

## Research Article

# Simulating Systems Thinking under Bounded Rationality

Mark W. Sellers <sup>1,2</sup>, Hiroki Sayama <sup>1,2</sup> and Andreas D. Pape<sup>1,3</sup>

<sup>1</sup>Center for Collective Dynamics of Complex Systems, Binghamton University, State University of New York, Binghamton, NY 13902-6000, USA

<sup>2</sup>Department of Systems Science and Industrial Engineering, Binghamton University, State University of New York, Binghamton, NY 13902-6000, USA

<sup>3</sup>Department of Economics, Binghamton University, State University of New York, Binghamton, NY 13902-6000, USA

Correspondence should be addressed to Mark W. Sellers; [mseller1@binghamton.edu](mailto:mseller1@binghamton.edu)

Received 22 February 2020; Revised 27 July 2020; Accepted 5 September 2020; Published 16 September 2020

Academic Editor: Quanmin Zhu

Copyright © 2020 Mark W. Sellers et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Brian Arthur's El Farol bar model of bounded rationality provides a simple computer model of decision making in a complex, dynamic, and self-organized environment. Can systems thinking provide a viable alternative strategy to traditional methods for dealing with these types of problems? Nine different agents, designed from both traditional and systems perspectives, compete in fifteen variants of the El Farol environment and their performance in 4 categories—Winner, Top Performers, Competitive, and Vulnerable—is compared. We show that systems thinking is a competitive strategy that is, at least, on par with traditional strategies and may be less vulnerable to elimination or ruin. However, there are two consequential elements that emerge. First, all strategies have some environments where they succeed and others where they fail. Second, as the population of practitioners adopts these adaptive, systems-based strategies, the environment exhibits new behaviors with a new set of unintended consequences.

## 1. Introduction

The modern world is a complex, interconnected network of natural and man-made systems. We choose our actions within them every day. Many of these systems are dynamic, adaptive, and self-organized. These types of systems often exhibit nonlinear, random, chaotic, or emergent behaviors. The El Farol bar in Santa Fe, New Mexico (Figure 1), is a real-world example of this interaction with complexity which is famously described in Arthur [1] (more on that below). As per Arthur, patrons want to go to the bar on nights when there is live music—but not when it is crowded. The agents in Arthur's description make their decision to go or not applying a random set of heuristics to imperfect historical data. The number of seats is fixed, but the aggregate of the decisions and the resulting level of crowding change significantly week to week under the influence of these decisions. The dynamic behaviors of complex, interactive, social, and economic systems similar to this can be consequential.

Modern thinkers such as Nasim Nicholas Taleb, Karl Popper, Dietrich Doerner, Benoit Mandelbrot, Edward Lorenz, and others have discussed, in great detail, the uncertainty in social and economic systems, as well as the inadequacy of modern stochastic and analytical tools, to predict the behaviors of these systems. The mathematics of fractals, networks, chaos, feedback, and systems dynamics, among others, demonstrate that the understanding of such systems does not improve with increasing precision in the data or models. This is a shift away from our traditional model of a deterministic, mechanical world and illuminates significant shortcomings in our traditional approach. In order to address this limitation in traditional understanding of the behavior of the world, systems researchers over the last 100 years defined a new approach to explore these problems. That perspective is captured under a wide and vague term—"Systems Thinking."

"Systems thinking" means many things to many different people. Sellers [2] surveyed a broad sample of 20th century



FIGURE 1: El Farol bar, Santa Fe, NM.

authors and found that the collected set of descriptions of systems thinking is a complex array of skills in math, computers, systems theory, networks, dynamic systems, domain expertise, and personal characteristics. This definition of systems thinking is so complicated that many researchers have come to the conclusion that it may not be accessible to most people. Popular, optimistic renderings of this comprehensive type of systems thinking envision it as a panacea that will avoid or reverse the unintended global consequences of our legacy short-term, linear decisions. This utopian promise of this broad systems thinking may or may not be true; however, neither its legitimacy nor application has been substantiated by empirical research. Even if it had been proven, this extensive version of “Systems Thinking” is not practical for the common type of interaction between people and complexity that is studied in the El Farol bar problem. That problem requires a pragmatic version of systems thinking.

Stroh [3], Sweeney and Meadows [4], and others use simple counterintuitive games, guided group discussions, and collective systems analysis of past failures to illuminate flaws in our traditional ways of thinking and provide some understanding of the value and perspectives of systems thinking. Senge [5], Ackoff [6], Klir [7], and others propose that acquiring a theoretical understanding of the behavior, relationships, and structure in a system in addition to traditional component reductionism is the key to effective systems thinking. However, while both analyses of past failures and expanded systems perspectives provide important new insights into systems, they do not necessarily provide simple, actionable recommendations that individuals can apply to their daily interactions with complexity.

Pragmatic approaches to systems thinking are very different from the theoretical approaches mentioned above. Doerner [8], Meadows and Wright [9], and others advocate a simple and conservative, yet dynamic, approach to the interactions with complex systems. Doerner demonstrated that the most successful strategy when interacting with unknown complex systems is to assume the long-term behaviors are not understood and every decision is subject to change. Doerner recommends a strategy of “act and measure” to gain an understanding of the unique characteristics of the system through simple piecewise action. Meadows advocates that we “get the beat of the system” and “stay

humble-stay a learner,” which supports the belief that we may never fully understand the dynamics of the system or future consequences—unintended or otherwise—of our decisions. Taleb [10–12] states that our world is random, antifragile, and driven by black swans. His version of systems thinking is intended to guide individual actions in a complex world where history is marginally useful, linear projections are inaccurate, bell-curve analysis is illusory, and outcomes are unpredictable. Taleb [12] states that “You get pseudo-order when you seek order; you only get a measure of control when you embrace randomness.” Doerner’s, Meadows’, and Taleb’s approaches are all inherently similar. They (1) do not require that we fully understand the complex, dynamic, random, or fat-tailed system, (2) do not require that we develop a working theory of the system, and (3) do not require that we develop a complex array of “system thinking” skills. They do require that we acknowledge the limits of our understanding and mental models while learning how to act in the presence of these dominating and consequential unknowns.

*1.1. Related Literature.* The El Farol bar problem as framed by Brian Arthur is a simple simulation that demonstrates the capacity for stability in a closed, social, complex, and adaptive system with participants that have limited access to data and limited data processing capability. This simple simulation has found applications in a wide array of disciplines including disciplines as widely diverse as statistical mechanics [13], computer networking [14], shared resource optimization [15], 3<sup>rd</sup> party mediation [16], social networking [17], dynamic learning processes [18, 19], and self-organized networks [20]. All of these applications are possible because at the core of the El Farol model is a stable, dynamic system that approaches equilibrium but never converges on any solution. Researchers can easily modify the system to test the desired characteristic of interest while leaving all other parameters unchanged. In this paper, the El Farol model provides a consistent complex environment test platform where various simple mental strategies compete. Other researchers have investigated the interactions of multiple strategies using this model. St Luce and Sayama, using a novel approach, recently evaluated the phase spaces of self-organizing suites using only 2 or 3 very simple strategies and found among other things that “The distribution of strategies used within the El Farol bar problem played a vital role in the behavior and success of the agents [21].” The common thread through all of this research is that the El Farol model of bounded rationality provides a simplistic, stable, and flexible simulation of human decision making under the constraints of limited data with limited processing of that data and that behavior is the platform for this research.

## 2. Methods

*2.1. Testing Practical Systems Thinking.* The technique used in this paper to explore pragmatic systems thinking is straightforward in principle: Simulate a complex system with

interacting agents making decisions, some traditional thinkers, and some systems thinkers, and see who wins. Of course, and importantly, we must carefully define some reasonable ways to represent “systems thinking” in this environment.

The environment we choose is Brain Arthur’s [1] exploration of bounded rationality in the El Farol bar extension 3 [22] simulation as implemented in NetLogo [23], which is an ideal platform for this study. It is easy to program, reproducible, stable, a model of human decision making under uncertainty, and has been explored by a large number of researchers.

We represent two aspects of practical systems thinking methods in this model. First, following the suggestions of Doerner, an agent type was designed that, like the El Farol traditional agents, monitored the system behavior, chose its best strategy, and made a decision. However, unlike the original agents, after each decision, they have an additional step where they can review and change their set of strategies. These agents act, measure their success, and then, apply small random variants to their set of decision strategies, replacing the least accurate each time. We call these agents “Double Loop Learners” below (Strategic Competitors). Second, following a more structural and behavioral approach as suggested by Klir, Ackoff, and others, an agent was designed to exclusively monitor the relationships in the system. This agent did not observe the system directly, but observed the behavior and prediction accuracy of a random sample of other agents interacting with the system, and based its decision on the successes and failures observed in that sample. This agent does not know the architecture or history of the system and does not collect a personal history of success. The understanding of the system was exclusively an aggregate of the immediate observed relationships between the system and a random set of agents. In keeping with spirit of Doerner, Taleb, and others, neither of these agent types attempts to create a model of the system. Their strategies evolve from current observations only. We call these agents “Observers” below (Strategic Competitors).

As a traditional thinking baseline, we develop a pair of rational actors using weighted temporal observation of the recent system history. One of these agents had an option to purchase “insurance” each week. They are called “Rational Actors” and “Rational Actors, Insured” below (Strategic Competitors).

**2.2. Complex Environment.** The El Farol bar game provides a complex market environment populated with 100 agents acting based on a small set of unique random strategies assigned to each agent. Each agent’s goal is to go to an uncrowded bar and avoid a crowded bar. Each week, each agent chooses whether to go or not, using the learning method described below. The 100 “go” and “no-go” decisions are counted and compared to the number of seats in the bar. Each agent choosing correctly is, then, rewarded.

In the original model, based on how the agents learned, the number of agents going to the bar each week never diverged or converged, but continuously varies (the original

El Farol learners are referred to as “Single Loop Learners” below.) The system behavior in the original model varies significantly with each new random number seed where a new set of 5,000 random numbers is generated and assigned to the agents as their “strategies” (see below). This independent, unpredictable, yet moderately stable, system provides an ideal platform to test various other strategies.

In addition to varying the random seed with each run, two independent variables in the system were varied to create a total of 15 test environments. Each of the 15 environments was exercised 400 times for a total of 6000 individual runs. Each run was 250 steps in the simulation, and the statistics for all of the competing agents were recorded at the end of the 250 steps. The 2 independent environment variables varied for the study are as follows:

- (1) The 100 core agents are aggregates of 2 types of agents (defined below)—traditional and “systems thinking.” The ratio of traditional “Single Loop Learners” to systems thinking “Double Loop Learners” was tested at 3 conditions: 5% and 95%, 50% and 50%, and 95% and 5%.
- (2) The ratio of agents to available seats was varied by setting the total number of seats available in the bar to 30, 40, 50, 60, and 70 while holding a constant 100 competing agents.

**2.3. Strategic Competitors.** In these 15 environments, 9 unique types of agents compete for seats in the bar. To simplify the statistics and provide a consistent platform, only the decisions of the 100 core environment agents were counted against the number of seats available in the bar. The other 7 individual agents’ decisions were compared with the outcome of the core system.

- (1) “Single Loop Learners” (1 type, multiple random individuals): original El Farol model agents with a set of 10 random strategies that remain fixed for 1 complete 250 step run. Each week, these agents only “learn” which one of their fixed solutions is their best available strategy for that week. Their strategy repertoire never changes no matter how good or bad it may be. However, since they do have a random repertoire, they have a limited capacity to adapt as the system evolves over time by selecting a different strategy each week.
- (2) “Double Loop Learners” (1 type, multiple random individuals): these are the “single loop learners” that, in addition to selecting a best strategy each week as above, learns in a 2nd loop by replacing their worst strategy for that week with a new random strategy. The new strategy then competes for their “best” strategy the following week. These agents are very effective at adapting to the system and have significant impact on the system dynamics as their number is increased.
- (3) “Observer” (1 type, 2 unique): these 2 agents only use observed relationships. They do not observe the bar

or carry a personal history. They make their decision based exclusively on the observed behaviors of a randomly chosen set from the double and single loop learner agents. That list is constant for 1 complete run. They observe whether these agents are historically successful and what they chose last week. Only the behaviors of the most and least successful of those observed agents are combined into a go/no-go decision. Their strategy process never changes. One of these 2 agents observes a group of 30 random agents while the other observes a smaller group of 5.

- (4) “Rational Actor” (1 unique): this agent tracks the history of the last 10 weeks and derives these 3 data points to calculate a utility score of going and not going—(1) “Probability of crowding,” (2) average excess patrons when crowded, and (3) average extra seats when uncrowded. “Probability of crowding” is derived by scaling average attendance. For example, an average attendance of 50 with 60 seats available has a 1% probability of crowding while an average of 70 with 60 seats available has a 99% probability of crowding. The agent decides to go based on the net utility scores of going and not going. The utility of going is calculated as  $((1 - \text{Pcr}) * \text{Seats}) - (\text{Pcr} * \text{Standers})$ . The utility of not going is mathematically the negative of the utility of going. Essentially, the agent goes if the utility of going is positive and stays home if negative.
- (5) “Rational Actor, Insured” (1 unique): this agent is the “Rational Actor” mentioned above with an option to buy “insurance.” This agent calculates the same two uninsured go/stay utility scores as the “Rational Actor” along with two additional utility scores for insured options. The agent, then, selects the largest net value from the four options. The utility score for insured options is more complicated to calculate because an agent that chooses the insurance option essentially always chooses correctly and is always initially rewarded, however, at a cost (i.e., the insurance premium). The insurance cost is deducted from their reward plus an extra deduction for time lost if they chose wrong. An important point here is that the agent’s decision is based on history, probabilistic, and utility prediction, but the reward or loss of that decision is calculated from the actual results. If the agent’s costs exceed the “reward” (number of empty seats or excess patrons), the net value will fall below zero, and the week is counted as a failure. For example, the insured utility calculation for going is  $((1 - \text{Pcr}) * (\text{Seats} - \text{Insurance})) + (\text{Pcr} * (\text{Standers} - \text{Insurance} - \text{TimeLost}))$ . Because of the added cost of “Time lost” with a bad decision, it is in the best interest of the insured agent to choose correctly even with insurance. It is important to note that only over a critical price range will the agent judiciously choose between insured and uninsured options. Four levels of insurance costs were tested, and the frequency of selecting an insured option was monitored.

- (6) “Simple strategists” (3 unique): three extremely simple strategies were tested: “Always go,” “Same as last week,” and “Opposite of last week.”

*2.4. Scoring.* All agents receive either a 1 or 0 each week depending on the correctness of their prediction—both crowded and uncrowded nights. If they were scored only on the number of uncrowded nights that they attend, then a strategy of “Always go” would yield a perfect score. Therefore, staying home on crowded nights must be included in their score. Because of insurance, there are two different algorithms for computing prediction accuracy:

- (1) For any uninsured decision, the prediction accuracy is the percentage of predictions that are correct. The extent of crowding or available seats is not considered.
- (2) For insured agents, “correct” is defined as a net positive reward calculated from a gross reward minus insurance costs. An insured agent not only makes a prediction and acts on that prediction like all other agents but also buys insurance. An insured agent’s gross reward is based on the number of excess patrons or empty seats. Three excess patrons or 3 empty seats have a similar gross reward of +3. The insurance cost is subtracted from this gross reward, as well as a “time lost” penalty if the prediction was incorrect. If the net reward is positive, the agent was “correct” and the prediction accuracy is increased.

Each agent’s score is recorded as a percentage of the block of 400 simulations in a single environment where the agent met the “Winning Criteria” defined below. Each agent has 4 independent scores. Note that the winning criteria are not statistical bins; rather, they represent 4 distinct real-world decision points. For example, the score for “Competitive” includes agents in both “Win/Tie” and “Top Performer” categories. For “Competitive” agents, “making the cut” may be the only important life requirement. Each score represents the probability that one might be rewarded or punished by life.

*2.5. Winning Criteria.* Winning in a “many-versus-many” competition is not a simple raw score of how often the agent made a good decision; rather, it is whether the agent was better or worse at it than the other agents. In some competitions, winning agents made correct predictions 100% of the time, while in the most competitive run, the best agent had a prediction accuracy of only 55.6%. In order to compare multiple runs with diverse prediction accuracies, “winning” was determined as a percentile ranking at the end of each run. Four different levels of percentile ranking were considered in the study.

- (1) *Win/Tie.* First place: this is the group of winners who are publicly acknowledged where winning is perceived as the most important identifier of a skilled individual.

- (2) *Top Performers*. Performance above the 98%: competitions such as enrolment in prestigious universities and job promotions that have real-world consequences.
- (3) *Competitive*. Performance above the 50%: competition such as stock traders “beating the market average.”
- (4) *Vulnerable*. Performance below 15%: susceptible to extinction or elimination. Agents in this group may not continue in a competition because they perish, are ruined, or quit.

### 3. Overall Results

In this study, each competitor is a computer algorithm in a complex, self-organizing environment constructed entirely from other competitors. Other than the number of seats at the bar, there is no reference to an external environment. This type of competitive architecture, where the participants are both competitors and the playing field, is not uncommon in the world we live in. The dynamics of man-made systems, such as societies, economies, markets, law, traffic, and sports, emerge from social interaction and competition between human beings—with each competitor following some individual, internal, and mental strategy. These virtual, man-made interactive systems define our social environment. For the most part, they exist only within the mind of man and only interact with natural systems incidentally. As individual actors in a large, complex society, our daily lives are dominated by our participation in these virtual systems. Success in these competitions, or at the very least learning not to fail, is an important life skill that we cultivate throughout our lives either consciously or unconsciously.

Table 1 presents a complete picture of the success rates across all 15 environments. The behaviors of individual agents are compared across all 6000 runs in all 4 categories of winning criteria. There is, on average, 50 single and 50 double loop learners in each of the 6000 runs, which results in the 300,000 agent scores to compare. Accordingly, there is only 1 instance of each of the 7 other agents in each run. The total number of agents of each type that actually competed is used to calculate the net probabilities for the 4 winning criteria. Read the table horizontally to choose the best strategy for each category of winning and vertically to understand the expected outcome for each strategy.

The 100 single and/or double loop learners in each simulation are the basis for the percentile distribution and, as a result, are distributed evenly about the competitive category. The double loop learner is far more effective at regulating to the expected number of seats in the bar (see Figure 2 below) and as a result has fewer extreme results in Win, Top Performer, and Vulnerable categories. Single loop learners have a broader spread with an enticing 6.5% chance of being a Top Performer that is offset with an enhanced 18% chance of being Vulnerable. These are aggregate probabilities across all 15 environments, and as shown below, winning expectations for agents in both of these core groups are significantly affected by altering the number of seats or

mixture of single and double loop learners. Those effects are discussed in “Behavior of The Environment and Core Agents” below.

The 4 specially designed agents as shown in Table 1—“Observer 5,” “Observer 30,” “Rational No Insurance,” and “Rational with Insurance”—are significantly less vulnerable to elimination and, in general, exhibit far more successful behavior in every category. In the aggregate, or when competing in an unknown system, any one of these strategies is a good choice. However, as discussed in “Designed Agents Performance”, their success also varies significantly across the 15 environments.

The 3 simplistic agents with fixed strategies exhibit mixed behaviors. The “Always Go” strategy is appealing, Winning 4.2% of the time, but Vulnerable 31% of the time. The “Follow Last Week” strategy provides a slight statistical advantage in vulnerability with no possibility of winning. The “Opposite of Last Week” strategy is a slightly milder mixed bag than the “Always Go” strategy. There is a slightly improved chance to Win, coupled with an increased vulnerability. Of the 3, “Follow Last Week” is the pragmatic best choice since it is the least vulnerable of the 3 choices.

#### 3.1. Behaviour of the Environment and Core Agents

*3.1.1. Traditional vs. Systems Thinking Environment.* The data suggest that a world with all systems thinkers is not necessarily a better place. While an individual agent’s competitive predictive performance improved with applied system thinking, the final system does not eventually stabilize as the number of systems thinkers in the system increases. Even an initially stable system is destabilized as the number of these better, more accurate predictors are added to the system (as shown in the top graph in Figure 3). In the two plots in Figure 3, each simulation begins with 1 systems thinker (gray line) and, then, 97 additional traditional single loop agents are enhanced with the second learning loop at the midpoint and begin to evolve their strategies. Note how both stable and volatile traditional agent attendance behavior are transformed into a new type of “system thinking” instability. This unstable behavior is characteristic of all “systems thinking-”adapted El Farol systems regardless of the initial systems behavior. The implication of this is clear. The systems thinker must always be aware that it is part of the system within which it is making a decision and that its individual behavior does impact the system behavior.

*3.1.2. Prediction Accuracy and Winning.* The exact crowd prediction accuracy needed to win over the 250 simulated steps changes with every simulation. In some competitions, the winning agent accurately predicted crowding 250 out of 250 times while in the most competitive simulations, the winning agent was correct only 139 times. The distribution of winning prediction accuracy of the entire sample set is as shown in Figure 4. Across all 15 environments, the expectation for the prediction accuracy in each category is—Winner: 78.5%; Top Performer: 72.5%; Above average: 50.4%; and Vulnerable: <38.4%. This mean value does not

TABLE 1: Winning/vulnerable rates for a single agent by agent type.

	Single loop learner	Double loop learner	Observe 5	Observe 30	Rational without insurance	Rational with insurance	Always go	Follow last week	Opposite of last week
Count	300,000	300,000	6,000	6,000	6,000	6,000	6,000	6,000	6,000
Win	2447	153	142	1,494	344	1,645	250	0	133
Top	19,371	3,101	1,107	2,920	1,195	2,436	775	21	479
Comp	152,505	152,719	5,229	5,786	4,661	5,452	3,328	3,496	2,810
Vuln	52,709	33,460	225	42	235	97	1,822	724	1,078
$P_{Win}$	0.8%	0.1%	2.4%	25%	5.7%	27%	4.2%	0.0%	2.2%
$P_{Top}$	6.5%	1.0%	18%	49%	20%	41%	13%	0.4%	8.0%
$P_{Comp}$	51%	51%	87%	96%	78%	91%	56%	58%	47%
$P_{Vln}$	18%	11%	3.8%	0.7%	3.9%	1.6%	30%	12%	18%

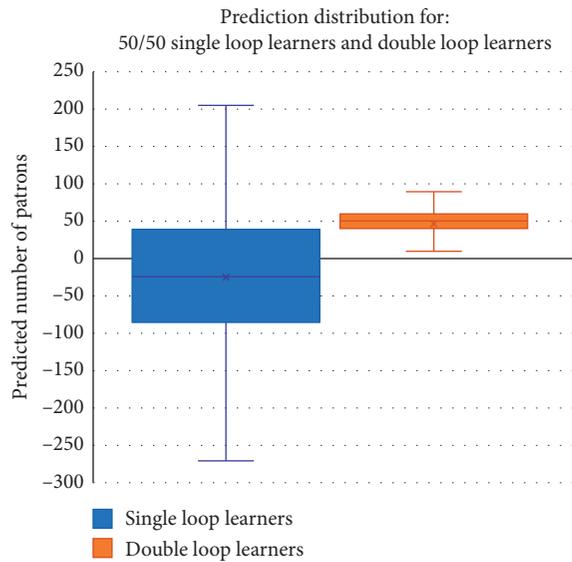


FIGURE 2: Prediction accuracy for single and double loop learners.

tell the complete story as there is a substantial distribution around those values across all simulation runs. Figure 4 shows the aggregate prediction accuracy distribution by winning criteria. Figure 5 unpacks the 4 levels of winning from Figure 4 across the 3 environments defined by the different mixtures of the single and double loop learning agents. In Figure 5, for example, in a “systems thinking” environment with 5/95 Single to Double loop learners (95 systems thinkers), on average, a low prediction accuracy of 64.4% will make you a Top Performer, while avoiding Vulnerability requires on average a prediction accuracy of at least 40.7%. In a traditional environment, with 95 single loop learners, the average prediction accuracy for Top Performers is 80.5%, while avoiding Vulnerability only requires a prediction accuracy of 27.0%. Note in Figure 5 that the spread of first place “Winner” category prediction accuracies for the 95 double loop learners appears to be inappropriately large. This broad “winner” distribution occurs because, as shown in Table 1, the double loop learners rarely win outright. One of those 95 agents only win 6.3% of their simulations. The spread in that category of winning in 5/95 environments is driven by the prediction accuracies of the other 8 types of winning agents. Remarkably, one of the 5 single loop learners in these 95 “system thinking” agent environments

win in 13.6% of the simulations for a net probability of winning with that strategy of 2.7% while the 95 double loop learners have only a net winning probability of 0.07%. However as shown before, the ever-present vulnerability must also be considered. The 5 single loop learning agents have an increased vulnerability of 26.7%, while the 95 double loop learners have a slightly reduced net 13.6% chance of landing in the vulnerable bottom 15%.

The difference between the winning thresholds across the different agent environments illustrated in Figure 5 is driven by the dramatic improvement in predictions of Double Loop Learners. Figure 2 clearly illustrates that the systems thinking “learning” process significantly decreases the diversity in the predictions of the core agents. The resulting uniformity means that blocks of system thinkers will arrive at the same predicted patronage simultaneously and as a result, the patronage can jump from week to week—as shown in the characteristic “systems thinking” volatility in Figure 3. Another important side effect of large number of agents converging to similar behaviors is that few agents will exhibit extreme success or failure rates and the overall distribution of prediction accuracies will shrink. This trend is clearly visible in Figure 5.

*3.1.3. Available Seats (Ratio of Agents to Seats).* Figure 6 shows the distribution of probability of open seats from the simulation, sorted by total seats in the bar. This impact to available open seats is an artifact of the ratio of seats to players and has important impacts on the average prediction accuracy of the agents. Consider the case where the bar has only 1 seat. If everyone chooses not to go, then everyone fails. If exactly 1 agent chooses to go, then only that agent predicts correctly. In all other cases, the bar will have no open seats, in other words, a very low probability of available seats, and all of the agents choosing not to go will predict correctly. Any strategy with a bias for staying home will have very high prediction accuracy—driving up the threshold prediction accuracy for winning while driving down the prediction accuracy necessary to be vulnerable. The same logic holds for 99 seats and a bias for going. If this is true, then the most difficult competition will occur at the minima where the probability of crowding is near 50%. As shown Figures 7 and 8, the prediction accuracy needed to win as a function of the number of seats has a minimum at 50 seats. In Figure 7, an aggregate of all agent environments, winning with 30 or 70

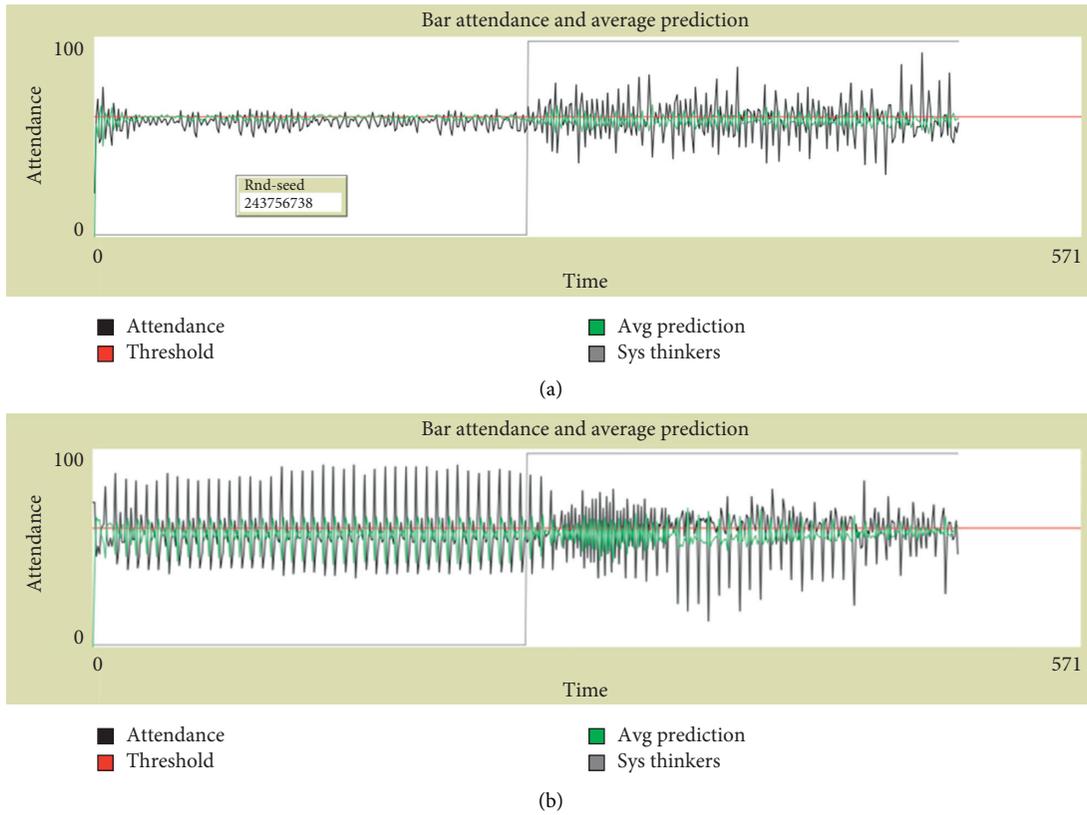


FIGURE 3: Bar attendance (black), prediction (green), seats (red), and DLL agents (gray).

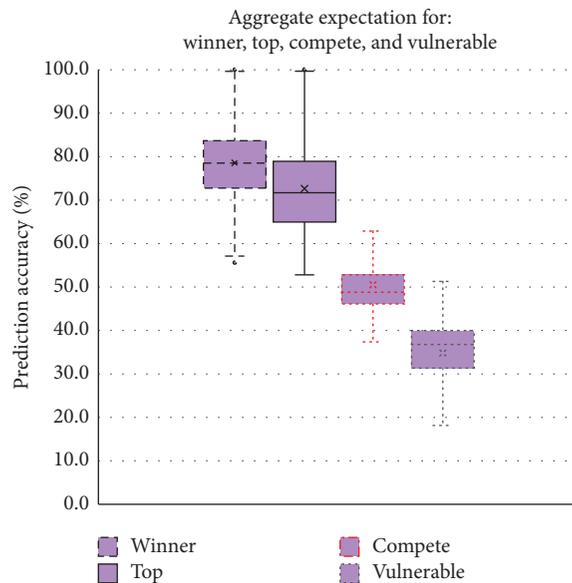


FIGURE 4: Distribution of winning criteria.

seats requires an average prediction accuracy of 82.1% and 82.6%, respectively, while 50 seats require an average prediction accuracy of only 74.4% to win. Figure 8 illustrates the impact of the number of seats combined with decreased diversity of predictions among 95 double loop

learners illustrated in Figure 2 on prediction accuracy. The combined effect of downward pressure on prediction accuracy for 50% open seats and compression due to improved prediction in the double loop learning distribution are apparent.

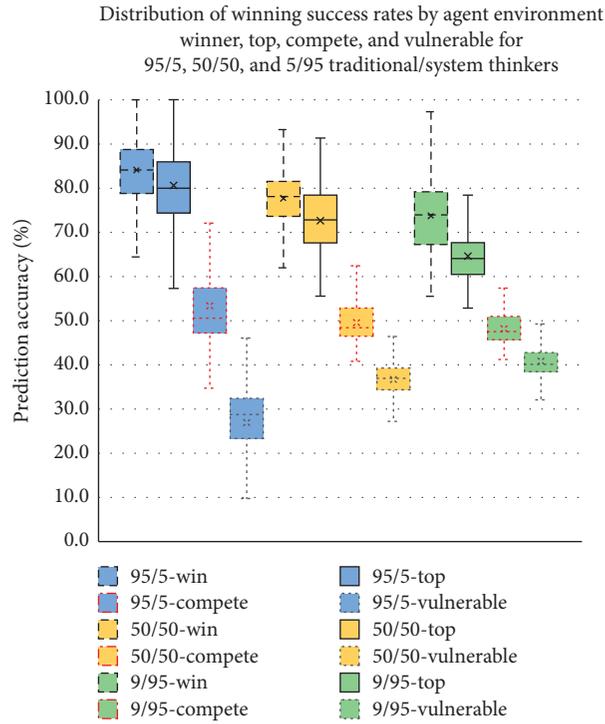


FIGURE 5: Distribution of winning criteria sorted by the agent ratio.

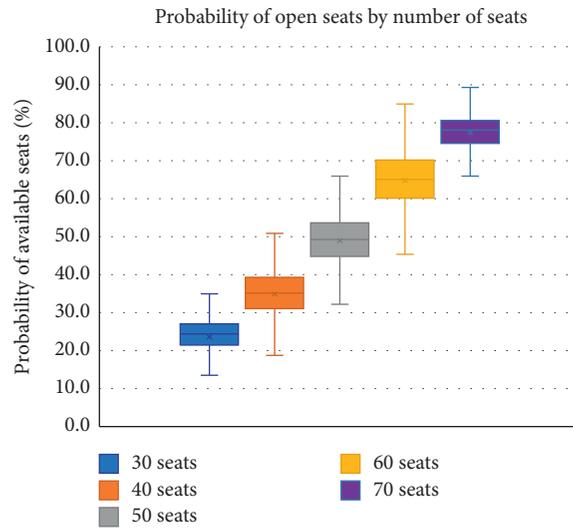


FIGURE 6: Probability of open seats by available seats.

### 3.2. Designed Agents Performance

3.2.1. *Observer Agents.* The size of the sample used by the 2 observer agents—Observe 5 and Observe 30—was chosen to represent 2 levels of commitment by the strategic agent. 5 other patrons is a number that a patron might have casually available, that is, without any serious investment of time. 30 represents a level of observation that an engaged patron wanting to compete well in the game might have.

From the general results in Table 2, even the casual observation of a random 5% of the other players is sufficient to

win more often (2.4%) than the single or double loop learners, and as a Top Performer 19% of the time with a vulnerability of only 3.8%. For a very small investment of effort, this is an impressive strategy. The Observe 30 requires more than just casual work but pays off as a dominant player winning 24.9% and Top performing 48.7% of the time with a vulnerability of 0.7%. Evaluating these 2 strategies across the 2 environment variables—Double loop learners and available seats—reveals an important dependency, as shown in Figure 9. Both of these strategies have an improved performance peak at 50 seats, and

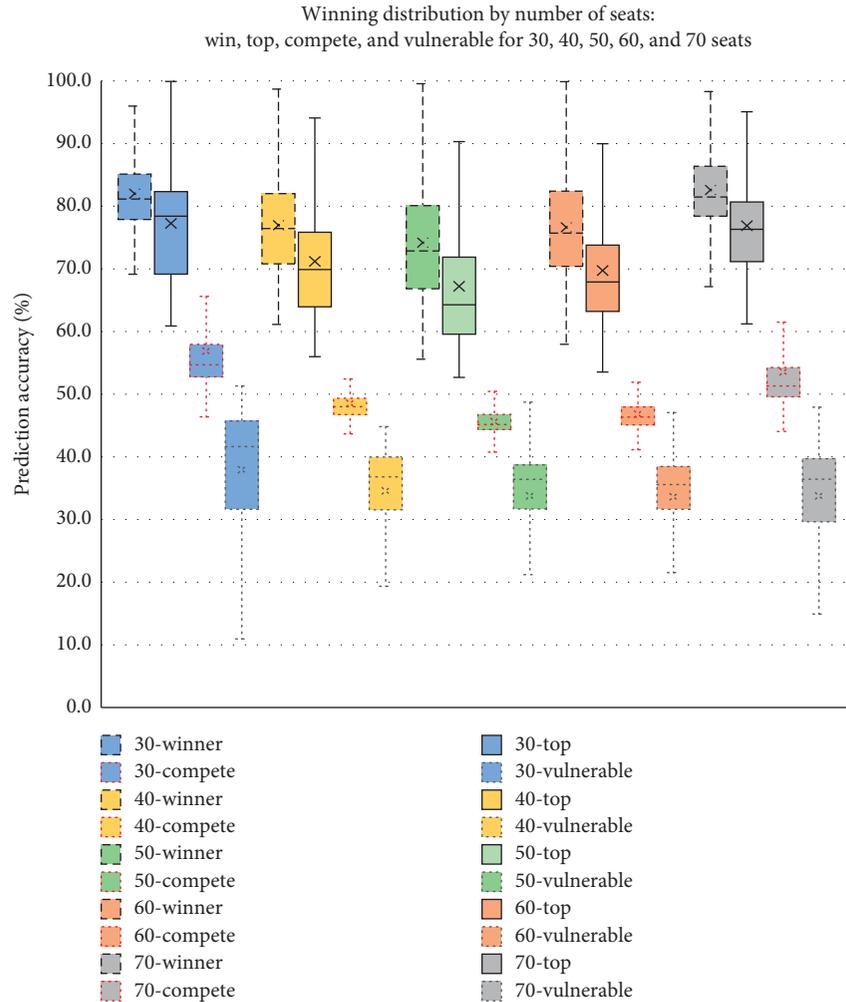


FIGURE 7: Winning prediction accuracy by the number of seats.

the Observer 30 also exhibits a slight performance peak at 50/50 double loop learners. As shown in Table 3, competing for 50 seats in the 50/50 learner environment, the Observer 30 agent wins a remarkable 68% of the time, while the Observer 5 peaks at 5.5% competing for 50 seats in the 5/95 environment. The cost of this high performance in the center is that, in extreme environments of 30 and 70 seats, the winning percentages fall dramatically. In fact, 234 of the total 267 Vulnerable rankings shared by these 2 strategies occur in the single test environment of 95 double loop learners with 70 seats.

**3.2.2. Rational and Rational Insured.** Rational agents use a traditional straightforward averaging of the last 10 weeks to predict the probability of crowding, and they also calculate a utility value for their decision based on the average number of free seats and excess patrons. Rational agents with access to insurance are able to purchase an option to take action after the event while all other types of agents must rely exclusively on prediction. This optionality gives any agent with access to insurance a huge advantage in mitigating their losses. However, the utility of that optionality is determined by the cost of insurance. Table 4 compares an uninsured

agent with 4 different levels of insurance cost. In the case of the most expensive insurance tested, on average, the agents elected to purchase the insurance only about 1 out of every 10 weeks (11% of decisions). Three lower prices were experimentally determined and then used to approximate buying insurance once per month (22%), every other week (43%), and 3 out of 4 weeks (72%). Winning success in the game is inversely correlated with the price of insurance. Decreasing price yields increasing winning success. The least expensive insurance in this simulation won outright 65% of the time.

Evaluating the effect of rationality and insurance across the 2 environment variables—double loop learners and available seats—reveals significant dependencies as shown in Table 5 below. Whereas the Observer 5 and Observer 30 did well in the very competitive environments of 50 seats, rational actors perform best in the high prediction accuracy environments of 30 and 70 seats. These extreme environments, as discussed above, reward agents with a consistent bias to their decisions. If the environments are, on average, crowded (30 seats) or uncrowded (70 seats), then the rational agents averaging of the last 10 weeks will more consistently be correct. Figures 7 and 8 clearly demonstrated that this bias

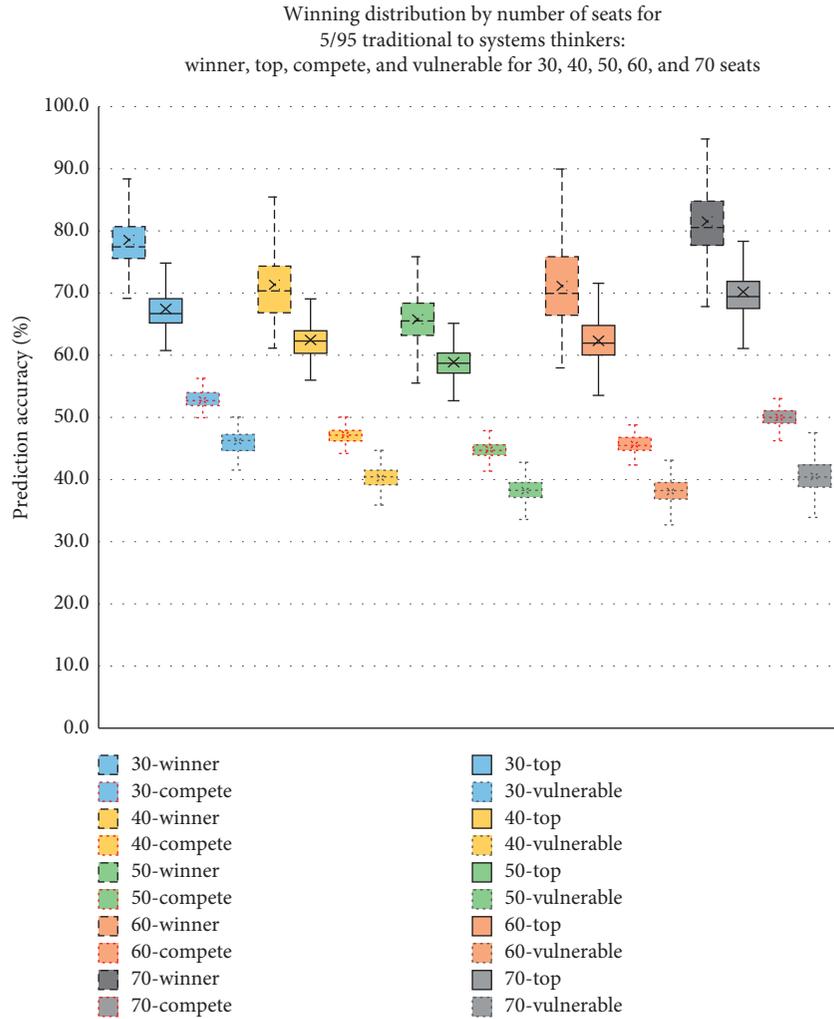


FIGURE 8: Winning prediction accuracy by the number of seats; double loop learners only.

TABLE 2: Observer agents overall winning rate.

	Agents	Win	Top	Compete	Vulner	$P_{Win}$	$P_{Top}$	$P_{Comp}$	$P_{Vln}$
Observe 5	6,000	142	1,107	5,229	225	2.4%	19%	87%	3.8%
Observe 30	6,000	1,494	2,920	5,786	42	25%	49%	96%	0.7%

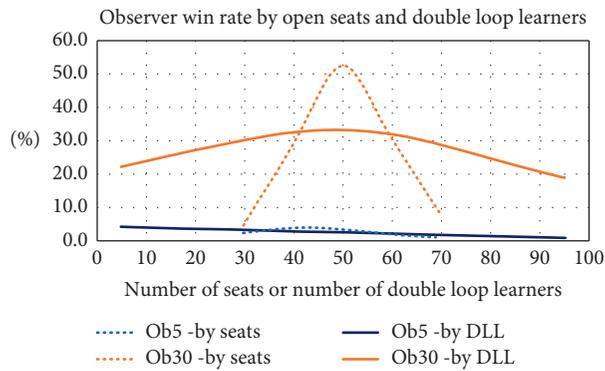


FIGURE 9: Observer winning by seats and double loop learners.

TABLE 3: Observer win rates by seats and environments.

	30 seats	40 seats	50 seats	60 seats	70 seats
DLL	Observer 5—2.4% overall winner				
5	4.3%	5.3%	5.5%	3.5%	1.3%
50	2.3%	3.8%	3.3%	1.5%	1.0%
95	0.5%	2.3%	1.3%	0.0%	0.0%
	Observer 30—24.9% overall winner				
5	3.5%	16.3%	46.0%	34.0%	11.8%
50	6.5%	34.5%	68.0%	47.3%	11.3%
95	4.8%	36.8%	43.8%	34.3%	0.0%

TABLE 4: Uninsured and insured rational agents overall winning rate.

	Agents	Win	Top	Compete	Vulner	$P_{Win}$	$P_{Top}$	$P_{Comp}$	$P_{VIn}$
Not insured	6,000	344	1,195	4,661	235	5.7%	20%	78%	3.9%
11% insured	1500	75	307	1247	46	5.0%	20%	83%	3.1%
22% insured	1500	152	368	1291	39	10%	25%	86%	2.6%
43% insured	1500	451	615	1418	11	30%	41%	95%	0.7%
72% insured (cheap)	1500	967	1146	1496	1	65%	76%	99%	0.1%

TABLE 5: Uninsured and insured rational agents by seats and environment.

	30 seats	40 seats	50 seats	60 seats	70 seats
DLL	Rational, no insurance—5.7% overall winner				
5	0.3%	0.0%	0.0%	1.5%	28.3%
50	6.5%	0.0%	0.0%	0.0%	23.0%
95	24.5%	0.5%	0.0%	0.0%	1.5%
	Rational, insured—27.4% overall winner				
5	4.8%	8.0%	2.8%	9.0%	33.0%
50	31.0%	24.8%	7.5%	27.8%	46.0%
95	58.3%	39.8%	23.0%	38.5%	57.3%

is present in the data. The other important observation here is that the insured agents perform extremely well with increasing numbers of double loop learners.

## 4. Conclusions

Systems Science, along with its expanding subdisciplines of network science, complexity science, dynamic systems, evolutionary science, and more, has brought into sharp focus the flaws in the traditional world model. Systems thinking offers an alternative to these traditional methods by viewing problems, as well as the methods to solve them from new perspectives. The question considered in this paper is whether systems thinking can help an individual successfully solve daily problems, as an alternative to the traditional deterministic or stochastic methods. The answer is yes, systems thinking is a competitive daily strategy for an individual that is at least on par with traditional disciplined strategies and may provide some enhanced resistance to extinction. However, there are 2 observations that temper these results. First, the win rate of any given strategy is highly dependent on the nature of the competitive environment which is a function of the spectrum of strategies of the other agents. Nominally successful strategies can fail significantly

in the wrong environment—and this is true for both traditional and systems techniques. Second, there is a complex relationship between the aggregate of the individual strategies and the collective environment. The behavioral characteristics of the collective environment will evolve over time as a sufficient proportion of the population adopts these adaptive, systems-based strategies. The net gain, loss, or impact for the individual or the collective environment is, therefore, unpredictable as we evolve through these complex mixtures of traditional and systems-based strategies. This simulation illuminates the predisposition of systems thinking—or any other widely applied strategy—to usher in a new set of unintended consequences. Future practitioners may be forced to address those consequences just as modern practitioners must address the unintended consequences of the actions of their forebearers.

A secondary conclusion is about the role of insurance. The initial purpose of the insured rational agent in this study was to provide a competitive high-water benchmark of traditional thinking to challenge the systems thinking agents. However, this strategy (1) significantly influenced the architecture of the overall test, and perhaps more importantly, (2) forced some exploration of the calculation and influence of insurance availability on any strategy. The treatment of the subject herein was not sufficient to draw any conclusions about any insurance methodologies, but it did highlight some important considerations for a study of that nature. Those considerations include (1) the 2 costs of buying and subsequently using insurance, (2) defining an entity that would profit from selling insurance, (3) optimizing the price, (4) optimizing the time and frequency of purchases, and (5) defining a fungible definition of seat value that all agents would utilize. In a perfect world, both traditional agents and system agents would have equal access to insurance. In this architecture, the systems thinking agents' strategies would require a 4-state decision matrix based on some predicted value of the transaction. It is important to

note that the availability of insurance for every agent may have significant impact on the specific numerical distributions of the competition results, but it is not likely to invalidate the conclusion that systems thinking provides useful life strategies since uninsured “systems thinkers” competed well.

Perhaps, the 3 most interesting phenomena to explore in more depth with this model are (1) eliminating and replacing vulnerable agents, (2) randomly changing environments (number of seats) during a simulation run, and (3) algorithms for cooperation among the agents. For vulnerability, are there any strategies that survive the long term if the bottom 15% are replaced after each run or will they all eventually succumb? Randomly changing environments (seats) will stress the strategies that perform poorly in some environments but not others. Which form of cooperation—cliques or a whole group—is most successful? Can this model be used to analyze solutions for the “Tragedy of the Commons?”

### Data Availability

Data were generated using a publicly available ABM simulation: Rand, W., Wilensky, U. (2007). NetLogo El Farol Extension 3 model. <http://ccl.northwestern.edu/netlogo/models/ElFarolExtension3>. Additions to the model are described in the paper.

### Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

### References

- [1] W. B. Arthur, “Inductive reasoning and bounded rationality,” *American Economic Review*, vol. 84, no. 2, pp. 406–412, 1994.
- [2] M. W. Sellers, “Why is systems thinking so rare?” 2017.
- [3] D. P. Stroh, “Systems thinking for social change: a practical guide to solving complex problems, avoiding unintended consequences, and achieving lasting results,” 2015.
- [4] L. B. Sweeney and D. Meadows, “The systems thinking playbook: exercises to stretch and build learning and systems thinking capabilities,” 2010.
- [5] P. M. Senge, “The fifth discipline: the art and practice of the learning organization,” 2006.
- [6] R. L. Ackoff, “Systems thinking and thinking systems,” *System Dynamics Review*, vol. 10, no. 2-3, pp. 175–188, 1994.
- [7] G. Klir, *Facets of Systems Science*, Vol. 7, Springer, Berlin, Germany, 2013.
- [8] D. Doerner, “The logic of failure,” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, vol. 327, no. 1241, pp. 463–474, 1990.
- [9] D. H. Meadows and D. Wright, *Thinking in Systems: A Primer*, Chelsea Green Publishing, Hartford, VT, USA, 2008.
- [10] N. Taleb, *Foiled by Randomness: The Hidden Role of Chance in Life and in the Markets*, Vol. 1, Penguin Random House LLC., New York, NY, USA, 2005.
- [11] N. N. Taleb, *The Black Swan: The Impact of the Highly Improbable*, Vol. 2, Penguin Random House LLC., New York City, NY, USA, 2007.
- [12] N. N. Taleb, *Antifragile: Things that Gain from Disorder*, Vol. 3, Penguin Random House LLC., New York City, NY, USA, 2012.
- [13] H. Lucs, C. O. Aydin, S. Keten, H. I. Ünsal, and A. R. Atilgan, “El farol revisited,” *Physica A: Statistical Mechanics and its Applications*, vol. 346, no. 3-4, pp. 651–656, 2005.
- [14] A. M. Bell and W. A. Sethares, “The El Farol problem and the internet: congestion and coordination failure,” in *Proceedings of the Fifth International Conference of the Society for Computational Economics*, Boston, MA, USA, 1999.
- [15] W. Rand and F. Stonedahl, “The El Farol Bar Problem and Computational Effort: Why People Fail to Use Bars Efficiently,” Northwest University, Evanston, IL, USA, 2007.
- [16] D. Mitsche, G. Saad, and J. Saia, “The power of mediation in an extended El Farol game,” in *Proceedings of the International Symposium on Algorithmic Game Theory*, pp. 50–61, Aachen, Germany, October 2013.
- [17] S.-H. Chen and U. Gostoli, “Coordination in the El Farol bar problem: the role of social preferences and social networks,” *Journal of Economic Interaction and Coordination*, vol. 12, no. 1, pp. 59–93, 2017.
- [18] R. Franke, “Reinforcement learning in the El farol model,” *Journal of Economic Behavior & Organization*, vol. 51, no. 3, pp. 367–388, 2003.
- [19] K. Tuyls, P. J. T. Hoen, and B. Vanschoenwinkel, “An evolutionary dynamical analysis of multi-agent learning in iterated games,” *Autonomous Agents and Multi-Agent Systems*, vol. 12, no. 1, pp. 115–153, 2006.
- [20] S.-H. Chen and U. Gostoli, “On the complexity of the El Farol bar game: a sensitivity analysis,” *Evolutionary Intelligence*, vol. 9, no. 4, pp. 113–123, 2016.
- [21] S. St. Luce and H. Sayama, “Phase spaces of the strategy evolution in the El farol bar problem,” *The 2020 Conference on Artificial Life*, vol. 32, pp. 558–566, 2020.
- [22] W. Rand, *NetLogo El Farol Extension 3 Model*, Center for Connected Learning and Computer-Based Modeling, Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA, 2007.
- [23] U. Wilensky, “NetLogo,” Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL, USA, 1999.