Research Article

Dynamic Decision Making and Race Games

Shipra De¹ and Darryl A. Seale²

¹ *Department of Economics, University of Nevada Las Vegas, 4504 Maryland Parkway, Las Vegas, NV 89154-6001, USA*

² *Department of Management, Entrepreneurship and Technology, University of Nevada Las Vegas, 4504 Maryland Parkway, Las Vegas, NV 89154-6001, USA*

Correspondence should be addressed to Darryl A. Seale; dseale@unlv.nevada.edu

Received 14 June 2013; Accepted 7 July 2013

Academic Editors: I. Ahmad and G. Dósa

Copyright © 2013 S. De and D. A. Seale. 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.

Frequent criticism of dynamic decision making research pertains to the overly complex nature of the decision tasks used in experimentation. To address such concerns, we study dynamic decision making with respect to a simple race game, which has a computable optimal strategy. In this two-player race game, individuals compete to be the first to reach a designated threshold of points. Players alternate rolling a desired quantity of dice. If the number one appears on any of the dice, the player receives no points for his turn; otherwise, the sum of the numbers appearing on the dice is added to the player's score. Results indicate that although players are influenced by the game state when making their decisions, they tend to play too conservatively in comparison to the optimal policy and are influenced by the behavior of their opponents. Improvement in performance was negligible with repeated play. Survey data suggests that this outcome could be due to inadequate time for learning or insufficient player motivation. However, some players approached optimal heuristic strategies, which perform remarkably well.

1. Introduction

A great deal of our understanding of judgment and decision making comes from a body of research that examines the dysfunctional consequences and systematic biases of adopting heuristics or “rules of thumb” in decision making. For example, decision makers (DMs) are known for ignoring base rate information, failing to revise opinions, having unwarranted confidence, and harboring hindsight biases, to name a few [1]. This seems to imply that humans are fairly incompetent beings [2]. Yet while this appears to be true in controlled settings, it is not so in real life. Toda points out that “man drives a car, plays complicated games, and organizes society” [1]. So why is there such a disconnect between experimentation and real-world phenomena?

While it is clear that people do make mistakes and can, under certain situations, exhibit systematic deviations from rational predictions, this research is criticized for concentrating on discrete incidents often lacking any form of meaningful feedback [1]. Critics contend that judgment is best viewed as a continuous and interactive process that enables DMs to cope with their environment. The claim,

made by Jungermann, is that decision makers who appear biased or error-prone in the short run may be quite effective in continuous or natural environments that allow for feedback and periodic adjustment in decision making [3].

Research in dynamic decision making (DDM) is well suited to advance our understanding of judgment and decision behavior in more natural environments characterized by continuous and interactive processes. In order for a task to be considered dynamic, it must be (i) characterized by multiple decisions that are interdependent and (ii) the environment must change as a result of the decision-makers' actions and/or additional external forces [4]. For example, consider a pedestrian walking down a busy street. Her goal is to reach the end of the block, perhaps as quickly as possible, without bumping into other pedestrians. At various points, our traveler may decide to veer right or left, or speed up or slow down to avoid oncoming traffic. These decisions may affect the course of other pedestrians who make similar decisions in an attempt to navigate the busy street. After each course correction, our traveler receives (potential) feedback in the form of new traffic patterns and makes subsequent course corrections toward her goal of reaching the end of the

block. Throughout this process, our traveler is not required to chart a complete course of action prior to initiating her walk, rather she makes a series of adjustments to her intended course of action while keeping the end of the block in sight.

Even as simple a task as walking down the street is clearly very complex to replicate and evaluate in the laboratory. As will be discussed shortly, researchers struggle to find decision tasks that are both dynamic and allow for feasible analysis of subjects' ability to complete tasks and learn from repeated actions. In this study, we introduce a novel DDM environment that satisfies these requirements while avoiding the criticisms that have been directed at previous research, which includes that decision tasks were too complicated for subjects to succeed and improve from repeated action. The simplicity of the decision task in this experiment if coupled with successful subject performance and discernible learning would suggest that under certain conditions, decision makers are capable of approaching optimal predictions.

To better understand DDM research and how the present study connects to the existing body of research, we review the relevant literature in the remainder of Section 1. We introduce the dynamic decision task and characterize its optimal solution in Section 2. Section 3 describes the experimental methods employed. The results, which address how well players performed and whether their performance improved over time, are reported in Section 4. A summary of our aggregate data was reported by Seale et al. The present study, however, (i) examines several new independent variables, (ii) analyzes data by experimental condition, and (iii) reports the results of decision making survey and "notebook" analysis. We discuss the findings of the study and suggestions for future research in Section 5.

2. Literature Review

Initial interest in dynamic decision making began in the early 1960s through the independent efforts of Toda and Edwards. In his seminal paper, Edwards introduced dynamic decision theory by first reviewing static decision theory, in which the DM makes a choice and receives the corresponding outcome or payoff [5]. In contrast, with dynamic decision situations, the DM makes a series of choices, each dependent on the last, and attempts to maximize payoffs in the long run. Due to its complicated nature, the DDM has not received the same attention as static decision making [6].

Brehmer categorizes the previous research into two distinct approaches [7]. The first approach, the individual differences technique, attempts to either predict behavior or "identify the demands" of the tasks performed [7]. This approach involves first separating subjects into two groups: those who succeed and those who do not. The groups are then compared in terms of behavior and psychological test scores that could potentially account for the disparity in performance. Unfortunately, results often fail to produce any significant correlation between the two [7]. Efforts to train subjects to adequately handle complex decision tasks have also proven unsuccessful, suggesting that heuristic competence cannot be engineered or taught in a general sense

[7]. The second approach, the standard method, involves analyzing the specific attributes of the system that could affect subject performance with the objective of understanding how people "develop mental models and formulate goals" [7]. Achieving such a comprehensive picture of decision making requires developing a classification system which can characterize different dynamic decision experiments and allow for more comparability between them. Then, by altering one characteristic at a time, systematically determine which has the greatest effect on performance, a process which carries with it the possibility of developing a general theory.

Gonzalez, Vanyukov, and Martin have revived the taxonomical approach to classify DDM problems, identifying four important characteristics.

- (i) *Dynamics*. The dynamic character of a task speaks of the degree and speed at which the system changes. Since dynamic decisions are made in a specific context and time, this includes whether the system changes endogenously and/or exogenously, as mentioned by Edwards, and whether decisions are made in real time, a criterion added by Brehmer [7, 8]. By real time, we mean that "decision makers are not free to make decisions when they feel ready to do so...they have to make the decisions when the environment demands decisions from them" [7].
- (ii) *Complexity*. Complexity refers to the situation in which decision makers are obliged to keep track of many factors and possibly conflicting goals [7]. Complexity can be difficult to gauge, as it is relative to a specific decision maker with particular cognitive abilities [8]. However, it is characterized by three attributes that work in conjunction: "(1) the number of components in the system, (2) the number of relationships among the components (i.e., the degree of coupling), and (3) the types of relationships among the components" [8].
- (iii) *Opacity*. Opacity measures how visible different aspects of the task are to the DM [8]. A situation is considered opaque if it does not deliberately make its characteristics known but rather requires the decision maker to "form and test hypotheses about [its] state" [7]. As with complexity, it is relative to a specific decision maker because even if certain information about the state of the system may be determined, the decision maker must know how to obtain it [8].
- (iv) *Dynamic Complexity*. Dynamic complexity focuses on the nature of feedback provided by the system [8]. This can be further bisected into the issues of feedback quality and feedback delays. If a system is prone to nonlinearities or side effects, where a deliberate change in one variable leads to unintended consequences in other variables, decision makers may face difficulty in prioritizing goals [8]. Also, if there are significant time gaps between decisions and their outcomes, this can complicate the decision maker's ability to assess the system [7].

An example of a typical DDM experiment is the Beer Distribution Game [9]. In this role-playing experiment, Sterman arranged several teams of four players including a producer, distributor, wholesaler, and retailer [8]. Each participant attempted to “manage a simulated inventory distribution system” [9]. Dynamics in this game were low since state changes occurred only with players’ decisions and players had ample time to make them; likewise, with only three variables (backlog, inventory, and current demand), complexity was also low [8]. On the other hand, because consumers’ demands remained unknown to most players, opaqueness was high, as was dynamic complexity, since feedback delays were frequent and large [8]. Performance in the Beer Distribution Game was generally poor; on average, team costs were ten times greater than the optimal benchmark cost [9]. The study concluded that this was due to misperception of feedback [9]. Specifically, players seemed to attribute the fluctuations in the system to external factors such as customer demand (which was actually constant), instead of the endogenous interaction among the other players [9].

Other experimental studies of DDM have investigated the importance of feedback. Diehl and Sterman had subjects manage inventory while having to contend with fluctuating sales. They varied feedback strength and delays and monitored the effect on decision making processes [10]. If misperceptions were truly the root cause of poor performance, adjusting these variables should have had little effect on deviations from optimality [10]. The results indicated considerably suboptimal performance, which appeared to deteriorate rapidly with time delays and feedback strength [10]. In yet another study, Atkins et al. experimented with feedback methods (i.e., tabular versus graphical data) finding that graphical feedback leads to better performance, although tabular data suggests greater learning [4].

The performance of heuristics is also an important topic in DDM research. Kleinmuntz tested various heuristic strategies in a simulation of medical decision making [2]. He concluded that the success of a heuristic strategy is dependent upon the dynamic characteristics of the task itself (i.e., the availability of feedback and opportunities for taking corrective actions) [2]. He showed that if these criteria were met, it is possible to observe successful performance by using less complicated decision strategies [2]. The research in feedback and heuristics appears to underscore the cognitive limitations of decision makers and to call for less complicated experiments with more opportunities for learning. Thaler predicts that future studies will “[make] their agents less sophisticated and [give] greater weight to the role of environmental factors, such as the difficulty of the task and the frequency of feedback” [11].

Studies in DDM cite a number of limitations and shortcomings in the field, including difficulty of arriving at optimal solutions [12]; insensitivity to deviations from optimal solutions [13]; accurate and timely feedback processes [9]; and the complexity of task environments that (i) yield little understanding of how decisions affect performance objective and (ii) make it difficult to generalize results across experiments [14]. One explanation for the somewhat slow progress of empirical work in DDM is “the difficulty of extending the

experimental methods used to study individual decisions to aggregate, dynamic settings” [9]. This slow progress is echoed by Hey and Knoll who argue that, despite several decades of experimental study, “little is known about the way people actually tackle stochastic dynamic decision problems” [15].

To address these shortcomings, we study DDM in a new paradigm, that of race games with computable decision strategies. By “race,” we imply that players compete to become the first to achieve a designated point threshold. Such a game has low dynamics, complexity, opaqueness, and dynamic complexity as states change only as a consequence of the players’ decisions, variables are few, and feedback is clear and immediate. By “computable decision strategy,” we require that there be an optimal play that can be determined computationally, a key factor since dynamic decision theory focuses on evaluating human decision making against a “series of temporally related decisions with optimal solutions yielded by mathematical models” [1].

A noteworthy difference between studies in race games and the previous research in DDM is the change in the overarching goal from maximizing revenue and/or minimizing costs to that of winning—a key objective in business, sports, politics, and a host of other arenas. Therefore, despite this difference, the race game paradigm can offer various insights into human behavior and adaptability. Its simplicity, in addition to its ability to satisfy the requirements of a dynamic task, makes such a game a rather attractive dynamic decision-making experiment.

3. The Game of Hog

To further understand dynamic decision making in the realm of race games, we study one of the many variants known as Hog. Hog is a race game in which each player has only one roll per turn. However, the player may choose to simultaneously roll as many dice as he pleases. For practical purposes, we typically impose a maximum number of dice, $d_{\max}$. If a one appears on any of the dice, no points are earned for that turn; otherwise, the sum of the numbers appearing on the dice is added to the player’s total score. The first player to reach some designated point threshold wins the game.

The beginnings of a sample game may be found in Table 1. Player 1 begins to play and decides to roll four dice. Since he does not roll any ones, his score is the sum of the numbers which do appear on the dice, in this case 16. It is then Player 2’s turn, and he also chooses four dice. Not as lucky as his opponent, Player 2 receives zero points for his turn for having rolled a one. The play returns to Player 1 who chooses fewer dice; he again obtains a positive turn total which is added to his overall score, now 25. In his next turn, Player 2 finally makes the scoreboard and after another round of play even takes a slight lead at 27 to 25.

Because players must state at the outset of each turn, the number of dice they will roll can be compared to the optimal number at every turn to determine if the decision is optimal, conservative (fewer than the optimal number of dice), or aggressive (greater than the optimal number of dice). Such information is valuable in determining trends in behavior.

TABLE 1: Hog sample game: first six turns.

Turn	Quantity of dice	Numbers appearing on dice	Turn total	P_1 's score	P_2 's score
P_1	4	5, 6, 3, 2	16	16	0
P_2	4	2, 1, 1, 4	0	16	0
P_1	3	4, 2, 3	9	25	0
P_2	5	5, 2, 2, 4, 6	19	25	19
P_1	8	6, 4, 1, 1, 4, 4, 1, 3	0	25	19
P_2	2	5, 3	8	25	27

3.1. Origins and Optimal Policy. The roots of Hog may be traced back to a 1993 publication by the Mathematical Sciences Education Board. The qualities that make Hog ideal as a dynamic decision-making task are that the “rules of the game are straightforward,” yet, “optimal strategies are not at all obvious,” and participants “will not come to the Hog Game task with an a priori idea of what is “supposed to happen” [16].

The literature related to Hog tends to focus on optimal performance in regard to maximizing the expected value of points, which can be accomplished by rolling either five or six dice as indicated by Figure 1. It is apparent that players face an inherent tradeoff when choosing the number of dice to roll; while increasing the quantity of dice raises the average nonzero score, it necessarily decreases the probability of achieving it.

Perhaps not immediately obvious is the realization that in Hog, maximizing points and maximizing the probability of winning are not the same thing. To illustrate, consider an extreme situation. Say it is Player 1's turn and both he and his opponent are tied at 99 points each. In this situation, Player 1 should clearly roll only 1 die since any quantity larger than that would merely decrease the probability that he will obtain a nonzero score. From this example, it is clear that determining the optimal number of dice with respect to maximizing the probability of winning depends on a player's score and his or her opponent's score.

Provided that the goal threshold is 100 points, the optimal solution for Hog is determined in the following manner. We adopt the notation used by Neller and Presser, who originally solved the game. Let $\pi(d, k)$ be the probability of rolling a score of k points with d dice and allow $P_{i,j}$ to denote the probability that a player with i points (Player 1, who plays optimally) will win, given that his opponent (Player 2, who is also playing optimally) has j points [17]. If $i \geq 100$, then $P_{i,j} = 1$, since Player 1 has achieved enough points to win. Similarly, if $j \geq 100$, then $P_{i,j} = 0$, since Player 2 has won. However, in general, when $0 \leq i < 100$ and $0 \leq j < 100$, we know that Player 1's optimal choice will be the quantity of dice d , $0 < d \leq d^*$, which will maximize the expected probability of winning. Here $d^* = \min\{d_{\max}, \lceil(100 - i)/2\rceil\}$, where $d_{\max}$ is an artificial limit on the quantity of dice to be rolled. No rational player would wish to roll more than $\lceil(100 - i)/2\rceil$ dice, since any nonzero score with this quantity would ensure a win, and a greater number of dice would merely decrease the probability of obtaining a nonzero score. For each d , the expected probability of winning is determined by the summation of the probability of rolling each possible

FIGURE 1: Expected scores for rolling d fair, six-sided dice, $1 \leq d \leq 25$.

score times the probability that Player 2 will not win by rolling optimally in his following turn, that is,

$$P_{i,j} = \max_{0 < d \leq d^*} \sum_{k=0}^{6d} \pi(d, k) (1 - P_{j,i+k}). \quad (1)$$

Although Neller and Presser have already solved Hog, we provide a more thorough and efficient calculation of the solution in Appendix A. We offer a less computationally intensive but equivalent calculation of $\pi(d, k)$ and present a proof of its correctness, an exercise that Neller and Presser omit. Specifically, Neller and Presser's solution for $\pi(k, d)$, seen below,

$$\pi(d, k) = \begin{cases} \frac{1}{6}, & d = 1, k \in \{0, 2, 3, 4, 5, 6\}; \\ 0, & d = 1, k \notin \{0, 2, 3, 4, 5, 6\}; \\ \pi(d-1, 0) + \frac{1}{6} \sum_{k=2}^{6(d-1)} \pi(d-1, k), & d > 1, k = 0; \\ \frac{1}{6} \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r), & \text{otherwise,} \end{cases} \quad (2)$$

does not take advantage of a great number of “impossible” point values [17]. Consider the 3750 unique combinations of d and k when $0 < d \leq 25$ and $0 < k \leq 150$. In Case 2, Neller and Presser indicate that if $d = 1$ and $k \notin \{0, 2, 3, 4, 5, 6\}$, the value of $\pi(d, k)$ must be zero, which is clear since no other value can be obtained with 1 die. However, for any given d , there are also upper and lower limits of $2d$ and $6d$, respectively, for which any other point value (other than zero, which is handled in Case 3) cannot be obtained. This means that in the Neller and Presser solution, there are 2256 d and k pairs of $\pi(d, k) = 0$ that are captured in Case 4 through a summation of zeros, which is a significant waste of computing time. Also, Case 3 can be improved with a closed form solution which eliminates the summation calculation for the 24 d and k pairs that fall in that category. Thus, we present the following improved solution:

$$\pi(d, k) = \begin{cases} \frac{1}{6}, & d = 1, k \in \{0, 2, 3, 4, 5, 6\}; \\ 0, & 0 < k < 2d \text{ or } k > 6d; \\ \frac{5\pi(d-1, 0) + 1}{6}, & d > 1, k = 0; \\ \frac{1}{6} \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r), & \text{otherwise.} \end{cases} \quad (3)$$

Furthermore, our use of d^* in the calculation of $P_{i,j}$ is an even greater improvement over the Neller and Presser solution that instead uses $d_{\max}$; while the difference in the actual calculation time of $\pi(d, k)$ for $d_{\max} = 25$ may be negligible, the time saved in the computation of (1) is quite considerable. Of the 9801 game states, d^* is less than $d_{\max}$ 4752 times. Clearly, the savings depend on the game state and increase with $i + j$. If we use $d_{\max}$, we must calculate the summand in (1) a total of 245,025 times to determine all game states. With d^* , we reduce it to 185,625. Finally, while Neller and Presser simply present their solution, we also specify two different methods for its determination.

The optimal roll decisions for the two-player game of Hog are presented in the graphical form in Figure 2. Player 1's score runs along the i -axis, while Player 2's score is on the j -axis. The d -axis indicates the number of dice that Player 1 should roll, given that the game state is (i, j) . The solution is particular to $d_{\max} = 25$, and Neller and Presser informed us that the optimal solution remains the same for all $d_{\max} \geq 26$ [17].

4. Experiment

To secure ease and accuracy of data collection and to avoid the cumber of real dice, we opted to devise software to allow two subjects to anonymously play the game via computer network. At any time, a player can see the number of points he and his opponent have earned during the game. When it is his turn, the player chooses the number of dice he would like to roll by moving a slider to the desired quantity, and both he

FIGURE 2: Optimal solution for the two-player game of Hog and 25 dice maximum.

and his opponent can view the numbers rolled. Players were informed that all simulated dice were six-sided and fair.

We used two different treatments in the study. In the first treatment, both players began the game with zero points and one of them was randomly chosen to go first. The game ended when one player reached 100 points, and the loser of the game played first in the next game, if any. In the second treatment, one of the players was randomly chosen to start the game with fifty points, while his opponent, who was chosen to play first, began with zero points. Thereafter, players alternated who got the 50-point advantage in each game. Note that the optimal decision does not depend on how players arrived at a particular game state. By artificially placing players in a 50-point advantage or disadvantage, we hoped to generate additional observations, where players were substantially ahead or behind their opponent's score, and to eliminate any biases which may develop from playing the initial stages of the game.

4.1. Methods. Subjects were recruited from two sections of a Principles of Management and Organizational Behavior course at a large southwestern university and were provided with both a monetary and an academic incentive (extra credit) to participate. Students in the class received a flier advertising the opportunity and were asked to sign up for a particular session via E-mail correspondence with the experimenter.

Subjects were invited into the computer laboratory and seated at any open workstation. They were told that they would be randomly paired to play a series of five of the same dice game against someone else in the room. They were provided with ample time to read the informed consent form and experimental instructions. Experiment proctors were available to answer questions both before and throughout the duration of the games. Upon completion, players were asked to fill out a questionnaire, after which they received payment for their participation before exiting the laboratory.

The questionnaire was used to collect demographic information and other data that might help explain subjects' performance. It included questions designed to assess knowledge of probability theory and to discover other qualities and personality traits such as risk tolerance. We developed a risk

score per player using a publicly accessible risk scale provided by the International Personality Item Pool (IPIP), an agency specializing in the development of personality differences. We adopt their risk-taking scale which is determined by a set of 10 Likert-style questions; these questions are similar to those used by the Jackson Personality Inventory which is considered a psychometrically sound measure of personality. The questions were purported to have an alpha of 0.78, indicating a high level of internal consistency. Our analysis of data obtained from all 122 subjects yielded less reliability, that is, an alpha of 0.7071 after removing one question. Although such an alpha is not ideal, a value ≥ 0.70 is considered acceptable [18].

Data was collected over the course of five experimental sessions, two of treatment 1, collectively consisting of 60 subjects, and three of treatment 2, providing another 62. Subjects ranged from age 19 to 62 (on average, 25) with slightly more males than females. They earned \$4 per game won plus a \$5 participation bonus. However, although not announced in advance, because the sessions lasted an hour or more, subjects were paid a minimum of \$10 regardless of their actual earnings. Earnings spanned \$10 to \$25 with the vast majority earning \$13 or \$17 (for winning 2 or 3 games, plus the \$5 participation bonus).

During all sessions, keyboards were removed from each computer station so that players would not be distracted or have access to any calculator, spreadsheet, or internet resource. However, each subject was given a piece of paper on which he could record any information he thought might be important or advantageous for successful performance. Subjects could request more paper if necessary. This paper was collected with each player's survey and later examined. Although qualitative analysis is limited in the extent that writing down a strategy does not guarantee that it is carried out, experimenters have found "notebook analysis" to be a "valuable addition to the researcher's tool kit of process methods" [10].

5. Results

In this experiment, there are two main issues we wish to address. (1) How do players generally perform? (2) Does player performance change over time? First, we define a measure of performance for capturing departures from optimality.

Let δ be defined as $\delta := A - O$, where A is the actual number of dice chosen by a player at a particular game state and O is the optimal number of dice that the player should have chosen at that game state to maximize the probability of winning. We can calculate an average δ per player to capture his overall performance, or calculate an average δ per player, per game, to observe his performance over time. The smaller the $|\delta|$, the more optimal the player. A shortcoming with this metric, however, is that taking the average causes us to lose information. A player who makes both conservative ($A < O$) and aggressive ($A > O$) roll decisions may average a δ close to zero, indicating an optimal player, when in fact δ per decision could have been quite large.

FIGURE 3: Treatment 1: percentage of rolls characterized by delta.

We address this shortcoming by taking the magnitude of delta. Then we must consider x and $-x$ to be equally suboptimal for we cannot discriminate between the two. While this does correct the difficulty at hand, $|\delta|$ has its own flaws. Specifically, it is an ordinal measure only; being two dice away from optimality is not twice as "bad" as being one die away from optimality. Furthermore, a difference of x dice from the optimal solution could have different consequences depending on the state of the game.

The remaining option for a metric is α , defined as $\alpha := P_O - P_A$, where P_A is the probability of winning with the actual number of dice chosen by a player at a particular game state and P_O is the probability of winning with the optimal number of dice. Then, α is a positive number which provides a clearer view of how suboptimal a particular decision is. Despite this attribute, even α is not a faultless measure. When we scrutinize a particular α , we have no idea how many choices of dice were available between it and that which would lead to an optimal α score of zero.

While none of our metric choices are perfect, each has a particular advantage. From δ , we can capture direction of the deviation from optimal; from $|\delta|$, we can observe the incremental degree of the deviation (especially from the perspective of the player); from α , we can obtain the actual consequence of the deviation at a particular game state. We will thus require the use of δ , $|\delta|$, and α at various stages of our analyses depending upon the nature of the question we are attempting to answer.

5.1. Performance. In Treatment 1, players made a total of 3,608 roll decisions, 565 of which were optimal. Figure 3 shows the aggregate distribution of all roll decisions for all players from Treatment 1. The horizontal axis indicates δ , while the vertical axis describes the percentage of the total rolls marked by δ . Similarly, players in Treatment 2 made 411 optimal decisions out of a total of 2,346. The distribution for Treatment 2 may be found in Figure 4.

A Q-Q plot of both distributions indicates that delta for both treatments is approximately normally distributed. We thus calculate a t -statistic for each individual player's average delta to test for statistically significant departures from optimality. If significantly less than zero at the 5%

FIGURE 4: Treatment 2: percentage of rolls characterized by delta.

level, we classify the player as conservative. If significantly greater than zero at the 5% level, we classify the player as aggressive. The remaining players are considered neutral. Given the shortcoming with the δ metric discussed earlier, we use it only to determine a direction and make no assessments regarding the degree of the conservativeness or aggressiveness displayed by players. In Treatment 1, 33 players (55.0%) were conservative, 19 (31.7%) were neutral, and 8 (13.3%) were aggressive. Likewise, in Treatment 2, 38 players (61.3%) were conservative, 22 (35.5%) were neutral, and only 2 (3.2%) were aggressive. Thus, both treatments indicate that the majority of players were conservative in their observed decisions.

Another interesting indicator of overall performance arises from an analysis of player awareness of the discrepancy between his own and his opponent's points totals. From our examination of the optimal solution, we understand that a player must continually raise his choice in number of dice if he falls behind his opponent. Therefore, we take the aggregate data from each treatment and partition it at 20 point intervals in terms of how far behind or ahead of his opponent a player is at a particular game state. For example, $(-60, -40]$ includes the decisions from all game states in which any player is at least 40 points behind his opponent but at most 59 points behind. Likewise, $(20, 40]$ captures all game states in which a player leads his opponent by at least 21 points but not more than 40 points. For each interval, we calculate the average optimal number of dice for the given game states and determine the actual average number of dice rolled for those same game states.

A graph of this data for Treatment 1 and Treatment 2 is available in Figure 5. The dark markers represent the optimal averages, while the white markers designate the observed averages. We use a one-way analysis of variance (ANOVA) to confirm the downward trend in average number of dice as a player catches up to and exceeds his opponent's score. The results suggest that players took the game state into account when making decisions. Although choices may not have been optimal, they follow the optimal trend. Responses from survey data confirm this phenomenon. When asked about their strategies, the majority of subjects indicated that they consciously increased their choice in number of dice when they found themselves falling behind.

Finally, we assess how observed and optimal play differed. Here, we focus on players' average $|\delta|$ across all five games. Results for both treatments appear in Figure 6. To obtain discrete numbers, we round down to the nearest integer so that $|\delta| = 1$ is an abbreviation for $|\delta| \in [1, 2)$, $|\delta| = 2$ is an abbreviation for $|\delta| \in [2, 3)$, and so forth. In both treatments, the majority of players were on average only two dice away from optimality. Note that we cannot comment on how "bad" two dice away from optimality really is (in terms of the difference in the probability of winning) because that depends on the game state in which it occurred.

Thus, overall, we find that players are far more likely to be conservative than aggressive in play, in which they adjust their decisions based upon the state of the game, and on average, they tend to be 2 dice away from the optimal solution. Next, we examine if observed behavior approached optimal play with experience in the game of Hog.

5.2. Improvement. Given the complex nature of the optimal solution to Hog, one would not expect players to know how to compute it. However, if a player's decision approaches the optimal policy with repeated play, the player is "improving" his chances of winning. Detection of improvement over the course of the five games might signify that feedback and experience aided in learning.

To test for improvement with repeated play, we conducted a one-way repeated measures (ANOVA) for each condition, with mean α per game as the repeated measure. Mean α serves as the best measure for observing differences in performance since it quantifies the consequence of each decision. The repeated measures (ANOVA) for Treatment 1 indicate a statistically significant difference in performance across games ($F = 4.133$, $P = 0.004$). The ANOVA for Treatment 2 failed to show improvement in play ($F = 0.703$, $P = 0.572$). This lack of improvement for subjects in Treatment 2 may be due to the experimental design, where players alternated turns, starting some games with zero points, and some with fifty points. The mean α values per game are displayed in Figure 7. The results for Treatment 1 are shown in the left panel; the results for Treatment 2 appear in the right panel.

5.3. Survey Data. The quantitative data studied thus far paints only a partial picture of decision making in Hog. Although it reveals the actual decisions that subjects made, alone it explains little about why or how these decisions were made. To better understand the observed behavioral patterns noted previously, we turn to qualitative data gathered from a survey administered to all subjects.

Recall that the postexperimental survey collected general demographic data, as well as subjects' risk tolerance and understanding of probability concepts. Including both treatments, 122 subjects participated in the survey; 119 answered all of the questions found in Table 2, a response rate of 97%. Although we asked players the number of classes in which they have learned some probability theory/statistics, we specifically tested their knowledge by asking three straightforward multiple-choice questions pertinent to the game.

FIGURE 5: Average observed and optimal numbers of dice by score discrepancy.

FIGURE 6: Per player frequency of the average magnitude of differences between optimal and observed numbers of dice.

The vast majority of respondents correctly identified $5/6$ as the probability of not rolling a 1 using one die. However, less than half of the subjects knew that this number must be squared to obtain the probability of not rolling any 1s with two dice. Although the correct answer was the most popular choice, a third of the subjects appear unaware of the fact that increasing the number of dice rolled will increase

the probability of rolling a 1. Slightly more than $3/4$ of all subjects seem to understand the concept of independent events. To assess if additional classes taken in probability theory led to increased knowledge of probability theory, we tested the correlation between the number of classes taken and the number of questions answered correctly in the survey. The correlation coefficient ($P = 0.067$) was not significant

FIGURE 7: Subject performance across games.

at the 0.05 level, connoting negligible correlation between the number of probability classes taken and knowledge of probability concepts.

Another important component of performance in Hog is a person's risk tolerance. An extremely risk-averse or risk-seeking player may fail to achieve the optimal policy due to these personality traits. In order to isolate the various effects contributing to performance, we must control for a person's risk tolerance. Subjects' risk scores are reported in Table 3. The risk score is merely a relative measure; a player with risk score 1.5 is more risk averse than a player with risk score 4.8. The remainder of Table 3 notes players' responses to a number of additional characteristics which have the potential to affect performance.

Finally, the subjects were asked to hand in their scratch paper with their surveys. Two-thirds of subjects returned blank sheets of paper, suggesting that the majority of subjects felt there was little that they needed to keep track of in order to succeed in the game. Of those which were not blank, many consisted mainly of doodles. However, a handful of subjects wrote out several calculations such as different probabilities, and even more subjects kept a running tally of the numbers he and/or his opponent rolled, perhaps to determine their own set of *de facto* probabilities. Of interest is the observation that if we separate the use of scratch paper by treatment, 38% of subjects in Treatment 1 "used" their scratch paper, while only 27% of the subjects in Treatment 2 "used" theirs. This might suggest that because of the large discrepancy in points, players in Treatment 2 felt that the outcome of the game was inevitable, and there was little the player from behind could do to bridge the point gap.

To determine if any of the aforementioned characteristics had a perceivable effect on performance, we ran a multivariate regression with players' average α as the dependent variable and 10 independent variables: opponent's alpha, treatment type, age, gender, risk score, probability knowledge, and whether the player gambles, feels lucky, enjoyed the game,

or used his scratch paper. Treatment, gender, gamble, lucky, enjoy, and scratch are all indicator variables; "1" when the answer of the question is "yes", and "0" for Treatment 2. Results are presented in Table 4. Since not every participant answered every survey question, the regression is based upon 88 observations. Each coefficient is reported to the right of its regressor, with the standard error directly beneath it in parentheses. We star those coefficients which are significant at the 0.05 level.

Although the regression coefficients are jointly significant, only Opponent Alpha and Gamble have statistically significant coefficients at the 0.05 level. This implies the more optimal the performance of an opponent, the more optimal the performance of a player. Several written responses support this finding; a dozen subjects either admitted to copying their opponents' strategies or noted that their opponents copied them. Furthermore, a player who admits that he gambles performs worse than a player who does not. All other variables have no statistically significant effect on performance. Together, the two statistically significant predictors of player performance in Hog account for 25.27% of the variation found in average α .

5.4. Heuristic Solutions. With no expectation that players can compute the optimal policy to the game of Hog, we might expect them to adopt heuristic strategies. We examine below six types of heuristic strategies. Within each type, we consider twenty-five unique strategies that correspond to the number of dice that a player might select on her opening turn. For each strategy, we estimate the probability of winning by simulating the decision heuristic for one million repeated games and compare the proportion of games won to the 0.53 threshold—the probability that Player 1 wins if both players behave optimally.

One type of heuristic that a player might adopt is the *Consistent Roll* strategy. Using this heuristic, the player decides to roll h dice on each turn, regardless of past outcomes or

TABLE 2: Aggregate survey data: knowledge base questions.

Response rate: 0.97; correct answers are starred				
(i) How many classes (post high school) have you taken in which you have learned some statistics and/or probability theory?				
None	One	Two	Three	Four or more
0.04	0.15	0.41	0.29	0.11
(ii) What is the probability of not rolling a 1 using one die?				
1/2	2/3	3/4	4/5	5/6*
0.02	0.02	0.00	0.01	0.95
(iii) What is the probability of not rolling any 1s using two dice?				
1/2	7/12	2/3	25/36*	5/6
0.05	0.10	0.13	0.39	0.33
(iv) If you rolled a 1 during your last turn, how would that affect the probability that you will roll a one during your current turn?				
More likely to roll a "1"	Has no effect on the current turn*	Less likely to roll a "1"		
0.19	0.78	0.03		

TABLE 3: Aggregate survey data: individual characteristics.

Response rates are provided in parentheses.			
(i) Risk-taking score (scale: 1–5, the lower the score, the more risk averse.) (1.00)			
[1, 2)	[2, 3)	[3, 4)	[4, 5]
0.03	0.30	0.56	0.11
(ii) Do you gamble? (0.99)			
Never	Sometimes	Often	
0.36	0.59	0.05	
—			
Yes	No		
(iii) In general, do you feel lucky? (0.99)			
0.53	0.47		
(iv) Did your strategy change from game to game? (0.98)			
0.59	0.41		
(v) Did your strategy ever change based upon the behavior of your opponent? (0.99)			
0.44	0.56		
(vi) Do you feel your performance improved from game to game? (0.93)			
0.36	0.64		
(vii) After playing the game once did you want to play again? (0.92)			
0.81	0.19		
(viii) Did you enjoy playing the game? (0.93)			
0.81	0.19		

the number of points the player leads or trails her opponent. However, with this and all subsequent heuristic strategies, we assume that the player is somewhat forward thinking; as she approaches the game-winning threshold of 100 points, she amends her decision and selects d dice, where $d = \min\{h, \lceil(100 - s)/2\rceil\}$, where s is the player's score. For example, if the player is using a *Consistent Roll* strategy with $h = 6$ dice and $s = 94$ points, she selects $d = 3$, rather than $d = 6$, as no more than three dice are needed to reach the game-winning threshold. The performance of this type of heuristic is displayed as the solid line in the top panel of Figure 8, where the horizontal axis marks the possible choices of d , for $0 < d \leq 25$, and the vertical axis reports the proportion of wins for each value of d . The dashed, horizontal line at 0.53 shows the proportion of wins expected from optimal play.

With a judicious choice of d ($4 < d \leq 8$), the *Consistent Roll* strategy wins between 45% and 50% of the games played. As d increases beyond eight, the probability of winning gradually declines. With $d = 25$, the probability of winning is approximately 12%. Although this winning percentage may seem high, note that an opponent playing the optimal strategy would need, on average, five successful turns to win the game. Thus, our heuristic player would have a number of opportunities (turns) to reach the 100-point threshold.

The top panel of Figure 8 shows the results of two other heuristic strategies. The *Consistent Roll Plus* strategy is similar to *Consistent Roll*, except the heuristic selects $d + 1$ dice if the player is behind her opponent. Also, similar is the *Consistent Roll Plus/Minus* strategy. Here, the heuristic selects $d + 1$ dice if behind and $d - 1$ dice if ahead. Both strategies

FIGURE 8: Various heuristic performances in 1M games.

TABLE 4: Regression results.

Regressor		Regressor	
Opponent alpha	0.522* (0.120)	Gamble	0.00338* (0.00149)
Treatment	0.000618 (0.00137)	Lucky	-0.00117 (0.00137)
Age	-0.0000612 (0.000119)	Enjoy	-0.00228 (0.00220)
Gender	0.00108 (0.00137)	Scratch	-0.000106 (0.00158)
Risk	-0.00151 (0.00118)	Constant	0.0110 (0.00631)
Probability	-0.0000378 (0.000696)		
Dependent variable: alpha			
Observations: 88, $\bar{R}^2$: 0.2527, SER: 0.0000349			

*indicates significance at the 5% level.

adjust to d as they approach the winning threshold. Notice that the outcomes between these three types of strategies are inconsequential. Adjusting ± 1 die hardly matters in the proportion of games won; carefully selecting d does.

The bottom panel of Figure 8 displays the results of three other types of heuristic strategies. Using the *Turn Plus/Minus* strategy, the heuristic player selects d dice on her first turn of each game, then $d + 1$ dice on her next turn, if successful (did not roll a 1), and $d - 1$ dice if unsuccessful. Using the *Copy Opponent* strategy, the heuristic player selects d dice on her first turn, and then simply copies her (optimal) opponent on successive turns. Finally, the *Copy Opponent Plus/Minus* strategy selects d dice on the first turn, then copies

the opponent's selection of dice, making an adjustment of $d + 1$ dice if behind and $d - 1$ dice if ahead. This adjustment process continues for each heuristic strategy until the game is won or lost, provided $0 < d \leq 25$. Again, with a careful selection of d , each of these three strategy types performs well, winning nearly 45% of games played, depending on initial choice of d . Notice, however, that the two heuristic strategies that copy their (optimal) opponent perform better than the *Turn Plus/Minus* strategy over a broader (initial) choice of d . Thus, while it may be advantageous to adjust ones' strategy, depending on the success of previous outcomes, copying a successful strategy and allowing for additional adjustment based on whether the player is ahead or behind ones' opponent can be even more successful.

6. Conclusions

In this study, we addressed some of the criticisms directed at the complicated and often haphazard nature of dynamic decision making in experimental tasks. We introduce Hog as a method to study decision making under uncertainty because of its simplicity, timely feedback, and clear objectives. In addition, the game of Hog has an analytical solution that allows us to compare observed behavior to optimal benchmarks. Yet despite the simplistic nature of the game, immediate feedback, and availability of fairly successful heuristic strategies, player success was minimal and performance showed only modest improvement with repeated play. Likely causes, which we will discuss later, include (1) insufficient time for learning and (2) the slow pace of the game that at times may have impacted player motivation.

First, in both treatments, we found that players were sensitive to the number of points they were ahead or behind their opponents. Consistent with the optimal policy, players chose

to roll fewer dice when they were ahead of their opponent, and more dice when they were behind. Improvement in play, however, was only noticed in Treatment 1. This is likely due to the limited number (five) of repeated games. As Doubllass notes, without “adequate time to integrate, experiment, and reflect,” learning may fail to occur [19]. Considering that subjects played only five games, one immediately after another, it is possible that there was not sufficient time to allow for substantial learning. This sentiment was explicitly expressed by a survey respondent who wrote, “I do not think there were enough number of games for me to practice and learn and improve upon.”

Second, the survey data reveal that the most frequent complaint about the game was the pace. The Hog software was programmed to reveal the outcome of a roll decision one die at a time. Thus, if a player chose a large number of dice, he and his opponent would have to wait for each die to appear in a sequential manner. Although some players enjoyed the suspense (e.g., “watching those dice populate was nerve racking”), many felt it was too slow. Unfortunately, this might have had a negative effect on performance. Notice from Tables 3 and 4 that there were players who sometimes rolled 24 dice fewer than the optimal strategy. On first glance, we might assume they were poor decision makers, but actually they could have been reacting to the time cost of the game. If the optimal solution calls for all 25 dice, the player must be severely behind his opponent; knowing that the chances of winning are low, the player might have chosen one die just to end the game sooner. One player suggested “if you hit a 1 the game should not kept[sic] track of dice just gone to the next turn.” While this option is not appropriate since it restricts the quantity of feedback received by the players, if replicating the study, we might choose to display all dice simultaneously to speed up the game.

Another change in the design of the study would be to have subjects play against the computer (which we can guarantee will always make the optimal decision) instead of another individual. When determining the solution to Hog, it does not matter how a particular game state is reached; however, we do assume that the opponent will be playing optimally from any given game state onwards. Since opponents were not playing optimally, it is not clear how poorly players really performed. Further investigation in this area is warranted since applying our decision data to an opponent-dependant optimal strategy might alter our conclusions regarding player performance. Players may perform either better or worse than what is revealed when violating the optimal opponent assumption.

Given the limited evidence of subject learning, and design considerations noted previously, it may be difficult to characterize Hog as an ideal experimental task. However, it remains one of the few dynamic decision tasks we have encountered that scores well on complexity and opaqueness and provides clear and unambiguous feedback to experimental subjects. Future research could build on our results by (i) increasing the number of repeated trials, thus, giving subjects more opportunity for learning; (ii) varying the nature or quality of feedback (i.e., reporting the number of points an opponent earned, but not the number of dice selected), to determine

how this impacts decision behavior; (iii) allowing for simultaneous play, where each subject anonymously selects the number of dice to roll, thus making the game even more competitive; or (iv) by examining other individual traits that might offer more value in predicting behavior in dynamic decision tasks. We believe that there are many directions and opportunities for further research using this simple DDM paradigm.

Appendices

A. Obtaining a Score k , with d Dice

Let $0 < d \leq d^*$ and $0 \leq k \leq 6d^*$ so that $\pi(d, k)$ is the probability of rolling a score of k with d fair, six-sided dice. Neller and Presser offer the following as the solution to $\pi(d, k)$ [17]:

$$\pi(d, k) = \begin{cases} \frac{1}{6}, & d = 1, k \in \{0, 2, 3, 4, 5, 6\}; \\ 0, & d = 1, k \notin \{0, 2, 3, 4, 5, 6\}; \\ \pi(d-1, 0) + \frac{1}{6} \sum_{k=2}^{6(d-1)} \pi(d-1, k), & d > 1, k = 0; \\ \frac{1}{6} \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r), & \text{otherwise.} \end{cases} \quad (\text{A.1})$$

We suggest an improved solution which presents a closed form for the third case and moves several members from group four to group two, thereby, reducing the number of calculations which must be performed.

Theorem A.1. *Let $0 < d \leq d^*$ and $0 \leq k \leq 6d^*$. Then $\pi(d, k)$, the probability of rolling a score of k with d fair, six-sided dice, is given by the following:*

$$\pi(d, k) = \begin{cases} \frac{1}{6}, & d = 1, k \in \{0, 2, 3, 4, 5, 6\}; \\ 0, & 0 < k < 2d \text{ or } k > 6d; \\ \frac{5\pi(d-1, 0) + 1}{6}, & d > 1, k = 0; \\ \frac{1}{6} \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r), & \text{otherwise.} \end{cases} \quad (\text{A.2})$$

Proof. Let $\pi(d, k)$ denote the probability of rolling a score of k points with d fair, six-sided dice, where $0 < d \leq d^*$ and

$0 \leq k \leq 6d^*$. We determine $\pi(d, k)$ by considering the following four cases:

Case 1 ($d = 1$ and $k \in \{0, 2, 3, 4, 5, 6\}$). Since rolling the number 1 corresponds to a score of 0, our sample space of scores for one die is $S = \{0, 2, 3, 4, 5, 6\}$. Thus,

$$\pi(1, k) = \frac{1}{6}, \quad \forall k \text{ such that } k \in S. \quad (\text{A.3})$$

Case 2 ($0 < k < 2d$ or $k > 6d$). The least number that can possibly be rolled on any one die that will result in a nonzero score is two. Therefore, the smallest number that could possibly be rolled with d dice and results in a nonzero score would be $2d$, so that $\pi(d, k) = 0, \forall k$ such that $k < 2d$. Similarly, the greatest number that can possibly be rolled on any one die is six. Thus, the greatest number that could possibly be rolled with d dice would be $6d$ so that $\pi(d, k) = 0, \forall k$ such that $k > 6d$.

Case 3 ($d > 1$ and $k = 0$). Given a finite number of d identical dice, without loss of generality, we may number the dice $n_1, n_2, \dots, n_d$, so that when we refer to $d-1$ dice, we refer to the same $n_1, n_2, \dots, n_{d-1}$ dice. Since the events of rolling at least one 1 and not rolling any 1s are mutually exclusive, we may partition the set of all scores possible with d dice by the events which lead to a score of zero with $d-1$ dice and those which lead to a nonzero score with $d-1$ dice. This gives us the following:

$$\pi(d-1, 0) + \pi(d-1, \tau) = 1, \quad (\text{A.4})$$

where $\pi(d-1, \tau)$ is the probability of obtaining a nonzero score with $d-1$ dice and is equal to $\sum_{i=2}^{6(d-1)} \pi(d-1, i)$. Then,

$$\begin{aligned} \pi(d, 0) &= \pi(d-1, 0) \pi(d, 0 | d-1, 0) \\ &+ \pi(d-1, \tau) \pi(d, 0 | d-1, \tau). \end{aligned} \quad (\text{A.5})$$

Since $\pi(d, 0 | d-1, 0) = 1$ (the value of n_d is irrelevant if a one has already appeared on one of dice $n_i, i \in \{1, 2, \dots, d-1\}$), and $\pi(d, 0 | d-1, \tau) = \pi(1, 0)$ (the probability of rolling a 1 on dice n_d , or $1/6$), using (A.4) and substitution in (A.5), we obtain

$$\begin{aligned} \pi(d, 0) &= \pi(d-1, 0) + \frac{1}{6} (1 - \pi(d-1, 0)) \\ &= \frac{5\pi(d-1, 0) + 1}{6}. \end{aligned} \quad (\text{A.6})$$

Case 4 (all other k, d). We use similar logic to that of Case 3. The number appearing on n_d must belong to $\{1, 2, 3, 4, 5, 6\}$. Therefore, we may partition the set of all scores possible by the events that lead to a score of $k-2, k-3, \dots, k-r, r \in \{2, 3, 4, 5, 6\}$ and $k-r \geq 2$, on $d-1$ dice ($k-1$ would result in a score of zero). Thus,

$$\pi(d, k) = \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r) \pi(d, k | d-1, k-r). \quad (\text{A.7})$$

Since $\pi(d, k | d-1, k-r) = \pi(1, r) = 1/6 \forall r$, we have

$$\pi(d, k) = \frac{1}{6} \sum_{r=2}^{\min(6, k-2)} \pi(d-1, k-r). \quad (\text{A.8})$$

□

B. Solving Hog

With the solution to $\pi(d, k)$ complete, we turn to solving the following equation:

$$P_{i,j} = \max_{0 < d \leq d^*} \sum_{k=0}^{6d} \pi(d, k) (1 - P_{j,i+k}). \quad (\text{B.1})$$

With a goal threshold of 100 points, we know that $i, j \in \{0, 2, 3, 4, \dots, 98, 99\}$ are the only game states that we must calculate, since as mentioned in the text, $P_{i,j} = 1$ for $i \geq 100$ and $P_{i,j} = 0$ for $j \geq 100$. This yields a total of 9801 games states to be computed.

Note that at any game state, the probability of winning at that state is only dependent upon states, where the sum of the two players' scores is greater than or equal to the sum of the two players' scores in the current state [20]. This is due to the fact that players cannot lose points. In other words, the sum of the players' scores at any future state cannot be less than the current sum of scores. Therefore, the efficient way to proceed is to partition the games states into independent groups and work backwards. That is, we calculate the probability of winning for all states with $i+j = 198$ then $i+j = 197, i+j = 196, \dots, i+j = 2, i+j = 0$. Without loss of generality, we may replace 0 with 1 in each of the games states. This now corresponds to Figure 9, in which every ordered pair in each diagonal from left to right sums to the same number as the other ordered pairs on the diagonal. If we let i be the x -axis and j the y -axis, the ordered pairs (game states) of each diagonal belong to the lines $i+j = s, s \in \{2, 3, 4, \dots, 197, 198\}$.

We start with the right most diagonal (99, 99) and move to the left, solving each game state on the diagonal before continuing to the next diagonal partition. Let us call the player with i points Player 1 and the player with j points Player 2. If Player 1 receives zero points after which Player 2 also receives zero points in the proximate turn, we will have traversed from game state (i, j) to (j, i) and back to (i, j) . Therefore, every game state $(i, j), i \neq j$ is dependent on one other game state within its same partition, namely, (j, i) .

Then, let us rewrite (B.1) as follows:

$$\begin{aligned} P_{i,j} &= \max_{0 < d \leq d^*} \left\{ \left[\sum_{k=2}^{6d} \pi(d, k) (1 - P_{j,i+k}) \right] + \pi(d, 0) (1 - P_{j,i}) \right\}. \end{aligned} \quad (\text{B.2})$$

Recall that we are solving each partition in reverse order. So naturally, by the time we attempt to solve for $P_{i,j}$ and $P_{j,i}$, all game states $P_{j,i+k}$ and $P_{i,j+k}$ will have already been solved for $k \geq 2$. The section in square brackets, therefore, is a known value and will henceforth be denoted by $f(i, j, d)$. That is,

$$P_{i,j} = \max_{0 < d \leq d^*} \{ f(i, j, d) + \pi(d, 0) - \pi(d, 0) P_{j,i} \}. \quad (\text{B.3})$$

FIGURE 9: Visual partitioning of game states by the sum of players' scores.

Let us consider the case when $i = j$ and fix d at $\bar{d}$. Then, $P_{i,i} = f(i, i, \bar{d}) + \pi(\bar{d}, 0) - \pi(\bar{d}, 0)P_{i,i}$, or $P_{i,i} = [f(i, i, \bar{d}) + \pi(\bar{d}, 0)] / [1 + \pi(\bar{d}, 0)]$. If we again let d vary, $P_{i,i}$ may easily be solved by finding the $\bar{d}$ which will maximize the probability of winning at state (i, i) . This is given by (B.4) as follows:

$$P_{i,i} = \max_{0 < d \leq d^*} \frac{f(i, i, d) + \pi(d, 0)}{1 + \pi(d, 0)}. \quad (\text{B.4})$$

In this way, all that remains is to solve the following system of equations for each pair of game states (i, j) and (j, i) for $i \neq j$:

$$\begin{aligned} P_{i,j} &= \max_{0 < d \leq d^*} \{f(i, j, d) + \pi(d, 0) - \pi(d, 0)P_{j,i}\} \\ P_{j,i} &= \max_{0 < \delta \leq \delta^*} \{f(j, i, \delta) + \pi(\delta, 0) - \pi(\delta, 0)P_{i,j}\}. \end{aligned} \quad (\text{B.5})$$

As with the case when $i = j$, let us fix d and δ at $\bar{d}$ and $\bar{\delta}$. We solve for $P_{i,j}$ and $P_{j,i}$ to obtain

$$\begin{aligned} P_{i,j} &= \frac{f(i, j, \bar{d}) + \pi(\bar{d}, 0) - \pi(\bar{d}, 0)[f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0)]}{1 - \pi(\bar{d}, 0)\pi(\bar{\delta}, 0)}, \\ P_{j,i} &= \frac{f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0) - \pi(\bar{\delta}, 0)[f(i, j, \bar{d}) + \pi(\bar{d}, 0)]}{1 - \pi(\bar{\delta}, 0)\pi(\bar{d}, 0)}. \end{aligned} \quad (\text{B.6})$$

With $0 < d \leq d^*$ and $0 < \delta \leq \delta^*$, this will lead to $d^* \cdot \delta^*$ different $\bar{d} - \bar{\delta}$ combinations. To determine which of these leads to the optimal solution, we use a game theoretic approach. Let us call the player with i points Player 1 and the player with j points Player 2. Given Player 2's choice of δ dice, Player 1 will choose d such that he maximizes $P_{i,j}$, that is,

$$\begin{aligned} P_{i,j} &= \max_{0 < d \leq d^*} \frac{f(i, j, d) + \pi(d, 0) - \pi(d, 0)[f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0)]}{1 - \pi(d, 0)\pi(\bar{\delta}, 0)}. \end{aligned} \quad (\text{B.7})$$

As we vary $\bar{\delta}$ from 1 to δ^* , we will find all of Player 1's optimal choices given every possible choice from Player 2. Likewise, we do the same for Player 2, and given Player 1's possible decisions, we determine

$$\begin{aligned} P_{j,i} &= \max_{0 < \delta \leq \delta^*} \frac{f(j, i, \delta) + \pi(\delta, 0) - \pi(\delta, 0)[f(i, j, \bar{d}) + \pi(\bar{d}, 0)]}{1 - \pi(\delta, 0)\pi(\bar{d}, 0)}. \end{aligned} \quad (\text{B.8})$$

for each possible $\bar{d}$, $0 < \bar{d} \leq d^*$. This will lead us to find all combinations of $\bar{d} - \bar{\delta}$ for which neither player would choose to alter his choice of dice quantity, a Nash equilibrium. A Nash equilibrium is a situation in which given all other agent's strategies, no agent can improve his condition by changing his own strategy. Careful inspection of all game states and dice combinations leads to the conclusion that there is in fact one and only one Nash equilibrium for each set of (i, j) and (j, i) game states so that our problem is well defined.

Let us consider a specific example and turn to the states (90, 77) and (77, 90). Table 5 depicts Player 1's possible dice choices on the horizontal axis and Player 2's choices on the vertical axis. Each ordered pair is the calculation of $(P_{i,j}, P_{j,i})$ for the combination of dice to which it corresponds based on its placement in the table. Each player's optimal decision, given the choice of his opponent, is bolded and underlined. We find the Nash Equilibrium at $d = 3$ and $\delta = 7$.

Since we let $d_{\max} = 25$, for all games states in which $i, j \leq 50$, d^* and δ^* will be 25. Therefore, for approximately 1/4 of our game states, solving for $P_{i,j}$ and $P_{j,i}$ game theoretically will lead to 625 calculations of each probability. Traversing these solutions to find the unique Nash equilibrium is also costly. We hence provide an alternative numerical approach which converges to this solution.

For linear systems of equations, there are numerous iterative techniques for solving $\mathbf{Ax} = \mathbf{b}$. A common approach is Gauss-Seidel iteration, where we calculate "a sequence of approximate solution vectors $x^{(0)}, x^{(1)}, x^{(2)}, \dots$ " which will converge to the actual solution provided that our system satisfies certain conditions [21]. Since this is merely a cursory glance at iterative solutions, we refer interested readers to [21] for appropriate convergence criteria. We continue generating

TABLE 5: Finding the Nash equilibrium for $P_{90,77}$ and $P_{77,90}$.

$\delta \setminus d$	1	2	3	4	5
1	(0.6535, 0.2567)	(0.7572, 0.2394)	(0.8794, 0.219)	(0.8860, 0.2179)	(0.8678, 0.2210)
2	(0.6460, 0.3012)	(0.7478, 0.2701)	(0.8742, 0.2315)	(0.8799, 0.2297)	(0.8587, 0.2362)
3	(0.6398, 0.3386)	(0.7397, 0.2965)	(0.8699, 0.2417)	(0.8748, 0.2396)	(0.8506, 0.2498)
4	(0.6343, 0.3715)	(0.7323, 0.3207)	(0.8656, 0.2517)	(0.8696, 0.2497)	(0.8422, 0.2638)
5	(0.6222, 0.4439)	(0.7109, 0.3909)	(0.8394, 0.3141)	(0.8348, 0.3168)	(0.7970, 0.3394)
6	(0.6075, 0.5324)	(0.6830, 0.4822)	(0.8018, 0.4032)	(0.7839, 0.4151)	(0.7304, 0.4507)
7	(0.6040, 0.5535)	(0.6773, 0.5006)	(0.7971, 0.4143)	(0.7767, 0.4290)	(0.7181, 0.4712)
8	(0.6074, 0.5329)	(0.6859, 0.4727)	(0.8139, 0.3744)	(0.7991, 0.3858)	(0.7439, 0.4281)
9	(0.6118, 0.5068)	(0.6963, 0.4386)	(0.8335, 0.3280)	(0.8258, 0.3342)	(0.7762, 0.3742)
10	(0.6156, 0.4838)	(0.7055, 0.4084)	(0.8510, 0.2864)	(0.8502, 0.2871)	(0.8062, 0.3240)
11	(0.6188, 0.4645)	(0.7134, 0.3826)	(0.8662, 0.2504)	(0.8717, 0.2456)	(0.8331, 0.2790)
12	(0.6215, 0.4482)	(0.7201, 0.3606)	(0.8793, 0.2193)	(0.8904, 0.2095)	(0.8570, 0.2392)

these $x^{(i)}$ s until a predetermined level of precision is achieved. This means that we repeatedly solve

$$x_i^{(k)} = \left[- \sum_{j=1, j < i}^n \left(\frac{a_{i,j}}{a_{i,i}} \right) x_j^{(k)} - \sum_{j=1, j > i}^n \left(\frac{a_{i,j}}{a_{i,i}} \right) x_j^{(k-1)} + \left(\frac{b_i}{a_{i,i}} \right) \right] \quad (\text{B.9})$$

(assuming nonzero diagonal elements) until $\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| < \epsilon$, where ϵ is our error tolerance [21].

Returning to our problem, we have a nonlinear system $\mathbf{x} = \max_d \{\mathbf{A}_d \mathbf{x} + \mathbf{b}_d\}$ for which there is no general solution method [20]. Thus, such a system could not usually be solved iteratively using either of the two methods mentioned previously. Be that as it may, our particular system does in fact converge to the true solution (obtained game theoretically) if we use Gauss-Seidel iteration using an initial vector $\mathbf{x}^{(0)}$, where $x_i = 1.0$ for all i . We therefore recommend its use since it saves a considerable amount of computation time.

References

- [1] R. M. Hogarth, "Beyond discrete biases: functional and dysfunctional aspects of judgmental heuristics," *Psychological Bulletin*, vol. 90, no. 2, pp. 197–217, 1981.
- [2] D. N. Kleinmuntz, "Cognitive heuristics and feedback in a dynamic decision environment," *Management Science*, vol. 31, no. 6, pp. 680–702, 1985.
- [3] H. Jungermann, "Two camps on rationality," in *Decision Making under Uncertainty*, R. W. Scholz, Ed., pp. 63–86, North Holland, New York, NY, USA, 1983.
- [4] P. W. B. Atkins, R. E. Wood, and P. J. Rutgers, "The effects of feedback format on dynamic decision making," *Organizational Behavior and Human Decision Processes*, vol. 88, no. 2, pp. 587–604, 2002.
- [5] W. Edwards, "Dynamic decision theory and probabilistic information processing," *Human Factors*, vol. 4, no. 2, pp. 59–73, 1962.
- [6] B. E. Bakken, "On improving dynamic decision-making: implications from multiple-process cognitive theory," *Systems Research and Behavioral Science*, vol. 25, no. 4, pp. 493–501, 2008.
- [7] B. Brehmer, "Dynamic decision making: human control of complex systems," *Acta Psychologica*, vol. 81, no. 3, pp. 211–241, 1992.
- [8] C. Gonzalez, P. Vanyukov, and M. K. Martin, "The use of microworlds to study dynamic decision making," *Computers in Human Behavior*, vol. 21, no. 2, pp. 273–286, 2005.
- [9] J. D. Sterman, "Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment," *Management Science*, vol. 35, pp. 321–339, 1989.
- [10] E. Diehl and J. D. Sterman, "Effects of feedback complexity on dynamic decision making," *Organizational Behavior and Human Decision Processes*, vol. 62, no. 2, pp. 198–215, 1995.
- [11] R. H. Thaler, "From homo economicus to homo sapiens," *Journal of Economic Perspectives*, vol. 14, no. 1, pp. 133–141, 2000.
- [12] P. Slovic and S. Lichtenstein, "Comparison of Bayesian and regression approaches to the study of information processing in judgment," *Organizational Behavior and Human Performance*, vol. 6, no. 6, pp. 649–744, 1971.
- [13] A. Rapoport, "Research paradigms for studying dynamic decision behavior," in *Utility, Probability, and Human Decision Making*, D. Wendt and C. Vlek, Eds., pp. 349–375, Reidel, Dordrecht, Netherlands, 1975.
- [14] A. J. Mackinnon and A. J. Wearing, "Systems analysis and dynamic decision making," *Acta Psychologica*, vol. 58, no. 2, pp. 159–172, 1985.
- [15] J. D. Hey and J. A. Knoll, "Strategies in dynamic decision making—an experimental investigation of the rationality of decision behaviour," *Journal of Economic Psychology*, vol. 32, no. 3, pp. 399–409, 2011.
- [16] Mathematical Sciences Education Board, "The hog game," in *Measuring up: Prototypes for Mathematical Assessment*, pp. 141–155, National Academy Press, Washington, DC, USA, 1993.
- [17] T. W. Neller and C. G. M. Presser, "Pigtail: a pig addendum," *The UMAP Journal*, vol. 26, pp. 443–458, 2005.
- [18] J. C. Nunnally, *Psychometric Theory*, McGraw Hill Publishers, New York, NY, USA, 1978.
- [19] C. K. Doublas, "Multimedia technology," in *The Praeger Handbook of Learning and the Brain*, S. Feinstein, Ed., vol. 2, pp. 315–320, Praeger, Westport, Conn, USA, 2006.
- [20] T. W. Neller and C. G. M. Presser, "Optimal play of the dice game pig," *The UMAP Journal*, vol. 25, no. 1, pp. 25–47, 2004.
- [21] W. Cheney and D. Kincaid, *Numerical Mathematics and Computing*, Brooks/Cole Publishing, Belmont, Calif, USA, 1980.