ten periods should remain at home because the addition of another agent will result in attendance of 61. The key issue is agents' ability to account for their own past behavior. The oft-repeated conjecture about the *El Farol* problem (that "no shared, or common, forecast can possibly be an accurate one; deductive logic fails" [6]) depends crucially on the assumption that agents cannot recognize their own attendance pattern in the aggregate.

A formal treatment of the knife edge case when attendance exactly equals 60 would alter the predictive rules to account for the agent's own behavior: Agents should attend if they predict 59 or fewer agents *other than themselves* will attend and stay home if they predict 60 or more agents *other than themselves* will attend. In this scenario, Arthur's formulation of the *El Farol* problem has well-defined steady states in which all agents can utilize the same successful predictive rule. The heterogeneity in agents' actions arises from the heterogeneity in information: Each agent's information set is unique because only the agent knows whether or not they were among the bar attendees at any point in time. When agents do not account for their own behavior, they must draw different conclusions from the same data set in order to produce average attendance of 60.

Moving to a stochastic framework that allows mixed-strategy equilibria requires explicit payoffs for the different outcomes. Each agents' mixed-strategy profile consists of a single parameter p_i , which indicates the probability that agent i attends. Let

M be the total number of agents, N be the total observed attendance, N^{-i} be the observed attendance exclusive of agent i , and \mathcal{N} be the maximum capacity of an uncrowded bar.

A mixed-strategy equilibrium must satisfy the condition

$$g \Pr(N^{-i} \leq \mathcal{N} - 1) + b \Pr(N^{-i} > \mathcal{N} - 1) = 0 \quad (1)$$

or

$$\Pr(N^{-i} \leq \mathcal{N} - 1) = \frac{b}{b - g}$$

which states that the expected return to the pure strategy of attending the bar exactly equals the expected return to the pure strategy of staying home. This must hold for all agents simultaneously. In addition, note that the indifference condition that determines a mixed strategy equilibrium depends on the distribution of total attendance, which, in general, depends on the probabilities for individual agents and not just on the mean of the entire distribution.

For a symmetric mixed strategy equilibrium, the probability that $\mathcal{N} - 1$ or fewer agents attend is

$$\sum_{N^{-i}=0}^{N^{-i}=\mathcal{N}-1} \binom{\mathcal{N}-1}{N^{-i}} \left(p^{N^{-i}} (1-p)^{\mathcal{N}-1-N^{-i}} \right). \quad (2)$$

When the symmetric mixed-strategy equilibrium is $p = 0.6$, then $g \approx -0.98b$.

The symmetric mixed strategy Nash equilibrium is not Pareto optimal because agents increase their probability of attending until the expected return to attendance exactly equals that of staying home: 0. In addition, the randomness in agents' choice of strategy will generate inefficient variance in attendance. Any attendance outcome that falls short of the maximum capacity of an uncrowded bar can be improved by increasing attendance, and vice versa. The Pareto optimal symmetric mixed-strategy profile⁴ can be calculated by

$$\max_p \sum_{N=0}^{N=\mathcal{N}} g N \Pr(N) + \sum_{N=\mathcal{N}+1}^{N=M} b N \Pr(N). \quad (3)$$

This p maximizes the total expected payoff to all agents, which also maximizes the expected return to individual agents. For example, when $g = 1$ and $b \approx -0.98$, the Pareto efficient symmetric mixed-strategy profile is $p \approx .5$, and the expected payoff to an individual agent is ≈ 0.48 . In contrast, the symmetric mixed-strategy Nash equilibrium is $p \approx 0.6$, and the expected payoff to an individual agent is 0. In this sense, the *El Farol* problem suffers from inefficient congestion similar to that observed in a standard public goods framework in a stochastic framework: In the mixed strategy (stochastic) Nash equilibrium, each individual agent's probability of attendance is just high enough that the expected return is 0.

There are no asymmetric mixed strategy equilibria. Consider two agents with differing probabilities of attendance and, without loss of generality, label them agents 1 and 2 with $p_1 < p_2$. The indifference condition (1) must hold for

⁴the mixed-strategy profile that maximizes the expected return to each agent given the constraint that the expected return be equal for all agents.

every agent, which implies that $\Pr(N^{-1} \leq \mathcal{N} - 1)$ equal $\Pr(N^{-2} \leq \mathcal{N} - 1)$. The density function for attendance exclusive of agent 1 can be expressed in terms of the density function for attendance exclusive of agents 1 and 2:

$$\begin{aligned}\Pr(N^{-1} = 0) &= \Pr(N^{-1,-2} = 0)(1 - p_2) \\ \Pr(N^{-1} = x) &= \Pr(N^{-1,-2} = x)(1 - p_2) \\ &\quad + \Pr(N^{-1,-2} = x - 1)p_2.\end{aligned}$$

By expanding and combining sums, the cumulative distribution that agent 1 faces can be expressed as

$$\begin{aligned}\Pr(N^{-1} \leq X) \\ = \sum_{x=0}^{x=X-1} \Pr(N^{-1,-2} = x) + \Pr(N^{-1,-2} = X)(1 - p_2).\end{aligned}$$

The cumulative distribution function that agent 2 faces differs only by the term $(1 - p_2)$, which is replaced by $(1 - p_1)$. Consequently, the indifference condition cannot hold simultaneously for two agents with different probabilities.

III. LEARNING RULE FOR MIXED-STRATEGIES

Arthur's inductive learning approach requires agents to explicitly predict how many others will attend. A mixed strategy Nash equilibrium requires knowledge of the entire distribution of attendance. Our boundedly rational adaptive learning rule does not rely on prediction of or inference about other agents' behavior; rather, agents adapt their probability of attending over time based on the history of their own experiences at *El Farol*.

It is tautological that people prefer to experience good times rather than bad, to repeat the enjoyable, and to minimize the unpleasant. Although the *El Farol* situation provides a simple setting in which good and bad are clearly defined, it is not possible to know in advance whether a trip to the bar will be good or bad since this depends on the actions of everyone else. Suppose that the agent initially attends p percent of the time. Consistent with the desire to maximize pleasure and minimize painful experiences, the agent goes more often (increases p slightly) if the bar is uncrowded but prefers to go less often (to decrease p) if the bar is crowded. Over time, the agent gathers information about the state of the bar and "remembers" this in the form of the parameter p . This learning rule can be interpreted as a kind of habit formation or stimulus-response and is directly analogous to certain adaptive algorithms from the signal processing

literature [12]. However, the current situation is more complicated since the "true" value of the unknown depends explicitly on everyone's behavior.

Suppose M agents compete for the \mathcal{N} spaces at *El Farol*. The probability that the i th agent attends is p_i . Let k be a time (iteration) counter and $N(k)$ be the number of agents attending at time k . Let μ be a characteristic parameter that defines how much each agent changes p_i in response to new information, and let $p_i(k)$ designate the instantaneous value of p_i at the time k . Let

$$N(k) = \sum_{i=1}^M x_i(k) \quad (4)$$

where the $x_i(k)$ are independent Bernoulli random variables that are 1 with probability $p_i(k)$ and zero otherwise. The evolution of the $p_i(k)$ is then defined by (5), shown at the bottom of the page.

The operation of the algorithm is uncomplicated. At each time k , the agent flips a biased coin, attending with probability $p_i(k)$. When the agent attends, then the parameter $p_i(k)$ is adjusted, increasing it proportionally to $N(k) - \mathcal{N}$ if the bar is uncrowded and decreasing it proportionally to $N(k) - \mathcal{N}$ if the bar is crowded. Since the $p_i(k)$ represent probabilities, they must be constrained to lie within 0 and 1. When the agent does not attend, $x_i(k)$ is zero, and $p_i(k+1) = p_i(k)$. Note that the stepsize does not decrease over time. The simplicity of the scheme makes it feasible to analyze the resulting behavior, as demonstrated in Section V.

In Arthur's formulation of the problem, agents have access to information about attendance at the bar even on evenings when they do not themselves attend. This can be incorporated into the algorithm (5), giving the update equation (6), shown at the bottom of the page, which mimics the information structure used by Arthur's agents. As will become clear, this information structure is a key element in the behavior of the algorithm. When agents base their updates on only their own experiences as in (5), they utilize "partial information." In contrast, (6) utilizes "full information" because agents base their decisions on the full record of attendance.

A related version of the stochastic algorithm updates according to whether the bar is crowded or not, as in (7), shown at the bottom of the next page.

None of these updating rules rely explicitly on payoffs. The reliance of the update rules on attendance rather than payoffs is

$$p_i(k+1) = \begin{cases} 0, & \text{if } p_i(k) - \mu(N(k) - \mathcal{N}) x_i(k) < 0 \\ 1, & \text{if } p_i(k) - \mu(N(k) - \mathcal{N}) x_i(k) > 1 \\ p_i(k) - \mu(N(k) - \mathcal{N}) x_i(k), & \text{otherwise.} \end{cases} \quad (5)$$

$$p_i(k+1) = \begin{cases} 0, & \text{if } p_i(k) - \mu(N(k) - \mathcal{N}) < 0 \\ 1, & \text{if } p_i(k) - \mu(N(k) - \mathcal{N}) > 1 \\ p_i(k) - \mu(N(k) - \mathcal{N}), & \text{otherwise} \end{cases} \quad (6)$$

not a crucial feature of our learning rule. An alternative specification of the game in which the payoff for attending is $(N(k) - \mathcal{N})$ and the payoff for staying home is 0 would make the first partial information algorithm (5) depend on payoffs. A game in which the payoffs were μ for attending an uncrowded bar, $-\mu$ for attending a crowded bar, and 0 for staying home would make the “signed” algorithm depend on payoffs. Asymmetric payoffs would merely entail different stepsizes for the crowded and uncrowded outcomes. The speed of convergence of a system with updates that depend on payoffs may depend on the magnitude (or units of measurement) of the payoffs.

Note that when the payoff to staying home is 0, a payoff-dependent updating rule as described above will correspond to partial information. Some treatments of the *El Farol* problem and of related problems like the minority game [7], [20] utilize a peculiar form of full information updating: Agents update according to which strategy performed best regardless of their own actions. So when 59 agents attend, for example, all 41 agents who remained at home assume that the strategy of attending would have had a higher payoff. Of course, had they all actually attended, the bar would have been crowded. The strategy updating in the minority game suffers from the same illusion: All agents who were in the majority assume that they would have been in the minority if they had chosen the other group. These updating schemes implicitly rely on agents’ ignorance of the fallacy of composition: Not everyone can attend an uncrowded bar, nor can everyone be in the minority.

IV. GENERIC BEHAVIOR OF THE ALGORITHMS

This section explores the generic behavior of the system when each of the $M = 100$ agents follows the strategy defined by (5). Although details of the various simulations differ, a typical case is illustrated in Fig. 1. The probabilities $p_i(0)$ were initialized randomly.

Perhaps the most striking aspect of these simulations is the rapid convergence to near the optimal value of $\mathcal{N} = 60$ and the associated decline in the variance of attendance. The outcome approaches that which would be chosen with centralized control, despite the fact that each agent is autonomous and makes the decision to go (or not to go) based on local information, that is, on its own experiences.

Fig. 2 shows values of the probabilities $p_i(k)$ over the course of a typical simulation run. By the final iteration, the agents have divided themselves into two groups. The probability parameter for 60 of the agents has risen very near 1, indicating that they attend nearly every time. The remaining 40 agents attend less and less frequently, with their probability parameter very near zero. This division of the population appears nowhere in the algorithm statement; rather, it is an emergent property of the adaptive solution to the *El Farol* problem. Despite the stochastic




Fig. 1. When all agents use the “partial information” adaptive solution, the number of attendees appears to converge rapidly and then only rarely exceeds the critical $\mathcal{N} = 60$.




Fig. 2. Probability parameters for each of the $M = 100$ agents in Fig. 1 as a function of iteration number (time). An emergent property of the adaptive solution is that the population divides itself into “regulars” and “casuals.”

nature of agents of the adaptive learning rule, it converges to a pure strategy Nash equilibrium.

In contrast, Fig. 3 uses the “full information” algorithm (6) to investigate the effect of allowing the agents to update their probabilities at every iteration, whether they have personally attended the bar or not. This reflects the information structure in Arthur’s simulations. Mean attendance is approximately 60, but the variance does not decline over time, indicating that seats in the bar often remain unfilled, and the bar is often overcrowded. Note that the transient behavior in the initial periods is masked by the long time scale. Fig. 4 should be compared with Fig. 2; the probability parameters for these agents continue to bounce randomly about some fixed value as their probabilities all increase or decrease simultaneously in response to the same signals.

Somewhat paradoxically, agents successfully coordinate their behavior and the system achieves a Pareto efficient outcome only when agents have access to less information. Several authors have noted a similar phenomena in transportation routing. Mahmassani and Jayakrishnan [17] use

$$p_i(k+1) = \begin{cases} 0, & \text{if } p_i(k) - \mu \operatorname{sgn}(N(k) - \mathcal{N}) x_i(k) < 0 \\ 1, & \text{if } p_i(k) - \mu \operatorname{sgn}(N(k) - \mathcal{N}) x_i(k) > 1 \\ p_i(k) - \mu \operatorname{sgn}(N(k) - \mathcal{N}) x_i(k), & \text{otherwise.} \end{cases} \quad (7)$$




Fig. 3. When all agents use the “full information” adaptive solution, the number of attendees fluctuates wildly about the optimal $\mathcal{N} = 60$.




Fig. 4. Probability parameters for 15 of the $M = 100$ agents in Fig. 3 (the full information algorithm) as a function of iteration number (time). Others exhibit similar behavior. In contrast to the partial information case, parameters remain diverse.

simulations to demonstrate that when individuals pursue a strict best-response strategy, changing their route, no matter how small the improvement over their current choice, the performance of the system as a whole degrades if more than 25% of drivers have access to real time information about congestion. Arnott *et al.* [1], [2] show that congestion can arise because of “concentration,” or similar responses to common information, and that consequently, more information can lead to increased congestion.

V. ANALYSIS OF THE ADAPTIVE SOLUTIONS

This section analyzes the steady states and convergence behavior of the proposed algorithms by comparing and contrasting the different algorithm forms (5) and (6). The first subsection describes the various possible steady states. The simulations in Figs. 1 and 3 suggest that these are quite different; the analysis describes these differences in a concrete way. Section V-C reviews the relevant technical background on weak convergence and states the theorem that will be used in Section V-D to describe the convergence and stability behavior of the algorithms about their steady states.

One simple way to understand the asymptotic behavior of the algorithms is to observe that each p_i evolves on a finite state space: a lattice with steps of size μ (because the updates are

always an integer $N(k) - \mathcal{N}$ times the stepsize). For algorithm (5), the zero state is absorbing (since $x_i = 0$ is guaranteed once $p_i = 0$). Since $p_i = 0$ is reachable, the algorithm can be viewed as a finite-state Markov chain with reachable absorbing states and, hence, must converge. In contrast, for algorithm (6), $p_i = 0$ is not an absorbing state, and no convergence (to zero) can be expected.

A. Steady States of the Algorithms

The first step in the analysis of the dynamic behavior of the algorithms is to determine the conditions under which the means of the $p(k)$ remain fixed, that is, to determine the steady states of the averaged system.

1) *Algorithm With Partial Information:* Taking the expectation of both sides of (5) gives

$$E\{p_i(k+1)\} = E\{p_i(k)\} - \mu E\{(N(k) - \mathcal{N}) x_i(k)\}$$

assuming that the $p_i(k)$ are not at the boundary points 0 or 1. This expectation remains unchanged exactly when the update portion is zero, that is, when

$$E\{(N(k) - \mathcal{N}) x_i(k)\} = 0.$$

Using (4) and the fact that $E\{x_i^2(k)\} = E\{x_i(k)\} = p_i$ (which follows directly from the definition of $x_i(k)$ as a Bernoulli 0–1 random variable), this can be rewritten

$$\begin{aligned} E \left\{ \left(\sum_{j=1}^M x_j(k) - \mathcal{N} \right) x_i(k) \right\} \\ = E \left\{ \left(\sum_{\substack{j=1 \\ j \neq i}}^M x_j(k) + 1 - \mathcal{N} \right) p_i(k) \right\}. \end{aligned}$$

Because the term in parenthesis is independent of $p_i(k)$, this becomes

$$\begin{aligned} &= \left(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq i}}^M E\{x_j(k)\} \right) p_i(k) \\ &= \left(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq i}}^M p_j(k) \right) p_i(k). \end{aligned} \quad (8)$$

Consider any candidate steady state p^* with \mathcal{N} ones and $M - \mathcal{N}$ zeroes. Let I_1 be the indices of the ones and I_0 be the indices of the zeroes. Then, there are two kinds of terms in (8). When $i \in I_1$, $\sum_{\substack{j=1 \\ j \neq i}}^M p_j^* = \mathcal{N} - 1$; therefore

$$\left(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq i}}^M p_j^* \right) p_i^* = (1 - \mathcal{N} + \mathcal{N} - 1) p_i^* = 0. \quad (9)$$

When $i \in I_0$, $\sum_{j=1, j \neq i}^M p_j^* = \mathcal{N}$, $p_i^* = 0$, and hence

$$\left(1 - \mathcal{N} + \sum_{j=1, j \neq i}^M p_j^*\right) p_i^* = (1 - \mathcal{N} + \mathcal{N}) 0 = 0.$$

Hence, p^* is a steady state.

Now, consider any p^* for which $\sum_{j=1}^M p_j^* = \mathcal{N}$ that is not of the form of \mathcal{N} ones and $M - \mathcal{N}$ zeroes. Thus, $0 < p_n^* < 1$ for at least one n . In this case, the relevant term in (8) is

$$\begin{aligned} \left(1 - \mathcal{N} + \sum_{j=1, j \neq n}^M p_j^*\right) p_n^* &= \left(1 - \mathcal{N} + \sum_{j=1}^M p_j^* - p_n^*\right) p_n^* \\ &= (1 - \mathcal{N} + \mathcal{N} - p_n^*) p_n^* = (1 - p_n^*) p_n^*. \end{aligned}$$

This cannot be zero, and hence, p^* is not a steady state. Hence, the only steady states of algorithm (5) are at p^* , consisting of \mathcal{N} ones and $M - \mathcal{N}$ zeroes. In particular, the symmetric mixed strategy Nash equilibrium at $p_j^* = .6$ for all j (for $g = 1$, $b \approx -0.98$) is not a steady state of this algorithm.

2) *Algorithms With Full Information:* In contrast, consider a similar analysis carried out for the ‘‘full information’’ algorithm. Taking the expectation of both sides of (6) gives

$$E\{p_i(k+1)\} = E\{p_i(k)\} - \mu E\left\{\sum_{j=1}^M x_j(k) - \mathcal{N}\right\}.$$

Steady states occur when $E\left\{\sum_{j=1}^M x_j(k) - \mathcal{N}\right\} = 0$, i.e., whenever

$$E\left\{\sum_{j=1}^M x_j(k)\right\} = \sum_{j=1}^M E\{x_j(k)\} = \sum_{j=1}^M p_j(k) = \mathcal{N}.$$

Hence, any p^* with $\sum_{j=1}^M p_j^* = \mathcal{N}$ is a steady state of this algorithm. Note that these are not mixed strategy equilibria of the *El Farol* game unless $p_i = .6$ for every agent.

B. Derivation of the Algorithms From a Global Cost Function

To further understand the global behavior of the system, we relate the algorithms utilized by individuals to a global cost function. The algorithm can be derived as an approximation to an instantaneous gradient descent for minimization of the cost function

$$J(k) = \frac{1}{2}(E\{N(k)\} - \mathcal{N})^2 \quad (10)$$

where

$$E\{N(k)\} = E\left\{\sum_{i=1}^M x_i(k)\right\} = \sum_{i=1}^M E\{x_i(k)\} = \sum_{i=1}^M p_i(k) \quad (11)$$

is the expected number of attendees at time k . The typical gradient strategy is to update the state using

$$p_i(k+1) = p_i(k) - \mu(k) \frac{dJ(k)}{dp_i(k)}. \quad (12)$$

With $J(k)$ as in (10)

$$\frac{dJ(k)}{dp_i(k)} = (E\{N(k)\} - \mathcal{N}) \frac{dE\{N(k)\}}{dp_i(k)}.$$

From (11), the derivative is $dE\{N(k)\}/dp_i(k) = 1$, and hence

$$\frac{dJ(k)}{dp_i(k)} = E\{N(k)\} - \mathcal{N}.$$

Replacing $E\{N(k)\}$ by its instantaneous value gives

$$\frac{dJ(k)}{dp_i(k)} \approx N(k) - \mathcal{N}$$

which is an instantaneous approximation to the gradient of $J(k)$. Substituting this into (12) gives

$$p_i(k+1) = p_i(k) - \mu(N(k) - \mathcal{N}). \quad (13)$$

In the limited information case, this update occurs only when $x_i(k) = 1$, and in the full information case, this update occurs every iteration, regardless of the agent’s attendance. Adding the *a priori* limits on $p_i(k)$ then gives the algorithms (5) and (6). For both algorithms, $E\{N(k)\} = \mathcal{N}$ in a steady state. However, because the limited information algorithm converges to a pure strategy equilibria, the actually observed costs will be 0, whereas with full information, the expected costs will be $(1/2)\text{Var}[N(k)]$.

Similarly, algorithm (7) based on the sign of $(N(k) - \mathcal{N})$ can be derived from the absolute value cost function $J(k) = |E\{N(k)\} - \mathcal{N}|$. By analogy, these algorithms are variants of the least mean square (LMS) algorithms, which are common in the context of linear system identification and adaptive filtering [21]. Algorithm (7) is an analog of the signed LMS algorithm [15].

C. Weak Convergence

The convergence of the algorithms to these steady states can be examined by looking at the stability properties of a related ODE. This requires considerably more technical machinery, which is reviewed here. The basis of the analytical approach is to find an ODE that accurately mimics the behavior of the algorithm for small values of μ . Studying the ODE then gives valuable information regarding the behavior of the algorithm. For example, if the ODE is stable, then the algorithm is convergent (at least in distribution). If the ODE is unstable, then the algorithm is divergent. We follow the approach of [5], which is based on the techniques of [8]. This approach is conceptually similar to stochastic approximation theory, but its assumptions (and hence conclusions) are somewhat different. First, the stepsize μ in (5) and (6) is fixed, unlike in stochastic approximations, where the stepsize is required to converge to zero [16]. Thus, the algorithms do not necessarily converge to a fixed vector; rather, they converge in distribution. Second, no continuity or differentiability assumptions need to be made on the update terms. Hence, (7) is as amenable to the method as (5) and (6).

To be specific, consider the algorithm as a discrete time iteration process

$$p(k+1) = p(k) + \mu G(p(k), U(k+1)) \quad (14)$$

where $p(k)$ is a vector of weights that define the probabilities, μ is the stepsize, and $U(k)$ is a (random) input vector. The function $G(\cdot, \cdot)$ represents the update term of the algorithm and is, in general, discontinuous, as in (7), although it may also be differentiable as in (5) and (6).

What is the nature of the random process $\{p(k)\}$? When is this process stable? How can we characterize its convergence to steady states? These questions can be addressed by relating the behavior of algorithm (14) for small μ to the behavior of the associated integral equation

$$p(t) = p(0) + \int_0^t \hat{G}(p(s)) ds$$

or, equivalently, to the associated deterministic ODE

$$\dot{p}(t) = \hat{G}(p(t)) \quad (15)$$

where $\hat{G}(\cdot)$ is a version of $G(\cdot, \cdot)$ that is smoothed, or averaged, over all possible inputs. Speaking loosely, the ODE $p(t)$ of (15) represents the ‘‘averaged’’ behavior of the parameters $p(k)$ in (14).

Suppose that $(p(k), U(k))$ is adapted to the filtration $\{\mathcal{F}_k\}$, and define

$$\hat{G}(p(k)) = E\{G(p(k), U(k+1)) | \mathcal{F}_k\} \quad (16)$$

to be a version of G that is smoothed by the distribution of the inputs $U(k+1)$. This smoothed version is often differentiable even if G itself is discontinuous. A time-scaled version of p is defined as

$$p_\mu(t) = p_{[t/\mu]}(t), \quad t \in [0, \infty)$$

where $[z]$ means the integer part of z . Note that $p(k)$ represents the discrete iteration process, whereas $p_\mu(t)$ represents a continuous time-scaled version. $p(t)$ (with no subscript) is the ODE (15) to which $p_\mu(t)$ converges weakly.

Let (E, r) denote a metric space with associated Borel field $\mathcal{B}(E)$, and let $D_E[0, \infty)$ be the space of right continuous functions with left limits mapping from the interval $[0, \infty)$ into E . We let $C_E[0, \infty)$ denote the subspace of continuous functions and assume that $D_E[0, \infty)$ is endowed with the Skorohod topology.

Let $\{X_\alpha\}$ (where α ranges over some index set) be a family of stochastic processes with sample paths in $D_E[0, \infty)$, and let $\{P_\alpha\} \subset \mathcal{P}(D_E[0, \infty))$ be the family of associated probability distributions (i.e., $P_\alpha(B) = P\{X_\alpha \in B\}$ for all $B \in \mathcal{B}(E)$). We say that $\{X_\alpha\}$ is relatively compact if $\{P_\alpha\}$ is relatively compact in the space of probability measures $\mathcal{P}(D_E[0, \infty))$ endowed with the topology of weak convergence. The symbol \Rightarrow will denote weak convergence, whereas the arrow \rightarrow will denote convergence under the appropriate metric. An excellent reference for all the mathematical terms and probabilistic constructs used in this section is [8].

Consider the following technical assumptions.

- 1) $\{\hat{G}(p(k)) : k \in \mathcal{Z}^+, \mu > 0\}$ is uniformly integrable.
- 2) $\mu^2 \sum_{k=1}^{[t/\mu]} E\{(G(p(k), U(k+1)) - \hat{G}(p(k)))^2\} \rightarrow 0$.

Theorem 5.1: Under assumptions 1–2, $\{p_\mu\}$ is relatively compact, and every possible limit point is a random process in $C[0, \infty)$. Furthermore, every limit point of $\{p_\mu\}$ satisfies (15).

This is a special case of Theorem 1 in [5]. Both the uniform integrability (technical assumption 1) and the mean convergence in assumption 2 follow directly from the boundedness of the $p(t)$.

The theorem asserts that the iteration (14) will behave like the ODE (15) for small enough μ . If the solution to the ODE is unique, then the sequence actually converges in probability (not just having a weakly convergent subsequence). The solution of the ODE is continuous, and the Skorohod topology for continuous functions corresponds exactly to uniform convergence on bounded time intervals. Hence, convergence in probability means that for every $T > 0$, $\epsilon > 0$, and $\lim_{\mu \rightarrow 0} P(\sup_{0 \leq t \leq T} |p_\mu(t) - p(t)| > \epsilon) = 0$. This is useful because the algorithm behaves like the relevant ODE, and the ODE can often be analyzed in a straightforward manner.

D. Convergence of the Algorithms

This section considers the convergence behavior of the algorithms (5) and (6) by finding the appropriate ODE (15) and examining its stability properties.

1) *Algorithm With Partial Information:* The appropriate smoothed update (16) for algorithm (5) is

$$\hat{G}(p_i) = E\{(N(k) - \mathcal{N}) x_i(k)\}, \quad \text{for } i = 1, 2, \dots, M.$$

This can be rewritten exactly as in Section V-A1 as

$$\left(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq i}}^M p_j(k)\right) p_i(k)$$

and the ODE (15) is then

$$\dot{p}(t) = \begin{pmatrix} \dot{p}_1(t) \\ \dot{p}_2(t) \\ \vdots \\ \dot{p}_M(t) \end{pmatrix} = - \begin{pmatrix} (1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq 1}}^M p_j(t)) p_1(t) \\ (1 - \mathcal{N} + \sum_{\substack{j=2 \\ j \neq 2}}^M p_j(t)) p_2(t) \\ \vdots \\ (1 - \mathcal{N} + \sum_{\substack{j=M \\ j \neq M}}^M p_j(t)) p_M(t) \end{pmatrix}. \quad (17)$$

The theorem of the previous section shows that the iteration (5) behaves like this ODE. The remainder of this section shows that this ODE (and, hence, the algorithm with partial information) converges to steady states p^* that consist of \mathcal{N} ones and $M - \mathcal{N}$ zeroes by showing that the ODE is stable about these steady states.

Let $\tilde{p}(t) = p(t) - p^*$ be the ‘‘error’’ term. Stability of the ODE (17) about p^* is equivalent to stability of

$$\dot{\tilde{p}}_i(t) = - \left(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j \neq i}} \tilde{p}_j(t) + \sum_{\substack{j=1 \\ j \neq i}} p_j^*\right) (\tilde{p}_i(t) + p_i^*) \quad \text{for } i = 1, 2, \dots, M \quad (18)$$

about the origin $\tilde{p}(t) = 0$.




Fig. 5. Numerical simulation of the “partial information” ODE (17) for the case where $M = 100$ and $\mathcal{N} = 60$.

If $i \in I_0$ (the set of indices of p^* with zero entries), then $\sum_{j \neq i} p_j^* = \mathcal{N}$ and $p_i^* = 0$, and (18) becomes

$$\dot{\tilde{p}}_i(t) = - \left(1 + \sum_{j \neq i} \tilde{p}_j(t) \right) \tilde{p}_i(t) \quad i \in I_0. \quad (19)$$

Thus, for small perturbations \tilde{p} away from steady state, these states are exponentially stable.

If $i \in I_1$ (the set of indices of p^* with entries equal to one), then $\sum_{j \neq i} p_j^* = \mathcal{N} - 1$, $p_i^* = 1$, and (18) becomes

$$\dot{\tilde{p}}_i(t) = - \sum_{j \neq i} \tilde{p}_j(t) [\tilde{p}_i(t) + 1] \quad (20)$$

which is not stable about $\tilde{p}(t) = 0$, as can be seen from a linearization argument. However, algorithm (5) clips $p(k)$, and hence, the ODE must clip $p(t)$, that is, $0 \leq p(t) \leq 1$ is enforced by the algorithm statement. (This is a result of the meaning of $p(t)$ as a probability.) Hence, $\tilde{p}_i(t) = p_i(t) - p_i^* \leq 0 \forall i \in I_1$, which implies that $\sum_{j \in I_1} \tilde{p}_j(t) \leq 0$ once (19) has converged. Hence, the right-hand side of (20) is positive for small $\tilde{p}_j(t)$, $\tilde{p}_j(t) > 0$, and $\tilde{p}_i(t) \rightarrow 0$ from below (equivalently, $p_i(t) \rightarrow 1$ from below for $i \in I_1$). Formally, this requires a decomposition into the exponentially stable states $i \in I_0$ and the clipped states $i \in I_1$. (See [11] for such a formalization.) The exponential stability of (19) guarantees that the I_0 states converge rapidly, after which, the I_1 states converge. The practical upshot is that the region of convergence about each steady state is correspondingly smaller.

Thus, the ODE converges to one of the steady states p^* , assuming it is initialized close enough and the stepsize is small enough. The algorithm (5) converges similarly. Fig. 5 shows a numerical simulation of this ODE (17) for the case where $M = 100$ and $\mathcal{N} = 60$.

2) *Algorithm With Full Information:* The appropriate smoothed update (16) for the algorithm (6) with full information is

$$\hat{G}(p_i) = E\{(N(k) - \mathcal{N})\} \text{ for } i = 1, 2, \dots, M.$$

As in Section V-A2, this can be rewritten

$$\sum_{j=1}^M p_j(t) - \mathcal{N}.$$




Fig. 6. Numerical simulation of the “full information” ODE (21) for the case where $M = 100$ and $\mathcal{N} = 60$.

The relevant ODE (15) consists of M identical copies of the scalar ODE

$$\dot{p}_i(t) = - \left(\sum_{j=1}^M p_j(t) - \mathcal{N} \right) \quad (21)$$

which has steady states at any p^* with $\sum_{j=1}^M p_j^* = \mathcal{N}$. The stability properties are easy to describe. Given any $p(0)$ with $\sum_{j=1}^M p_j(0) = p_0$, all entries in (21) increase or decrease together until $\sum_{j=1}^M p_j(t) = \mathcal{N}$. Thus, the ODE converges to $p^* = (\mathcal{N}/p_0)p(0)$. Accordingly, for each initial condition $p(0)$, there is a unique steady state to which the algorithm will converge.

Fig. 6 shows a numerical simulation of this ODE (21) for the case where $M = 100$ and $\mathcal{N} = 60$. The content of the weak convergence theorem is that the actual trajectories of algorithms (5) and (6) must, on average, follow the trajectories of the ODEs (17) and (21), at least for small μ . Since the steady states of the ODEs (17) and (21) are stable, these algorithms must also converge to a steady state if initialized close enough.

VI. DISCUSSIONS AND CONCLUSIONS

The *El Farol* problem initially explored the collective dynamics of boundedly rational agents, but we have shown that this model is also interesting as a simple model of congestion and coordination behaviors that occur with shared resources like Internet bandwidth.

Arthur [3] believed that any solution to the *El Farol* problem would require heterogeneous agents, that is, agents who pursue different strategies. In contrast, we have presented a simple adaptive solution that can be followed by all agents and can readily solve a decentralized resource allocation problem. Each agent in the adaptive solution is characterized by a parameter that determines how often the agent attends and a stepsize that determines how much to change the parameter in response to each visit to the bar.

The stochastic adaptive solution to the *El Farol* problem differs from previous treatments in several ways. We allow agents to proceed stochastically, as is commonly required for optimal game-theoretic decision making schemes. We do not require

agents to make explicit predictions of the state of the bar, and we allow them to use only the information that they have readily available, i.e., their own experiences. This makes our model more realistic (it is not clear how the agents in Arthur's model learn what happened at the bar when they are absent). Certainly, in applications like the Internet, such information is not available; the only way to know if a website is crowded is log on and try to use it.

Arthur's solution, in which each agent maintains a bank of strategies, leads to patterns of attendance that fluctuate considerably above and below the optimal. When crowded, none of the agents enjoys themselves. When undercrowded, there is a wasted resource represented by the empty seats at the bar. The stochastic adaptive solution, in contrast, leads to patterns of attendance with much smaller variance and, hence, much less waste. Generically, the attendance at the bar converges to an optimal solution: one where the bar is neither under nor overcrowded.

On the other hand, changing the information structure in the algorithm so that agents adapt their probabilities at every iteration causes the algorithm to no longer converge to such an optimal solution; rather, the attendance patterns continue to fluctuate wildly. Thus, we posit that the information structure is the crucial difference in our approaches. When agents each receive or utilize a subset of total information, then the system is far better behaved than when all act on complete information. In other words, the homogeneity of information may be the key ingredient driving the *El Farol* "problem." With more heterogeneous information, the problem may vanish.

The adaptive solution thus provides a simple mechanism whereby a large collection of decentralized decision makers, each acting in their own best interests and with only limited knowledge, can solve a complex congestion and social coordination problem. Moreover, convergence to the solution is relatively rapid (depending on the initial conditions) and robust.

Do we believe that customers of *El Farol* tick off the time until they can go again, increasing or decreasing the probability of a coin flip with each new visit? Of course not. But the incentives are in agreement with the common sense idea that people tend to minimize bad experiences and maximize good ones. Moreover, the global behavior of the population is consistent with certain kinds of coordination phenomena. For instance, users of an Internet provider can spread demand over much of the day, even though everyone might prefer (all else being equal) to log on in the middle of the afternoon. By developing certain habits (for instance, always logging on at the same time), users send signals to others to avoid these times. In this way, demand is smoothed.

There are many ways to generalize the adaptive solution to decentralized resource allocation problems. For instance, different people have different tolerances for what constitutes a crowd or an unacceptable delay. Each agent could also have a parameter that represents their tolerance for congestion. Additionally, to more closely model the Internet situation, one might incorporate time-of-day or day-of-week as a parameter in the process of logging on. It would also be instructive to create a hybrid situation in which a number of Arthur-like agents and a number of adaptive agents compete for spaces at the bar.

REFERENCES

- [1] R. Arnott, A. De Palma, and R. Lindsey, "Information and usage of free-access congestible facilities with stochastic capacity and demand," *Int. Econ. Rev.*, vol. 37, no. 1, pp. 181–203, Feb. 1996.
- [2] —, "Does providing information to drivers reduce traffic congestion?," *Transpn. Res.-A*, vol. 25A, no. 5, pp. 309–318, 1991.
- [3] W. B. Arthur, "Inductive reasoning and bounded rationality: The *El Farol* problem," *Amer. Econ. Rev.*, vol. 84, pp. 406–411, May 1994.
- [4] A. M. Bell and W. A. Sethares, "Avoiding global congestion using decentralized adaptive agents," *IEEE Trans. Signal Processing*, vol. 49, Nov. 2001.
- [5] J. A. Bucklew, T. Kurtz, and W. A. Sethares, "Weak convergence and local stability of fixed step size recursive algorithms," *IEEE Trans. Inform. Theory*, vol. 39, pp. 966–978, May 1993.
- [6] J. L. Casti, "Seeing the light at *El Farol*," *Complexity*, vol. 1, no. 5, pp. 7–10, 1996.
- [7] D. Challet and Y.-C. Zhang, "On the minority game: Analytical and numerical studies," *Phys. A*, vol. 256, pp. 514–532, 1997.
- [8] S. Ethier and T. Kurtz, *Markov Processes – Characterization and Convergence*. New York: Wiley-Interscience, 1986.
- [9] N. S. Glance and B. A. Huberman, "The dynamics of social dilemmas," *Sci. Amer.*, pp. 76–83, Mar. 1994.
- [10] A. Greenwald, B. Mishra, and R. Parikh, "The Santa Fe Bar problem revisited: Theoretical and practical implications," in *Mimeo*. New York: New York Univ. Press, 1998.
- [11] W. Hahn, *Stability of Motion*. Berlin, Germany: Springer-Verlag, 1967.
- [12] S. Haykin, *Adaptive Filter Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1991.
- [13] B. A. Huberman and R. M. Lukose, "Social dilemmas and internet congestion," *Sci.*, vol. 277, pp. 535–537, July 1997.
- [14] N. F. Johnson, S. Jarvis, R. Jonson, P. Cheung, Y. R. Kwong, and P. M. Hui, "Volatility and agent adaptability in a self-organizing market," *Phys. A*, vol. 256, no. 230, 1998.
- [15] C. R. Johnson Jr, *Lectures on Adaptive Parameter Estimation*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [16] L. Ljung and T. Soderstrom, *Theory and Practice of Recursive Identification*. Cambridge, MA: MIT Press, 1983.
- [17] H. S. Mahmassani and R. Jayakrishnan, "System performance and user response under real-time information in a congested traffic corridor," *Transpn. Res.-A*, vol. 25A, no. 5, pp. 293–307, 1991.
- [18] A. Rapoport, D. Seale, and E. Winter, "An experimental study of coordination and learning in iterated two-market entry games," *Hong Kong Univ. Sci. Technol. Dept. Marketing, MKTG 98.116*, 1998.
- [19] R. Rosenthal, "A class of games possessing pure strategy Nash equilibria," *Int. J. Game Theory*, no. 2, pp. 65–67, 1973.
- [20] R. Savit, R. Manuca, Y. Li, and R. Riolo, "The dynamics of minority competition," in *Proc. Second Int. Conf. Complex Syst.*, Nashua, NH, Oct. 1998.
- [21] W. A. Sethares, "Adaptive algorithms with nonlinear data and error functions," *IEEE Trans. Signal Processing*, vol. 40, pp. 2199–2206, Sept. 1992.
- [22] W. A. Sethares and A. M. Bell, "An adaptive solution to the *El Farol* problem," in *Proc. Thirty-Sixth Annu. Allerton Conf. Commun., Contr., Comput.*, Allerton, IL, Sept. 1998.
- [23] J. Sundali, A. Rapoport, and D. Seale, "Coordination in market entry games with symmetric players," *Organiz. Beh. Human Decision Process.*, vol. 64, no. 2, pp. 203–218, 1995.
- [24] E. Zambrano, "Rationalizable boundedly rational behavior," Santa Fe Inst., Santa Fe, NM, 97–06–060, 1997.



Ann M. Bell received the B.A degree from Cornell University, Ithaca, NY, and the M.A. and Ph.D. degrees from the University of Wisconsin, Madison, all in economics.

She currently works for Orbital Sciences, doing research in Advanced Life Support at the NASA Ames Research Center, Moffett Field, CA, and for the nonprofit organization Redefining Progress, analyzing the implications of new models of complex systems for public policy. She has also been an Assistant Professor of Economics at Vanderbilt University, Nashville, TN, and a Visiting Assistant Professor at the University of California, Santa Cruz.



William A. Sethares (M'87) received the B.A. degree in mathematics from Brandeis University, Waltham, MA, and the M.S. and Ph.D. degrees in electrical engineering from Cornell University, Ithaca, NY.

He was with the Raytheon Company, Waltham, as a Systems Engineer and is currently Associate Professor with the Department of Electrical and Computer Engineering, University of Wisconsin, Madison. His research interests include adaptation and learning in signal processing, communications,

and acoustics. He is the author of *Tuning, Timbre, Spectrum, Scale* (New York: Springer, 1998). His website at <http://eceserv0.ece.wisc.edu/~sethares/> contains further information about the El Farol problem and other fascinating topics.

James A. Bucklew (SM'02) received the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN, in 1979.

He is currently a Professor with the Departments of Electrical and Computer Engineering and Mathematics at the University of Wisconsin, Madison. His research interests are in the application of probability and statistics to signal processing and communication problems. He has published over 100 articles in these fields. He has published several fly fishing articles (in Spanish) and works as a translator.

Dr. Bucklew has served as an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY from 1990 to 1992 and for the IEEE TRANSACTIONS ON SIGNAL PROCESSING from 1997 to 1999.