

4-8-2011

## Dynamic decision making and race games

Shipra De

*University of Nevada, Las Vegas, des2@unlv.nevada.edu*

Follow this and additional works at: <https://digitalscholarship.unlv.edu/award>



Part of the [Behavioral Economics Commons](#), [Numerical Analysis and Computation Commons](#), and the [Theory and Algorithms Commons](#)

---

### Repository Citation

De, S. (2011). Dynamic decision making and race games.

Available at: <https://digitalscholarship.unlv.edu/award/5>

This Research Paper is protected by copyright and/or related rights. It has been brought to you by Digital Scholarship@UNLV with permission from the rights-holder(s). You are free to use this Research Paper in any way that is permitted by the copyright and related rights legislation that applies to your use. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/or on the work itself.

This Research Paper has been accepted for inclusion in Calvert Undergraduate Research Awards by an authorized administrator of Digital Scholarship@UNLV. For more information, please contact [digitalscholarship@unlv.edu](mailto:digitalscholarship@unlv.edu).

# Dynamic Decision Making and Race Games

by

Shipra De

B.A. Economics, B.S. Mathematics, and B.A. Computer Science

University Honors Scholar

Submitted to the Honors College

for the distinction of

Department Honors Scholar

University of Nevada Las Vegas

2011

University of Nevada Las Vegas  
Honors College

This thesis was presented

by

Shipra De

It was defended on

April 8th 2011

and approved by

Prof. Darryl Seale, Dept. of Management

Prof. Bradley Wimmer, Dept. of Economics

Russell Tinsley, Engineering Senior Specialist, JT3

Dean Peter Starkweather, Honors College

# Dynamic Decision Making and Race Games

Shipra De

Honors College

University of Nevada, Las Vegas

4505 Maryland Parkway

Las Vegas, NV 89154-7003

des2@unlv.nevada.edu

## Abstract

Frequent criticism in 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 the simple race game Hog, which has a computable optimal decision strategy. In the two-player game of Hog, 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, lack of player knowledge of key probabilistic concepts, or insufficient player motivation. Regardless, some players approached optimal heuristic strategies, which perform remarkably well. Results in Hog share similarities and differences with results in a predecessor dice game called Pig.

**Keywords:** behavioral economics; dynamic decision making; the game of Hog

## Acknowledgements

My Department Honors thesis has been the single most extensive and challenging task I have ever undertaken as a student. Spanning four semesters and requiring me to draw on knowledge acquired from three different disciplines, it has been the perfect culminating experience of my undergraduate career. The final product, the pages that follow, are not the fruits of my labor alone. I would not have been able to complete this research without the kindness and encouragement of many of my professors, friends, and family, who deserve recognition for their contributions.

I must first thank my thesis advisor, Dr. Darryl Seale, who has been an endless source of guidance, support, and optimism. Dr. Seale essentially allowed me to hijack the game of Hog from a set of experiments for which he had funding and intended to run himself. He let me make it my own and helped me every step of the way, from learning how to work with the IRB to the final write-up of the results. I could not have started or finished this project without him.

Special thanks to the remainder of my committee members who have been equally important and helpful. Specifically, Dr. Bradley Wimmer, who allowed me to audit his graduate research seminar where I discovered my interest for behavioral economics. He found me the perfect mentor in Dr. Seale and has continued his help by offering his knowledge and expertise. Also to Russell Tinsley, my former supervisor, whose mastery of statistics and the excel spreadsheet, as well as his willingness to impart these skills to me, made him an invaluable resource.

Thank you to the following professors who allowed me to consult them on various topics: Dr. Kim Barchard, who pointed me in the right direction for collecting qualitative data and introduced me to the IPIP, Dr. Monika Neda who shared her L<sup>A</sup>T<sub>E</sub>X expertise, Dr. Daniel Thompson who suggested the appropriate statistical tests to analyze my data, and Dr. Lawrence Larmore who availed himself for help with dynamic programming.

To the Honors College, thank you for this opportunity. I especially owe much gratitude and appreciation to Dr. Peter Starkweather, who has never failed to give me valuable criticism and advice despite the fact that I took shameful advantage of his open door policy on numerous occasions.

I could not have completed the necessary background research without the UNLV Lied Library and its access to thousands of books and journals (especially through the databases JSTOR and ScienceDirect). Having the ability to check out an unlimited quantity of books for a semester at a time has been one of the most useful and exciting privileges I have ever had.

Furthermore, an experiment of this nature cannot be done without subjects or an appropriately equipped laboratory. I am particularly grateful to the Office of Information Technology for providing the use of one its large computer labs and to Dr. Daniel McAllister and Dr. Rebecca Guidice who allowed me to recruit participants from their MGT 301 classes. Thank you to all of their students who participated. The actual sessions themselves would not have run as smoothly without the help of Management graduate students Nazia Charania, Chris Powell, and Nishant Sinha, and my friends Nicholas Glorioso and Ryan Shugars. Thank you all.

Last but not least, thank you to my sister, Tondra De, who spent numerous hours helping me use SPSS, and thank you to my parents for their endless love and support. If not for the faith and confidence of my friends and family I might not have survived the last two semesters.

# Contents

<b>1</b>	<b>Background</b>	<b>1</b>
1.1	Behavioral Economics . . . . .	3
1.2	Prospect Theory . . . . .	5
<b>2</b>	<b>Introduction</b>	<b>7</b>
2.1	Literature Review . . . . .	9
<b>3</b>	<b>The Game of Hog</b>	<b>14</b>
3.1	Origins and Strategy . . . . .	16
3.2	Optimal Policy . . . . .	18
<b>4</b>	<b>The Experiment</b>	<b>22</b>
4.1	Procedures . . . . .	24
<b>5</b>	<b>Results</b>	<b>25</b>
5.1	Performance . . . . .	27
5.2	Improvement . . . . .	32
5.3	Survey Data . . . . .	34
5.4	In Comparison to Pig . . . . .	38
5.5	Heuristic Solutions . . . . .	40
<b>6</b>	<b>Discussion</b>	<b>44</b>
	<b>References</b>	<b>49</b>
<b>A</b>	<b>Solving the Game Hog</b>	<b>52</b>
A.1	A Select Review of Probability Theory . . . . .	52
A.2	Obtaining a Score $k$ , with $d$ Dice . . . . .	53
A.3	A Brief Introduction to Dynamic Programming . . . . .	55
A.4	Solving $P_{i,j}, 0 \leq i, j < 100$ . . . . .	56

A.5 Solution Algorithms . . . . .	61
<b>B Hog Solution Code, C++</b>	<b>66</b>
<b>C Subject Recruitment Flier</b>	<b>70</b>
<b>D Instructions</b>	<b>71</b>
<b>E Survey</b>	<b>73</b>

## List of Tables

3.1	Hog Sample Game, First Six Turns . . . . .	15
3.2	Maximizing Expected Points in Hog . . . . .	17
4.1	Frequency of Numbers Appearing on Electronic Dice . . . . .	22
5.1	Aggregate Survey Data: Demographic Information . . . . .	35
5.2	Aggregate Survey Data: Knowledge Base Questions . . . . .	36
5.3	Aggregate Survey Data: Individual Characteristics . . . . .	37
5.4	Regression Results . . . . .	39
A.1	Finding the Nash Equilibrium for $P_{90,77}$ and $P_{77,90}$ . . . . .	59



## List of Figures

1.1	The Evolution of Homo economicus [34]	4
3.1	Expected Scores for Rolling $d$ Fair, Six-Sided Dice, $1 \leq d \leq 25$	17
3.2	Optimal Solution for the Two-Player Game of Hog and 25 Dice Maximum	19
3.3	Optimal Strategies with Opponent's Score Fixed at 0, 50, 80, and 99 Points	20
3.4	Optimal Solution for the Two-Player Game of Hog and 25 Dice Maximum, 2D View	21
4.1	Screenshot of Hog Dice Game Software at Initial Turn	23
4.2	Screenshot of Hog Dice Game Software Following Successful Roll	24
5.1	Treatment 1: Percentage of Rolls Characterized by Delta	27
5.2	Treatment 2: Percentage of Rolls Characterized by Delta	28
5.3	Treatment 1: Average Observed and Optimal Numbers of Dice By Score Discrepancy	30
5.4	Treatment 2: Average Observed and Optimal Numbers of Dice By Score Discrepancy	30
5.5	Per Player Frequency of the Average Magnitude of Differences Between Optimal and Observed Numbers of Dice	31
5.6	Treatment 1: Subject Performance Across Games	33
5.7	Treatment 2: Subject Performance Across Games	34
5.8	Cumulative Distributions, Hog vs Pig	40
5.9	Trial and Adjustment Heuristic Performance in 1M Games	42
5.10	Consistent Roll Heuristic Performance in 1M Games	43
A.1	Visual Partitioning of Game States by the Sum of Players' Scores	57

# 1 Background

Any standard text will define economics as some variation of “the study of how society manages its scarce resources” [24]. Although sociologists and political scientists too may be interested in such phenomena, what sets apart the economist from his fellow social scientists is the systematic framework under which he does his analysis [38]. He is not interested in ad hoc claims which correspond to a particular set of data, but rather general and refutable theories with explanatory power with respect to varying circumstances.

What exactly is a “theory” in an empirical science such as economics? According to Silberberg, a theory is composed of three separate elements. The first of these is the set of assertions or postulates “concerning the behavior of various theoretical constructs, i.e., idealized (perhaps mathematical) concepts, which are ultimately related to real-world objects” [38]. An example of such a postulate is the statement that firms maximize profits or that demand is downward sloping [38]. We denote this set as  $A$  and must bear in mind that since the statements in  $A$  cannot be observed, it is fruitless to question their exact conformance with reality [38]. As Milton Friedman mentioned in his famous text on positive economics, it is a theory’s predictive power by which it should be judged [14]. Since there is often an inverse relationship between realism and manageability, economists may very well give up some of the former to acquire the latter, provided that doing so does not compromise the predictive ability of the theory. On the other hand, the second component of a theory must be observable and realistic; it is a set denoted  $C$  as it consists of the assumptions or test conditions with which we will test our given collection of postulates [38]. This leaves the final part of a theory, the set of events,  $E$ , which the theory attempts to predict [38].

Returning to the example of the nature of demand, if we desired to test whether a rise in the price of a particular good, say  $x$ , reduces the quantity demanded of good  $x$ , our theory would stand as follows: given that demand is downward sloping ( $A$ ), if there exist data indicating that the price of good  $x$  has increased while real income and other prices have remained constant ( $C$ ), then it will be observable from the same data that the quantity demanded of good  $x$  has decreased ( $E$ ). Using propositional logic we symbolize this construct as  $(A \rightarrow (C \rightarrow E))$ , or the logical equivalent,  $((A \& C) \rightarrow E)$ , where “ $\rightarrow$ ” is read as “implies” [38]. Since it logically follows that

$(\neg E \rightarrow \neg(A \& C))$ <sup>1</sup>, we can test a theory and either refute or confirm its postulates by the truth or falsity of the event  $E$ . Explicitly, since  $\neg(A \& C)$  means  $\neg A$  or  $\neg C$ , if our event is false and we can verify that our test conditions are true, we may conclude that one or more of the assertions must be false<sup>2</sup>. Note that we can only confirm a theory, not prove it, since the converse of an implication is not necessarily true. However, the more frequently a theory is confirmed by empirical data, the more confidence we may have in its postulates [38].

One postulate of economics that has recently met with resistance is the assertion that individuals and firms behave rationally. What is meant by the the word “rational” has itself been mired in controversy, but most economists have reached the consensus that “rational behavior simply implies consistent maximization of a well-ordered function, such as a utility or profit function” [4]. Critics of this theory and empirical evidence suggest that economic agents do not behave rationally; a classic refutation of this postulate has been found in the ultimatum game, an experiment which has been repeated with various stakes and conditions.

The ultimatum game is played by two individuals, one of whom is randomly designated as the proposer while the other is known as the responder. The proposer is asked to determine the division of a specified sum of money between himself and his opponent. If the responder does not agree to the terms of the proposer, both players walk away without any earnings. The theory of a rational individual would suggest that the responder should agree to receive any nonzero amount since even a penny would leave him better off than accepting nothing [18]. Yet time and again, experimental research has found that responders often turn down offers unless they are at least twenty percent of the original sum [18]. These results have held whether the stakes are \$10, \$100, or even over a week’s pay for participants in a poor nation [18]. Responders are clearly willing to penalize their opponents for inequitable distributions even though that amounts to penalizing themselves in the process, thereby refuting the theory of rational behavior.

The ultimatum game is not the only experiment that has highlighted inconsistencies between economic theory and empirical evidence. Numerous other anomalies such as framing<sup>3</sup> and endow-

---

<sup>1</sup> “ $\neg$ ”, meaning “not”

<sup>2</sup> On the other hand if the assumptions are also false, the postulates may be either true or false.

<sup>3</sup> Situation “where extensionally equivalent descriptions lead to different choices by altering the relative salience of different aspects of the problem” [20].

ment<sup>4</sup> effects have also questioned the idea of a rationally behaving agent. These deviations from what conventionally has been considered the norm has led to the development of a special subfield of economics known as behavioral economics.

## 1.1 Behavioral Economics

Behavioral economics is a relatively new branch of economics, the relevance of which has only recently been widely-accepted. To define it explicitly, it is the branch of economics that attempts to develop models that more accurately depict real world phenomena through the integration of psychological principles.<sup>5</sup> Specific areas of study include heuristics and cognitive errors (prospect theory, money illusion, etc.), framing effects, and other anomalies (endowment effects, inequity aversion, etc.).

Mulainathan describes how traditional economics “conceptualizes a world populated by calculating, unemotional maximizers that have been dubbed Homo [e]conomicus” [26]. This rational agent model presumes that individuals have a strong grasp of their own preferences and will work towards maximizing them. Rabin explains that given his utility function, a rational individual will try to maximize his expected utility, or happiness [35]. This model has been defended with the argument that market forces such as competition and arbitrage should ensure an environment in which only rational agents can survive. However, as Mulainathan, Rabin, and many others have demonstrated, this is not necessarily true. Many individuals lack the ability to identify and articulate their own preferences, and some of the basic assumptions of utility theory are violated by cognitive biases such as framing effects [35]. In such cases the market is incapable of rooting out what Thaler terms the “quasi-rational” individual, or he who “[tries] hard but [is] subject to systematic error” [40]. Thus, as Figure 1.1 humorously depicts, behavioral economics requires that we scale back our faith in “rational man’s” propensity for maximizing his utility.


There are several ways in which individuals may stray from traditional economic theory. Be-

---

<sup>4</sup>Violation of consumer theory in which “the selling price for consumption goods is much higher than the buying price” [20].

<sup>5</sup>Behavioral economics, or the application of psychology to enhance economic models, should not be confused with experimental economics which uses laboratory experimentation to study that which is difficult to examine in naturally occurring economic phenomena.

Figure 1.1: The Evolution of Homo economicus [34]



Behavioral economists have identified three specific “bounds” on human behavior encapsulating how “real people” differ from [H]omo economicus” [18]. First there is bounded willpower, which refers to the concept that people do not make choices that will not maximize their utility in the long run [26]. Often there is inherent conflict between short run and long run goals. For example, while a particular level of saving will maximize consumption over a lifetime, individuals may fall well short of this optimum in order to increase immediate gratification. Second, bounded self-interest exists in that some people will risk their own happiness to benefit or spite others. We have already become acquainted with this bound through the ultimatum game, which highlights how people want to be treated fairly. Finally, we have bounded rationality, which is the underlying topic of this study. It is the idea that people are limited in their cognitive abilities, which hinders their capacity to solve problems [20, 26]. That is not to say that individuals cannot respond in a rational manner given modest computational skills and less than perfect memories [18]. For instance, a person who uses a heuristic, a cognitive simplification mechanism, may certainly minimize the opportunity cost of thinking [18]. Nevertheless, such “mental shortcuts and rules of thumb” lead to (often predictable) outcomes which differ from standard models [18]. These divergences fall into the categories of judgment and decision making.

It is important to recognize that there are certainly cases where “market forces are strong enough to make the three ‘bounds’ irrelevant for predictive purposes” suggesting that despite irrationality, “markets can sometimes lead to behavior consistent with conventional economic assumptions” [18]. For example, Becker has shown that even if we assume that individuals and firms behave completely erratically, as if their decisions are determined by the roll of a die, the

demand curve would still be negatively inclined [4]. However, the theory that the demand curve is downward sloping is hardly under dispute. Alternatively, behavioral economics is appreciated for its importance in law and other applications where one is involved in prescriptive analysis, or that which relates to the development of policies directed to elicit particular behaviors [18]. To illustrate, we consider the issue of retirement savings. The standard models for lifecycle consumption depict rational, far-sighted individuals who in retirement will consume their savings from youth and middle-age [22]. To the contrary, empirical evidence reveals that a great number of individuals do not in fact save much when they are younger. Studies have uncovered, though, that there is a type of inertia when it comes to 401(k) participation. That is, participation is appreciably greater under automatic enrollment and largely dependent upon default investment allocations [22]. Inexplicable by traditional methods, behavioralists can resolve such activity using psychological models.

Perhaps the greatest criticism of behavioral economics is that there is no clear definition of what it is. Posner stresses how it seems to be negatively defined, i.e. “economics minus the assumption that people are rational maximizers of their satisfactions” [33]. There is no clear forecast of what “behavioral man” will do in any given situation [33]. Thus, while behavioral economists have pointed out the flaws of the rational agent model, they have yet to present a fully developed alternative. Yet, human actions are not so random as the roll of a die. Experiments have shown that departures tend to be systematic [18]. In the ultimatum game we can understand that individuals do not always maximize in accordance to absolute levels, but in relative ones: the greater the inequity in the division of the sum, the more relative wealth the responder gains by rejecting the offer. Therefore with more experimentation and understanding of human psychology, this fledgling field may soon develop a more comprehensive behavioral model. Such a model will surely be harder to manage than the traditional one, but given that the rational agent model cannot always predict behavior that is observed, it is important that we attempt to establish one that can.

## 1.2 Prospect Theory

One of the most sophisticated theories rivaling that of rational choice economics, at least in the realm of decisions subject to risk, is that of prospect theory. First developed in the late seventies by Tversky and Kahneman, it rejects expected utility theory in favor of a model that attempts to

reconcile behavior contradictory to fundamental aspects of utility theory. Tversky and Kahneman highlight two results that are inconsistent with the predictions of expected utility theory.

The first of these they call the certainty effect. It is the tendency that people have to “overweight outcomes that are considered certain, relative to outcomes which are merely probable” [41]. For instance, Tversky and Kahneman find that in an experiment in which subjects are faced with the decision of *A*: \$2,500 with probability .33, \$2,400 with probability .66, and \$0 with probability .01, or *B*: \$2,400 with certainty, most respondents choose *B*, indicating that  $u(2,400) > .33u(2,500) + .66u(2,400)$ , or simplified,  $.34u(2,400) > .33u(2,500)$  [41]. On the other hand, when given the choice between *C*: \$2,500 with probability .33 and \$0 with probability .67, or *D*: \$2,400 with probability .34 and \$0 with probability .66, the majority of respondents choose *C* [41]. This yields the exact opposite inequality and hence defies one of the basic rules of expected utility theory. Interestingly, when the signs are reversed on all of the payoffs, the exact opposite choices are preferred, a phenomenon that Tversky and Kahneman term the reflection effect [41]. It is thus apparent that “certainty increases the aversiveness of losses as well as the desirability of gains” [41].

The second violation is the isolation effect, i.e. the notion that, for simplification’s sake, “people often disregard components [that] alternatives share, and focus on the components that distinguish them” [41]. The authors explain that often in two stage games, where the first involves a probability  $p$  of continuing to stage two and  $(1-p)$  of not progressing, people will only look at the possible outcomes in stage two. Ignoring shared aspects can lead to the same type of contradiction seen with the certainty effect. Therefore, this type of fickleness arising from a strong dependency between events has serious implications since “it violates the basic supposition of a decision-theoretical analysis, that choices between prospects are determined solely by the probabilities of final states” [41].

To contend with these discrepancies, Tversky and Kahneman devised prospect theory, which outlines two stages of decision making, namely editing and evaluation. Editing is further broken into four parts: (1) coding, in which decision makers organize their choices as gains or losses relative to a particular frame of reference; (2) combination, the simple combining of probabilities referring

to the same outcome; (3) segregation, where a riskless component may be isolated from the other choices for the sake of simplification; and (4) cancelation, a manifestation of the isolation effect [41]. During the evaluation process decision makers will assign a decision weight to each of the probabilities from the editing stage and a subjective value to each of the possible outcomes—in essence, the deviation from his or her point of reference [41]. With this information at hand, decision makers will make the choice leading to the highest attainable value. In this way utility is measured not by the final outcomes themselves but by the relative gains and losses associated with them. Accordingly, this formulation expects the anomalies mentioned before and leads to an S-shaped value function distinguished by three important characteristics: “(1) it is concave in the domain of gains, favoring risk aversion; (2) it is convex in the domain of losses, favoring risk seeking; (3) most important, the function is sharply kinked at the reference point, and loss-averse—steeper for losses than for gains by a factor of about 2-2.5” [20].

Prospect theory emphasizes the reference-dependent nature of perception [20]. It also gives rise to the importance of framing effects. If value is measured by change, and if a loss hurts by a quantity more than twice the amount of happiness felt from a gain of the same magnitude, then how questions are framed—whether they are presented as gains or losses—will have a significant impact on decision making. In fact, research has shown that sets of identical choices in which one set portrays an option as a gain while the other describes the same option as a loss leads to opposing decisions [20]. Unfortunately, while prospect theory “tells us that choices depend on the framing of a problem, [it] does not tell us how people will spontaneously create their own frames” [40]. As a consequence, even though prospect theory has made excellent progress in developing a model that more accurately predicts real-world behavior, the theory is still not complete.

## 2 Introduction

Within behavioral economics, one area of special importance is decision theory. The state of the economy is dependent upon the distinct decisions of millions of individuals, firms, and industries. In making these choices, economic entities are often required to process vast and/or complex amounts of data. However, due to limited cognitive powers, people often adopt heuristics or other simplifi-



cation techniques, the consequence of which may lead to deviations from optimal behavior. Various psychological studies have indicated that there are in fact “numerous cognitive and other bounds on human rationality, often producing systematic errors and biases” [39]. Therefore, experiments in decision making are necessary for identifying and predicting how economic agents make choices in real-world situations.

To date, many studies have focused on identifying detrimental effects of using judgment heuristics. Decision makers are known for ignoring base rate information, failing to revise opinions, having unwarranted confidence, and harboring hindsight biases<sup>6</sup>, to name a few [16]. This seems to imply that humans are fairly incompetent beings [21]. 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” [16]. So why is there such a disconnect between experimentation and real-world phenomena?

One of the major complaints against decision-making experiments is that they largely concentrate on discrete instances. In doing so they ignore the continuous and adaptive nature of decision making [16]. In actuality, individuals are constantly receiving feedback not only in the form of outcomes but also in regard to shifts in the environment and the conditions of the decision which are affected by his or her past actions [39]. As a result, these types of experiments are not necessarily an accurate representation of decision making outside of an artificial realm [21]. Kleinmuntz emphasizes that “costs and benefits of cognitive heuristics ought to be evaluated in tasks having a dynamic, continuous character” [21]. The claim, as 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” [36, 19].

To correct for this lack of continuous processing, some recent studies have begun to focus on dynamic decision making. In order for a task to be considered dynamic it must be characterized by multiple decisions [2]. These decisions must be interdependent, and the environment must change as a result of the decision-maker’s own earlier actions and additional external forces [2]. As an example, consider an individual who invests his money in the stock market. He will choose various

---

<sup>6</sup> “[T]he tendency of decisionmakers to attach an excessively high probability to an event simply because it ended up occurring” [18].

stocks to achieve his goal of maximizing profits, and may opt to buy, sell, or hold those shares at different points in time. His choices as well as those of millions of other individuals will affect stock prices which will in turn affect his and others' future decisions. Such a complex system dependent upon so many individuals is hardly ideal for studying individual decision making and the cognitive reasoning behind it. As such, we instead establish suitable laboratory experiments typically involving computer simulations called microworlds<sup>7</sup> that are manageable but still represent real-world situations.

## 2.1 Literature Review

Initial interest in dynamic decision making was generated during the early 1960s through the independent efforts of Toda and Edwards, the latter of whom is considered the father of behavioral decision theory, a subject concerned with how people *actually* make decisions as opposed to how they *should* make them. In his seminal 1962 paper, Edwards introduces dynamic decision theory by first reviewing static decision theory, which presents a decision maker faced with a well-defined set of choices [12]. Edwards describes the subjective expected utility model in which each of these choices, as well as the possible states of the world (which may or may not be associated with specific probabilities), is coupled with a value [12]. Together, these values form a payoff matrix, and in the static case, the decision maker takes one of the available courses of action and “receives the value or payoff associated with the intersection of that course of action and the state of the world which actually obtained—and then the world ends” [12]. In contrast, the dynamic world continues as the decision maker makes a series of choices, each dependent on the last, and attempts to maximize payoffs in the long run. With each decision, the decision maker gains new information and must now deal with the possibility of a changing environment, since the states may change due to his past decisions, autonomously, or both [12]. It is this collection of information and a changing environment which makes dynamic decision theory “so difficult and so much fun” [12].

Due to its complicated nature, dynamic decision making has not received the same attention as static decision making [3]. Brehmer summarizes how research that has been completed in the

---

<sup>7</sup>Microworlds sometimes are also known as synthetic task environments, high fidelity simulations, interactive learning environments, virtual environments, or scaled worlds [15].

field has been tackled using two distinct approaches. The first, the individual differences technique, attempts to either predict behavior or “identify the demands” of the tasks performed [7]. To this end, the approach first involves separating subjects into those who succeed and those who do not and then calls for comparing these groups in terms of behavior and psychological test scores that could potentially account for the disparity in performance [7]. Unfortunately, results have failed to produce any significant correlation between the two [7]. Efforts to train subjects to adequately handle complex systems 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 by which we can characterize different dynamic decision experiments and allow for more comparability between them. Then, by altering one trait at a time, we might systematically determine which has the greatest effect on performance, a process which carries with it the possibility of developing a general theory.

Gonzales, et al., have recently revived the taxonomical approach to reviewing dynamic decision-making problems. They feature four characteristics by which we may classify them:

(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<sup>8</sup>, a criterion added by Brehmer [4, 12, 15].

(ii) *Complexity*

Complexity refers to the situation in which decision makers are obliged to keep track of many factors and possibly conflicting goals [7]. It is difficult to gauge, as it is relative to a specific decision maker with particular cognitive abilities [15]. Regardless, it is characterized by three attributes that work in conjunction: “(1) the number of components in the system, (2) the

---

<sup>8</sup> “[D]ecision 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].

number of relationships among the components (i.e., the degree of coupling), and (3) the types of relationships among the components” [15].

(iii) *Opaqueness*

Opaqueness measures how visible different aspects of the task are to the decision maker [15]. A microworld 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 [15].

(iv) *Dynamic Complexity*

Dynamic complexity focuses on the nature of feedback provided by the system [15]. 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 [15]. 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 dynamic decision making experiment is the Beer Distribution Game [39]. In this role-playing experiment, Sterman arranged several teams of four players including a producer, distributor, wholesaler, and retailer. Each participant attempted to “manage a simulated inventory distribution system” [39]. 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 [15]. On the other hand, because consumers’ demand remained unknown to most players, opaqueness was high, as was dynamic complexity since feedback delays were frequent and large [15]. Performance in the Beer Distribution Game was pretty poor; on average, team costs were ten times greater than the optimal benchmark cost [39]. The study concluded that this was due to misperception of feedback [39]. 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 [39].

This and other experimental results have led feedback to be a popular topic for decision making research. Diehl and Sterman in a similar study also had subjects oversee an inventory while having to contend with fluctuating sales [10]. However, they varied feedback strength and delays to monitor its bearing in the decision making process. If misperception were truly the root cause of poor performance, adjusting these variables should have had little effect on deviations from optimality [10]. Yet results indicated considerably suboptimal performance all around, with “performance [deteriorating] dramatically with [increased] time delays and feedback effects” [10]. Feedback, therefore, appears to have a appreciable impact on performance. To explore this further, Atkins, Wood and Rutgers have experimented with feedback methods, i.e. tabular vs. graphical data, finding that graphical feedback leads to better performance, although tabular data suggests greater learning [2].

Researchers in these studies have admitted that an alternative explanation for poor performance could just be that subjects did not understand the full complexity of the game or had insufficient experience despite multiple trials [10]. Given that other scientific investigations have led to the conclusion that “decision-making expertise typically requires more than 10 years experience in environments where tasks are exacting and feedback unambiguous,” it may very well be unrealistic to expect subjects to achieve optimal performance in six hours time even with simplified and repeated scenarios [3]. Decision makers only have two methods of processing to make decisions, namely an analytic process or an intuitive process [3]. So given that ambiguous and infrequent feedback further complicates matters, it is hardly surprising that studies such as that by Diehl and Sterman have found that even if participants begin with analytical reasoning, they almost always move to the intuitive processing method, forsaking calculations in favor of heuristics [3, 10].

Thus, the performance of heuristics, too, is an important topic in decision making. Kleinmuntz tested the performance of various heuristic strategies in a simulation of medical decision making [21]. He concluded that the success of a heuristic strategy is dependent upon the dynamic characteristics of the task itself, such as feedback quality [21]. This is reasonable since optimal performance depends on “the availability of feedback” and “opportunities for taking corrective actions based

upon that feedback” [21]. Kleinmuntz showed that if these criteria were met, it is possible to enjoy successful performance by using less complicated decision strategies [21]. All of this 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” [40].

Criticisms directed at recent dynamic decision making experiments include that they are “overly complex, with (often) ambiguous feedback, lacking clearly delineated subject goals, and not amenable to analytical solution” [36]. For this reason, Seale, Rapoport and Stein study dynamic decision making “in a new paradigm—the jeopardy race game” [37]. By “jeopardy” we mean that at every turn a player puts his or her point total at risk, and by “race” we refer to the fact that the players compete to become the first to achieve a designated threshold of points [29]. Such a game has low dynamics, complexity, opaqueness, and dynamic complexity as states change only as a consequence of the player decisions, variables are few, and feedback clear and immediate. A highly noteworthy difference between studies in jeopardy race games and previous research in dynamic decision making 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 [37].

Seale, Rapoport and Stein studied the game of Pig. In Pig, each player strives to be the first to achieve 100 points. At each turn the player makes a series of decisions to roll a die (again) or hold. If the player rolls a one at any stage in the roll sequence, he immediately loses his point total for that turn and play shifts to his opponent. If he rolls anything other than a one, he adds that number to his point total for the turn. When he holds, his turn total is added to his overall score and play is transferred to his opponent [27]. For instance, suppose it is the start of a game and it is Player 1’s turn. He rolls a five and decides to roll again. On the second roll he receives a one and hence no points for his turn. It is now Player 2’s opportunity to score. He rolls a six, followed by a three, and then decides to hold. Player 2’s score is now nine points and play resumes with Player 1. They will continue in this way until one of them reaches 100 points. The two-player game of Pig was solved by Neller and Presser using a computationally intensive method [27]. This is highly

important because dynamic decision theory focuses on evaluating human decision making against a “series of temporally related decisions with optimal solutions yielded by mathematical models” [16]. Its simplicity, in addition to its ability to satisfy the requirements of a dynamic task, makes Pig a rather attractive dynamic decision making experiment.

Despite the simplicity of the game, Seale, Rapoport and Stein discovered that participants still did not roll often enough and failed to correct their actions (over the course of multiple games) [37]. They found that the vast majority of players were highly insensitive to the number of points by which they were ahead of or behind their opponents, and the only indication of learning stemmed from the discovery that there was a significant positive correlation between players’ and their opponents’ mean stopping threshold (average turn total at which a player decides to bank his points), suggesting that players tended to copy their opponents’ strategies [37]. Overall, their results did not support the view of an adaptive decision maker who approaches optimal behavior with repeated play [37].

### 3 The Game of Hog

To further understand dynamic decision making in the realm of jeopardy race games, we study one of the many variants of Pig called “Fast Pig” [30]. More commonly known as Hog, it is a race game in which each player has only one roll per turn. However, he or she may choose to simultaneously roll as many dice as he or she pleases<sup>9</sup>. 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 some designated point threshold wins the game.

The beginnings of a sample game may be found in Table 3.1. Player 1 begins 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. Play returns to Player 1 who chooses fewer dice; he again obtains a positive turn total which is added to

---

<sup>9</sup>For practical purposes we typically impose a maximum number of dice,  $d_{max}$ .

Table 3.1: 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

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.

Technically, Hog is not a jeopardy dice game since the turn total is never actually in jeopardy. However, we can understand it as the conceptual equal of playing Pig, since it is essentially the same game with the slight twist of “commit[ting] to the number of rolls in a turn before the turn begins” [28]. An additional important benefit of Hog over its porcine ancestor is that it also allows us to measure aggressive or conservative behavior [36]. Although the study by Seale, Rapoport, and Stein observed that players were five times more likely to choose a conservative hold (stopping before achieving the optimal point threshold) over an an aggressive roll (rolling past the optimal point threshold), the conclusion that players are more likely to play conservatively might be premature [37]. Their study could only analyze a subset of the data, i.e., only those decisions in which a player decides to hold. In cases where the subject was thwarted from continuing play due to the appearance of a one, it is not certain what the subject would have chosen to do had he approached the optimal level of points during that turn. Therefore, a player may not have had the opportunity to actually display aggressive tendencies because probability got in the way. In contrast, since players in Hog must state at the outset of each roll the number of dice they will play, this number 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).



### 3.1 Origins and Strategy

The roots of the dice game Hog may be traced back to a 1993 publication by the Mathematical Sciences Education Board. With the addition of an investigative worksheet into possible outcomes and strategies for ideal play, it served as a prototype of a fun and challenging way to educate grade school children about both basic numerical ideas and more advanced probabilistic concepts [25]. Since then it has been used as an instructional tool for teaching not only elementary and secondary school students but even undergraduate statistics majors, graduate mathematics majors, and Masters of Business Administration (MBA) students [13]. The nature of the game allows players to observe patterns and subsequently form and test hypotheses [5].

The same qualities that make Hog wonderful for classroom instruction also make it ideal as a dynamic decision-making task. Specifically, 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” [25]. Interested and curious to determine the best strategy to win, students and subjects alike are motivated to think critically with the ultimate desire of performing optimally.


Literature related to Hog tends to focus on optimal performance in regard to maximizing the expected value of points. These average values can and have been approximated through testing in the classroom [13]. More accurate numbers can be determined through computer simulation, while the most exact answers are easily calculated mathematically, as seen from Table 3.2. Since each die is independent of the others, to determine the probability of a nonzero score with  $d$  dice we calculate  $(5/6)^d$ . The average nonzero score is  $4d$  so that the product of columns 2 and 3 yield the expected value of rolling  $d$  dice. We graph column 4 in Figure 3.1. 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. Therefore, even though both five and six dice provide the largest expected value, the choice of five dice still commands the greater chance of a positive score while the option of six dice covers a larger range of scores.

Table 3.2: Maximizing Expected Points in Hog

d	$(5/6)^d$	$4d$	$(5/6)^d * 4d$	d	$(5/6)^d$	$4d$	$(5/6)^d * 4d$
1	0.833333	4	3.333333	14	0.077887	56	4.361648
2	0.694444	8	5.555556	15	0.064905	60	3.894328
3	0.578704	12	6.944444	16	0.054088	64	3.461625
4	0.482253	16	7.716049	17	0.045073	68	3.064981
<b>5</b>	<b>0.401878</b>	<b>20</b>	<b>8.037551</b>	18	0.037561	72	2.704395
<b>6</b>	<b>0.334898</b>	<b>24</b>	<b>8.037551</b>	19	0.031301	76	2.378866
7	0.279082	28	7.814286	20	0.026084	80	2.086724
8	0.232568	32	7.442177	21	0.021737	84	1.825884
9	0.193807	36	6.977041	22	0.018114	88	1.594025
10	0.161506	40	6.460223	23	0.015095	92	1.388734
11	0.134588	44	5.921871	24	0.012579	96	1.207595
12	0.112157	48	5.383519	25	0.010483	100	1.048260
13	0.093464	52	4.860122				

Note: The expected value of a nonzero score per die is 4.

Figure 3.1: Expected Scores for Rolling  $d$  Fair, Six-Sided Dice,  $1 \leq d \leq 25$



### 3.2 Optimal Policy

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, we 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, as well as how close either is to the goal threshold. Feldman and Morgan recognize that these difficulties, coupled with the complexities mentioned earlier, make Hog an “interesting activity for students with respect to decision making in the face of uncertainty” [13].

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 of Hog. 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 [28]. 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.<sup>10</sup> For each  $d$ , this is determined by the summation of the probability of rolling each possible score times the probability that Player 2 will not win by rolling optimally in his following turn, i.e.

$$P_{i,j} = \max_{0 < d \leq d^*} \sum_{k=0}^{6d} \pi(d, k)(1 - P_{j,i+k}). \quad (3.1)$$

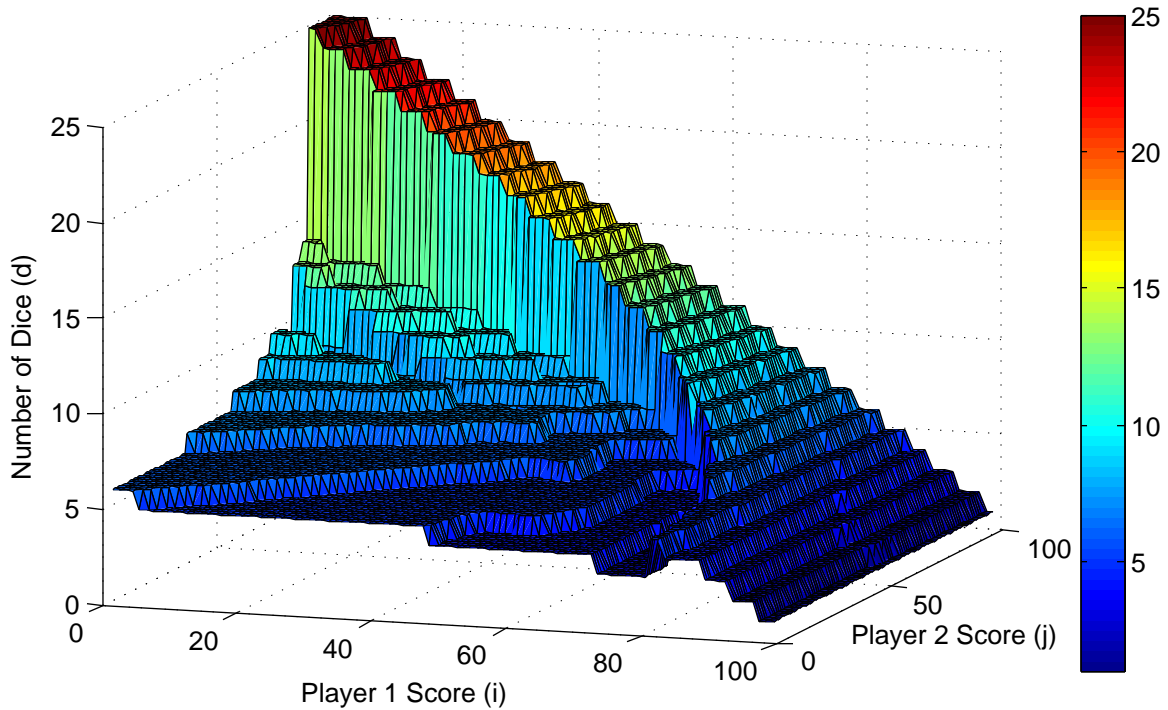
Interested readers may turn to Appendix A for an in-depth analysis of this solution, detailing how we computationally determine  $\pi(d, k)$  and solve Equation 3.1 for each game state. The optimal roll decisions for the two-player game of Hog is presented in graphical form in Figure 3.2. Player 1’s

---

<sup>10</sup>Here  $d^* = \min\{d_{max}, \lceil \frac{100-i}{2} \rceil\}$ , where  $d_{max}$  is an artificial limit on the quantity of dice to be rolled. No rational player would wish to roll more than  $\lceil \frac{100-i}{2} \rceil$  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.

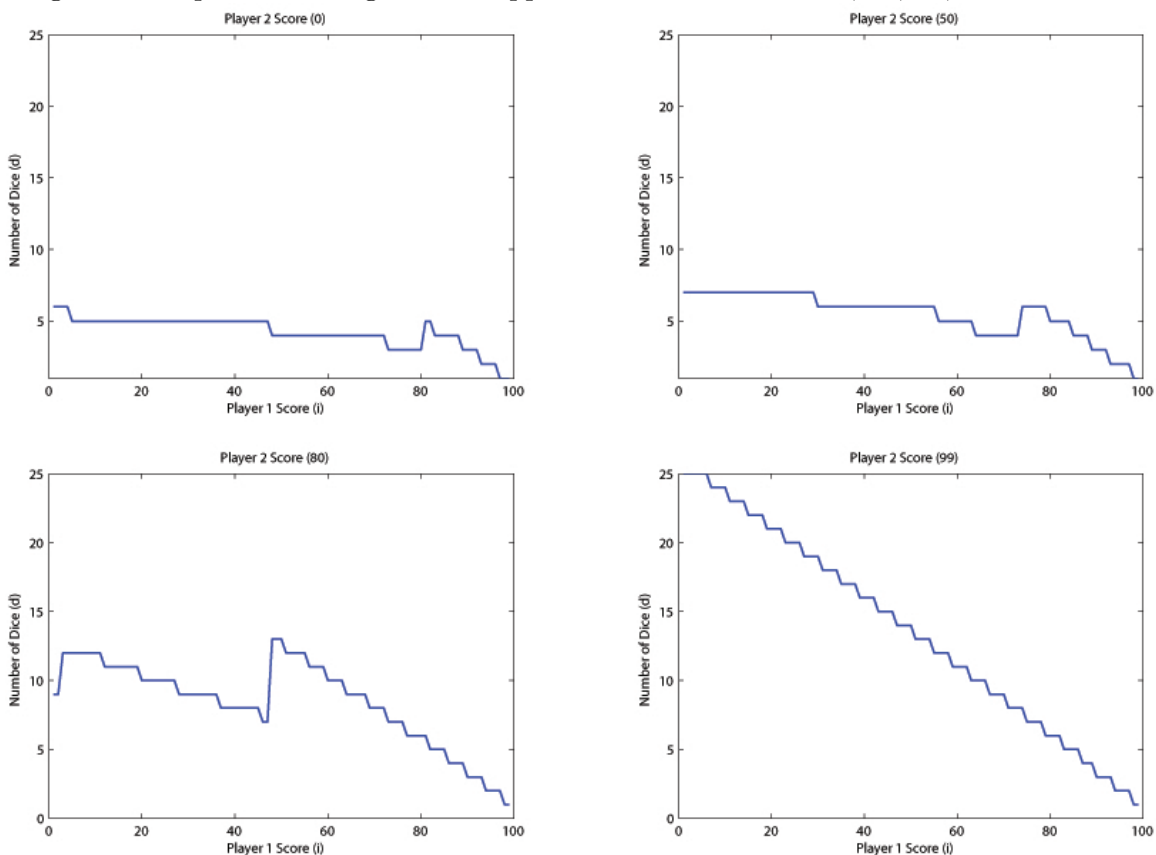
score runs along the  $i$ -axis while Player 2's score is on the  $j$ -axis. The  $d$ -axis indicates the number of dice that Player 1 should roll given that the game state is  $(i, j)$ . The solution is particular to  $d_{max} = 25$ , and Neller and Presser inform us that the optimal solution remains the same for all  $d_{max} \geq 26$  [28].

Figure 3.2: Optimal Solution for the Two-Player Game of Hog and 25 Dice Maximum



It is noteworthy that at state  $(0, 0)$ , Player 1 should roll 6 dice, not 5, which we previously determined to have the same expected value. This suggests that when scores are tied, even if the expected value is equal, it is ultimately preferable to choose the quantity of dice which leads to a greater variance of scores than a safer opportunity to score. In the horizontal plane, the line  $j = i$  consists of all game states in which the scores are tied. As we move to the right from  $(i, j)$  (i.e. we increase  $i$  while keeping  $j$  fixed) we enter game states in which Player 1 is in the lead. We find that  $d$  starts to fall since it is important for Player 1 to try and maintain the lead by choosing quantities of dice with a greater probability of obtaining a nonzero score. On the other hand if we move up from  $(i, j)$  (i.e. we increase  $j$  while keeping  $i$  fixed) Player 1 falls increasingly behind and

Figure 3.3: Optimal Strategies with Opponent's Score Fixed at 0, 50, 80, and 99 Points




must increment his choice in number of dice because he needs the added variance to catch up. The former situation is easier to observe by looking at cross sections of the optimal solution; such is provided in Figure 3.3.

It is initially a bit curious to find that in three of the graphs there is a sudden jump in the optimal number of dice. As it so happens, though, a player will always reach a game state in which from that point onward it is in the player's best interest to play as if the next turn will be his last chance to win (i.e. choose  $\frac{100-i}{4}$  dice). When that change occurs depends on his opponent's score since his opponent is trying to make the exact same decision. For instance, when Player 2 has zero points and Player 1 has 80, Player 2 will soon need to roll all twenty five dice or he will almost certainly lose. Once Player 2 makes that decision, Player 1 must try to win as quickly as possible (by rolling the minimum number of dice which will on average provide a win) since the more opportunities he allows Player 2 to roll all 25 dice, the more likely Player 2 might actually

win in one shot. So we would expect a kind of symmetry in these decisions; that is to say that if there is a jump at  $(0,80)$ , there should be a similar jump in the neighborhood of  $(80,0)$ . This is observable in the two-dimensional view of the optimal solution provided in Figure 3.4. Because of the interaction between the players' score and the fact the choice in dice must be discrete, it is not unreasonable for the change in  $d$  along the curve seen in Figure 3.4 to be quite large.

Figure 3.4: Optimal Solution for the Two-Player Game of Hog and 25 Dice Maximum, 2D View



Neller and Presser point out that the solution for Hog is very similar to the solution for Pig. In fact, we can actually approximate the optimal number of dice to roll in Hog by taking the optimal roll/hold boundary in Pig and dividing it by the average nonzero score per die (i.e. 4) [28]. The similar nature of the games and their solutions will provide for an interesting comparison of results because in absence of framing effects, one would expect behavior to be relatively the same. We, however, predict significant differences in performance. In Pig, every time a player is faced with a roll/hold decision, he is given the choice between  $H$ : *Current turn total,  $x$  points, with certainty* or  $R$ : *Another decision at  $x^*$  points,  $x^* \in \{x + 2, x + 3, \dots, x + 6\}$ , with probability .83 and 0 with probability .17*. Then given the insights of prospect theory, it is not surprising that players in Pig often prematurely choose  $H$  since they tend to overweight sure outcomes. In contrast, a player

in Hog faces a single choice between numerous uncertain outcomes. Thus, in support of prospect theory, we expect players to be more risk-seeking in the game of Hog.

The nature of the games may also make Hog a more conducive environment for learning over multiple games. Although players in Pig did not approach the optimal strategy after repeated play, results from Feldman and Morgan indicate that even though their Hog players tended to “exaggerate the effect of getting a zero for any single turn”—on average they predicted that optimal number of dice would be three—after rolling various fixed numbers of dice ten times each, players’ revised predictions for the optimal number of dice verged on five [13]. Therefore if players in Hog spend some time in earlier games varying roll decisions to collect data, they may easily come closer to the optimal solution in later games.

## 4 The Experiment

To secure the 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. Two different screenshots of this computer-simulated version of the game can be viewed in Figures 4.1 and 4.2<sup>11</sup>. A player can at all times see the number of games either he or his opponent have already won. When it is his turn, each player chooses the number of dice he would like to roll by moving the slider to the desired quantity, and both he and his opponent can view the numbers rolled. Players should consider all simulated dice to be fair and six-sided. Data collected from over thirty two thousand rolls corroborate this claim (please see Table 4.1).

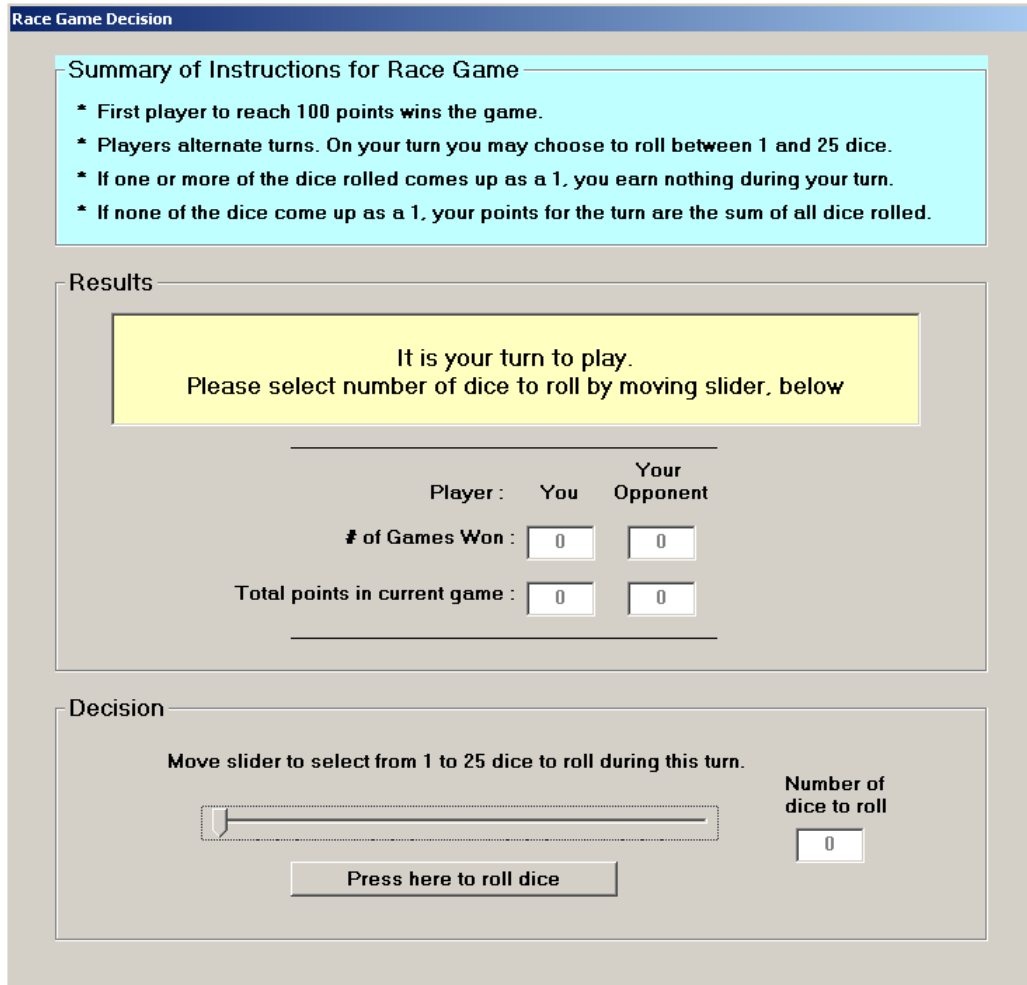
Table 4.1: Frequency of Numbers Appearing on Electronic Dice

	1	2	3	4	5	6
Frequency	5424	5301	5346	5386	5262	5376
Percent	0.168998	0.165166	0.166568	0.167814	0.163951	0.167503

*Note: Total number of rolls was 32095.*

<sup>11</sup>At game state (0,0) a player’s optimal strategy is 6 dice, which yields a probability of winning of 0.530023849. In Figure 4.2, the player chose to roll 16 dice which still yields a probability of winning of 0.511307475 since it is so early in the game. However, rolling a nonzero score with 16 dice only has probability 0.054088. So this player was pretty lucky!

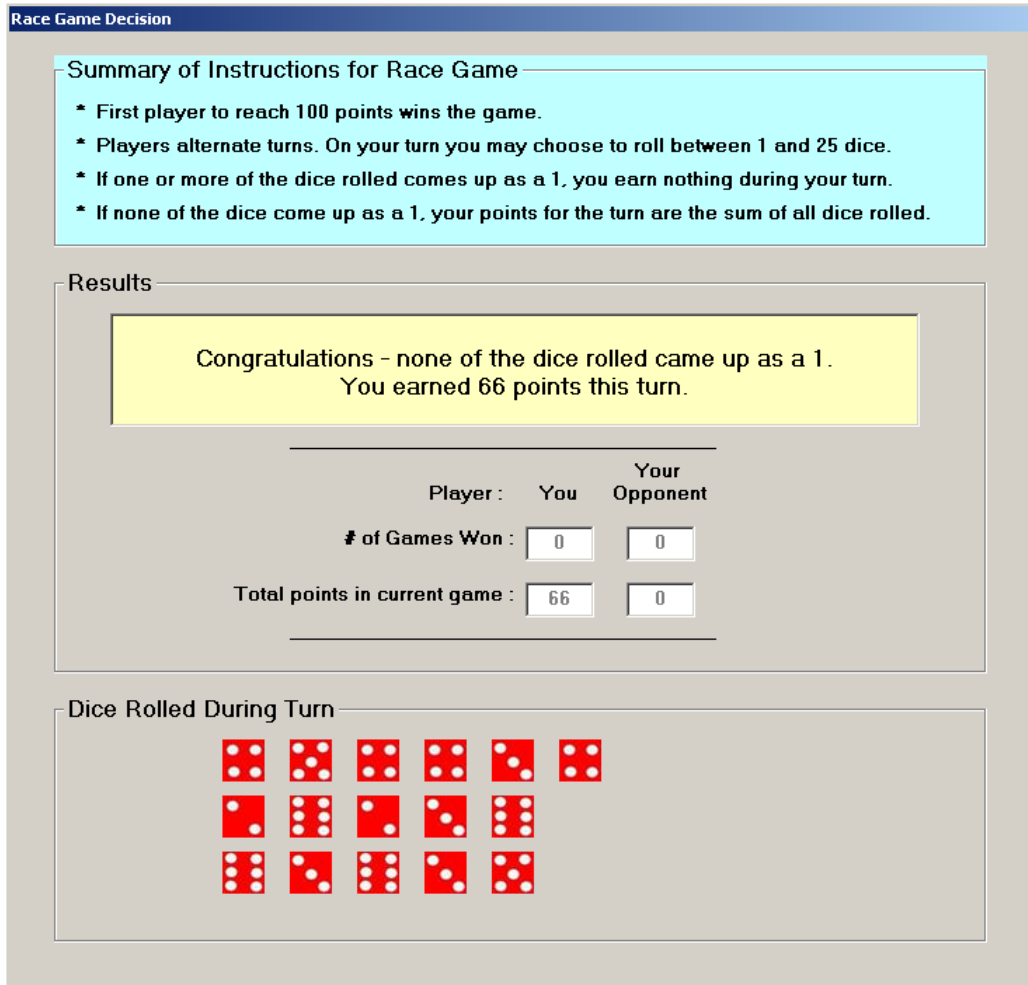
Figure 4.1: Screenshot of Hog Dice Game Software at Initial Turn



Seale, Rapoport, and Stein used two different treatments in their study involving the dice game Fig. We followed their example and used the same conditions. In the first, both players begin the game with zero points and one of them is randomly chosen to go first. The game ends when one player reaches 100 points, and in each subsequent game, the winner of the previous game allows his opponent to go first. In the second treatment one of the players is randomly chosen to start out with fifty points and his opponent goes first; they alternate who gets the 50-point advantage in each game thereafter. The optimal decision does not depend upon how players come to a particular situation, so by artificially placing players in a particular game state we hope to eliminate any biases which may develop from the initial stages of the game. From the cross-sectional data found in Figure 3.3 it is clear that a person starting with a 50 point deficit should roll a greater number



Figure 4.2: Screenshot of Hog Dice Game Software Following Successful Roll



of dice to try and make up for the point disparity. In Pig, however, Seale, Rapoport and Stein, found that players tend to be insensitive to the number of points they lead or trail their opponents [37].

#### 4.1 Procedures

Subjects were recruited from two sections of a Principles of Management and Organizational Behavior course at the University of Nevada Las Vegas and were provided with both a monetary and an academic incentive to participate. Students in the class received a flier advertising the opportunity and were asked to sign up for a particular session via email correspondence with the experimenter

(please see Appendix C).

We collected data 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 were paid the maximum of either \$10, or \$4 per game won plus a \$5 participation bonus; hence, earnings spanned \$10 to \$25 with the vast majority earning \$13 or \$17 (2 or 3 games).

Sessions lasted approximately an hour, with some subjects finishing far earlier and several staying much later. Participants were invited into the computer lab to be seated at any open computer and 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 over an informed consent form and to read the instructions provided in Appendix D. Experiment proctors were available to answer questions both before and throughout the duration of the games. Upon completion, players were asked to fill out the survey found in Appendix E, after which they received payment for their participation before exiting the lab.

During the length of the session, keyboards were removed from each computer station so that players could not be distracted or have access to any calculator, spreadsheet, or internet resource. However, each subject was afforded 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 with the hope that it may shed insight into how or why certain decisions were made. Although such 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 three main issues we wish to address: (1) How do players generally perform? (2) Does player performance change over time? (3) Is the performance of players in the

game of Hog superior to performance of players in the game of Pig? These in turn lead to further questions, all of which require us to first define a measure of performance, some metric for capturing departures from optimality.

One such metric is  $\delta$ , 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 the player should have chosen at that game state to maximize the probability of winning. In this way we might calculate an average  $\delta$  per player to capture his overall performance, or calculate an average  $\delta$  per player, per game, to observe his performance over time. The smaller the  $|\delta|$ , the more optimal the player. A problem with using 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  $\delta$  close to zero, indicating an optimal player, when in fact  $\delta$  per decision could have been quite large. Furthermore, consider a player who in game 1 has an average  $\delta$  determined from the set  $\{2, 4, 3, 4, 2, 4\}$  and in game 2 from the set  $\{11, -6, 7, -8, 9\}$ . Then  $\delta_1 = 3.17$  is greater than  $\delta_2 = 2.16$ , which could lead to an erroneous conclusion that player 1's performance has improved.

We might try to rectify this concern 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,  $|\delta|$  still fails as a perfect measure. There are inherent problems with  $|\delta|$  which cannot be fixed. Being two dice away from optimal is not twice as "bad" as being one die away from optimal, and a difference of  $x$  dice from the optimal solution could have vastly different consequences depending on the state of the game. To illustrate, if Player 1 rolls 6 dice at state  $(0, 0)$ , he maximizes his probability of winning, which is 0.5300. If he instead chose 5 dice his probability of winning would be 0.5299 (a difference of 0.0001) and 4 dice would be 0.5281 (a difference of 0.0019). Now consider the game state  $(92, 88)$ . At the optimal choice of 3 dice, Player 1's probability of winning is 0.7371. If he instead chose 1 die, only two fewer, it reduces the probability of winning to 0.4435 (a difference of 0.2936). Therefore,  $|\delta|$  does not provide a consistent picture of how well a player is performing.

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

While none of our metric choices are perfect, each has a particular advantage. From  $\delta$  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), and from  $\alpha$  we can obtain the actual consequence of the deviation at a particular game state. We will thus require the use of  $\delta$ ,  $|\delta|$ , and  $\alpha$  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 5.1 shows the aggregate distribution of all roll decisions for all players from Treatment 1. The horizontal axis indicates  $\delta$  while the vertical axis describes the percentage of the total rolls marked by  $\delta$ . 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 5.2.

Figure 5.1: Treatment 1: Percentage of Rolls Characterized by Delta





Figure 5.2: Treatment 2: Percentage of Rolls Characterized by Delta



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% level, we classify the player as conservative. If significantly greater than zero at the 5% level, we classify the player as aggressive. Remaining players are considered neutral. Given the problems with the  $\delta$  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 point total. From our prior study of the optimal solution, we understand that a player must continually raise his choice in number of dice if he falls increasingly 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 games

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 no more than 40 points. For each interval we calculate the average optimal number of dice for the given game states and also determine the actual average number of dice rolled for those same game states.

A graph of this data for Treatment 1 is available in Figure 5.3 while a graph for Treatment 2 can be found in Figure 5.4. The blue markers represent the optimal averages while the red 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. All assumptions for a one-way ANOVA are met or corrected for, i.e. the dependent variable, the average number of dice, is continuous and approximately normally distributed in each category as confirmed by a series of Q-Q plots; cases are independent; and although Levene's test indicates that equality of variances between independent groups is violated, we use Welch's test for significance which is robust to this failure. The ANOVA results indicate significant differences among the optimal means in both conditions at the 0.05 level. Further pairwise comparisons using Bonferroni's test indicate an adequate number of statistically significant differences in means to conclude that there is a downward trend through the first six intervals. The same can be said of the actual means in both conditions.

In Treatment 1, average observed number of dice was less than or equal to average optimal number of dice in all 10 intervals. Treatment 2 showed similar results save two intervals in which average observed number of dice exceeded average optimal number of dice. The results are promising; they suggest that players in aggregate were aware of the scoreboard and took the game state into account when making their 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 participants indicated that they consciously increased their choice in number of dice when they found themselves falling behind.

Lastly, we are interested in gauging how far from optimal players were on average. It is more logical to focus on players' average  $|\delta|$  than average  $\alpha$  because differences in alpha are small and

Figure 5.3: Treatment 1: Average Observed and Optimal Numbers of Dice By Score Discrepancy


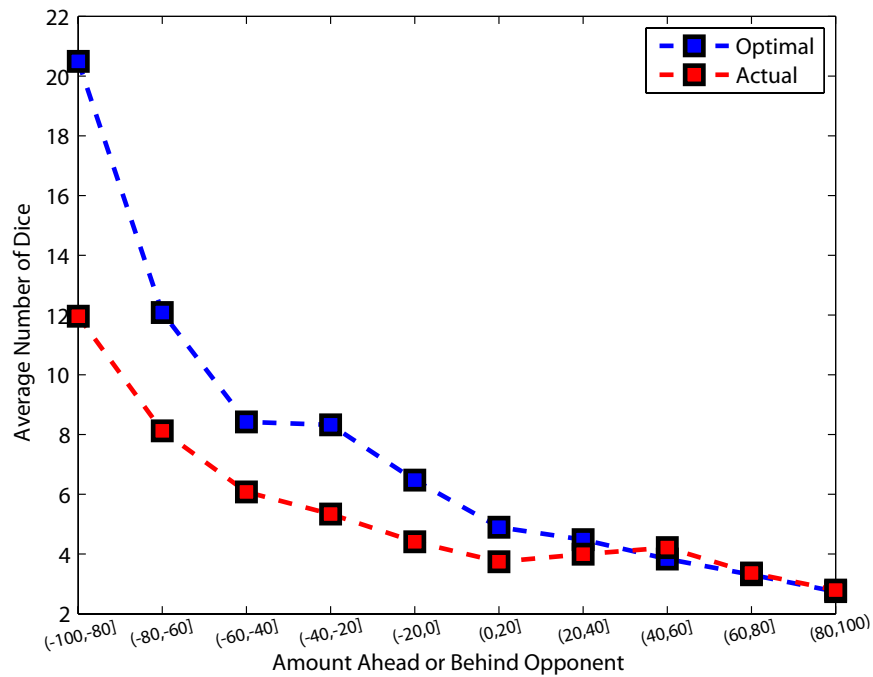



Figure 5.4: Treatment 2: Average Observed and Optimal Numbers of Dice By Score Discrepancy



players themselves are unaware of this measure. We therefore calculate  $|\delta|$  for each player across all five games and report the results for both treatments in Figure 5.5. 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  $|\delta| \in [2, 3)$ , etc. In both treatments the majority of players were on average only two dice away from optimal. Note that we cannot comment on how “bad” two dice away from optimal really is (in terms of the difference in the probability of winning) because that depends on the game state in which it occurred.

Figure 5.5: Per Player Frequency of the Average Magnitude of Differences Between Optimal and Observed Numbers of Dice



Thus, overall we find that that players are far more likely to be conservative than aggressive in play, that they do adjust their decisions based upon the state of the game, and on average tend to be no greater than 5 to 7 dice away from the optimal solution. While these discoveries offer important behavioral insights, they do not allow us to make any sort of general conclusions such as “player performance was poor” or “player performance was good.” Any such determination would ultimately be based on an arbitrary classification, i.e.  $|\delta| > x$  is “poor” while  $|\delta| \leq x$  is “good” for some discretionary  $x$ . To remain systematic, we will rate player performance depending on whether players approach the optimal solution after repeated play.



## 5.2 Improvement


Given the complex nature of the optimal solution to the game of Hog, we would not expect players to know and use it. However, since the optimal strategy maximizes a player's chance of winning the game, if a player's decisions get closer to the optimal decision strategy after repeated play, we call this improvement since the player is "improving" his chances of winning. Detection of improvement over the course of the five games would signify that feedback and experience were adequate for learning from prior suboptimal decisions.

Since each subject played all five games, a one-way repeated measures ANOVA is appropriate for discerning differences in subject behavior over time, provided we meet a few additional assumptions. The repeated measures ANOVA allows us to reduce error variance due to individual differences, thereby allowing for greater power in distinguishing the main effect. Our choice in dependent variable, average  $\alpha$ , is continuous. A Q-Q plot of the variable indicates that it is approximately normal, and in any case, the statistical tests we employ are considered robust to departures from normality. We employ Mauchly's test to check for the homogeneity of covariances; if the sphericity assumption is violated, we apply a Greenhouse-Geisser correction. For post hoc comparisons, we use a Bonferroni test to determine the pairwise interactions.

Average  $\alpha$  serves as the best measure for observing differences in performance since it quantifies the consequence of each decision. The repeated measures ANOVA results for Treatment 1 do indicate a statistically significant difference in performance. From the pairwise comparisons we discover two relationships to which this difference is attributed. Average  $\alpha$  in game 1 is statistically significantly different from both game 4 and game 5 at the 0.05 level. We graph average alpha per game in Figure 5.6; the game number lies on the horizontal axis and average  $\alpha$  per game spans the vertical axis.

For Treatment 2, we obtain statistically insignificant repeated measures ANOVA results. The graph of average  $\alpha$  for Treatment 2 (see Figure 5.7) supports the lack of distinguishable change in performance. However, it is possible that since different players have the 50 point advantage in games 1, 3, and 5 than in games 2 and 4, we lose power by testing all five games in conjunction. If


Figure 5.6: Treatment 1: Subject Performance Across Games



there is a psychological reaction to the point disparity which affects performance, it would be more appropriate to separate the games by parity and further segregate by whether the player went first or second. This generates four subsets of data: players who went first in games 1, 3, 5; players who went second in games 1, 3, and 5; players who went first in games 2 and 4; and players who went second in games 2 and 4. Each of these subsets also yield statistically insignificant ANOVA results for average  $\alpha$ .

Treatment 2 additionally obliges us to check a final measure. One could argue that since there are 10,000 possible game states, the chances that the same state is revisited over multiple games is fairly low. This may result in poor learning since prior experience may not directly apply to a given situation. However, in Treatment 2 we can guarantee that there is one game state that each individual will reach either two or three times—the original game state of (0, 50). Therefore, in lieu of the average  $\alpha$  metric, we might instead use the opening  $\alpha$  metric for game state (0, 50). In doing so we still find statistically insignificant results with the two subsets of data to which this metric applies. Thus, although we find slight evidence of learning in Treatment 1, we conclude that performance was stagnant over time in Treatment 2.

Figure 5.7: Treatment 2: Subject Performance Across Games



### 5.3 Survey Data

The quantitative data studied thus far paints only a partial picture of decision making in the game of 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 why we observed the behavioral patterns mentioned in prior sections, we turn to qualitative data gathered from a survey (please see Appendix E) administered to all subjects regardless of treatment type.

Table 5.1 summarizes demographic information reported by participants in the study. Response rates are provided in parentheses following each category. According to the most recent student profile made available by the Office of Institutional Analysis & Planning at UNLV, our sample is highly representative of the UNLV undergraduate population in all categories excepting gender (for which the numbers should have been reversed) [32].

In addition to collecting personal data, we used the survey to determine each player's knowledge of probability theory, since a lack of such understanding might help explain suboptimal behavior. Although we asked players the number of classes in which they have learned some probability

Table 5.1: Aggregate Survey Data: Demographic Information

---

*Response rates are provided in parentheses.*

---

- *Age (0.94):*

<b>19-22</b>	<b>23-26</b>	<b>27-30</b>	<b>31-34</b>	<b>&gt;34</b>
0.52	0.20	0.16	0.04	0.08

- *Gender (0.88):*

<b>Female</b>	<b>Male</b>
0.44	0.56

- *Ethnicity (0.91):*

<b>Hispanic/ Latino</b>	<b>Not Hispanic/ Latino</b>	<b>Unknown</b>
0.23	0.76	0.01

- *Race (0.97):*

<b>American Indian</b>	<b>East Asian</b>	<b>South Asian</b>	<b>White</b>	<b>Hawaiian/Pacific Islander</b>
0.01	0.13	0.03	0.47	0.09

<b>Black/African American</b>	<b>Multiple Race</b>	<b>Other</b>	<b>Unknown</b>
0.04	0.09	0.12	0.01

---

theory/statistics, we specifically test knowledge by asking three straightforward multiple choice questions pertinent to the game. Including both treatments, 122 subjects participated in the survey; 119 answered all of the questions found in Table 5.2, a response rate of 97%.

The vast majority of respondents correctly identified 5/6 as the the probability of not rolling a 1 using one die. However, less than half of the participants 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 participants 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 test the correlation between the number of classes taken and the number of questions answered correctly in the survey. The correlation coefficient obtained was 0.067 and was not significant at the 0.05 level, connoting

Table 5.2: Aggregate Survey Data: Knowledge Base Questions

---



---

*Response rate: 0.97; correct answers are starred*

---

• *How many classes (post high school) have you taken in which you have learned some statistics and/or probability theory?*

<b>None</b>	<b>One</b>	<b>Two</b>	<b>Three</b>	<b>Four or More</b>
0.04	0.15	0.41	0.29	0.11

• *What is the probability of not rolling a 1 using one die?*

<b>1/2</b>	<b>2/3</b>	<b>3/4</b>	<b>4/5</b>	<b>5/6*</b>
0.02	0.02	0.00	0.01	0.95

• *What is the probability of not rolling any 1's using two dice?*

<b>1/2</b>	<b>7/12</b>	<b>2/3</b>	<b>25/36*</b>	<b>5/6</b>
0.05	0.10	0.13	0.39	0.33

• *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?*

<b>More likely to roll a '1'</b>	<b>Has no effect on the current turn*</b>	<b>Less likely to roll a '1'</b>
0.19	0.78	.03

---

negligible correlation between the two.

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 constraints. Therefore, in order to isolate the various effects contributing to performance we must control for a person's risk tolerance. We develop 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  $\geq 0.70$  is considered acceptable [31].

Subjects' risk scores are reported in Table 5.3. These numbers do not allow us to make a

Table 5.3: Aggregate Survey Data: Individual Characteristics

---



---

*Response rates are provided in parentheses.*

---

- *Risk-taking Score [Scale: 1-5, the lower the score, the more risk averse.] (1.00)*

<b>[1,2)</b>	<b>[2,3)</b>	<b>[3,4)</b>	<b>[4,5]</b>
0.03	0.30	0.56	0.11

- *Do you gamble? (0.99)*

<b>Never</b>	<b>Sometimes</b>	<b>Often</b>
0.36	0.59	0.05

—

<b>Yes</b>	<b>No</b>
------------	-----------

- *In general, do you feel lucky? (0.99)*

0.53	0.47
------	------

- *Did your strategy change from game to game? (0.98)*

0.59	0.41
------	------

- *Did your strategy ever change based upon the behavior of your opponent? (0.99)*

0.44	0.56
------	------

- *Do you feel your performance improved from game to game? (0.93)*

0.36	0.64
------	------

- *After playing the game once did you want to play again? (0.92)*

0.81	0.19
------	------

- *Did you enjoy playing the game? (0.93)*

0.81	0.19
------	------

---

judgment such as “a player with a risk score between 1 and 2 is risk averse whereas a player with risk score between 4 and 5 is risk-seeking.” 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 5.3 notes players’ responses to a number of additional characteristics which have the potential to affect performance.

Finally, participants were asked to hand in their scratch paper with their surveys. Two thirds of participants returned blank sheets of paper. Of those which were not blank, many consisted mainly of doodles. Of interest is the observation that if we separate the use of scratch paper by treatment, 38% of participants in Treatment 1 “used” their scratch paper while only 27% of participants in

Treatment 2 “used” theirs.

To determine if any of the aforementioned characteristics had a perceivable effect on performance, we ran a multivariate regression with players’ average  $\alpha$  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<sup>12</sup>. The data are approximately normal and were examined for large outliers. A plot of predicted values against the residuals supports linearity and homoskedasticity, so that a multivariate regression is not unwarranted. Results are presented in Table 5.4. 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 that on average, an increase in 0.01 of a player’s opponent’s  $\alpha$  will increase a player’s  $\alpha$  by 0.00522, or the more suboptimal the performance of an opponent, the more suboptimal the performance of a player, a positive relationship. Several written responses support this finding; a dozen participants either admitted to copying their opponents or noted that their opponents copied them. Furthermore, a player who gambles will on average have an average  $\alpha$  that is 0.00338 greater than a player who does not gamble, i.e. a player who gambles on average performs worse than a player who does not. All other variables have no statistically significant effect on performance. The adjusted  $R^2$  value was  $\bar{R}^2 = 0.2527$ , meaning that the two statistically significant predictors of player performance in Hog account for 25.27% of the variation found in average  $\alpha$ .

## 5.4 In Comparison to Pig

Our results indicate that the performance of subjects in Hog is marked by both similarities and differences with the behavior exhibited by players of Pig. In particular, the vast majority of decisions made in both games were conservative and players seemed positively influenced by the decisions of

---

<sup>12</sup>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.

Table 5.4: Regression Results

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

\* indicates significance at the 5% level


their opponents. However, unlike Pig players who appeared incognizant of how much they trailed or led their opponents, Hog players seemed to account for the game state when making their decisions [37]. Furthermore, both treatments in Pig reveal that players on average failed to get closer to the optimal decision strategy, whereas at least in Treatment 1, Hog players demonstrated modest



improvement from repeated play [37].

What remains to be determined is which set of players, on average, made the more optimal decisions overall. In Pig, the optimal decision strategy is found by determining which is greater at every decision state: the probability of winning with a hold, or the probability of winning with a roll. In this way we can determine an  $\alpha$  of  $|P_{hold} - P_{roll}|$  which will capture the consequence of a particular hold decision. We thus compute an average  $\alpha$  per player in Pig based upon each player's hold decisions and compare it to the average  $\alpha$  levels found in Hog. Since average  $\alpha$  is approximately normal, as indicated by Q-Q plots of both sets of data, we use an independent t-test to check for significance. In both Treatment 1 and Treatment 2 we find statistically significant differences in performance. Figure 5.8 indicates that Hog players outperform their Pig-playing counterparts in both treatments. In Treatment 1, all Hog players had average  $\alpha$  scores less than or equal to 0.026 while only 27% of players in Pig were at or below the 0.026  $\alpha$  level. Similarly, in Treatment 2, all Hog players had average  $\alpha$  scores less than or equal to 0.037 while 56% of Pig players were at or below the same  $\alpha$  level.

Figure 5.8: Cumulative Distributions, Hog vs Pig



## 5.5 Heuristic Solutions

Across five games, players in Treatment 1 averaged 60 decisions and players in Treatment 2 averaged 38 decisions. Therefore, although the optimal solution to the game of Hog is not obvious, players clearly have an opportunity to test different strategies to win. Out of 122 subjects, only one reported

having deliberately rolled various quantities of dice in order to test the probability of rolling a one. Nevertheless, the majority of respondents indicated that they typically chose a quantity of dice within a specific range and revised this decision set as their experience grew.


In order to determine the effectiveness of various heuristic strategies, we ran multiple simulations using the most common and obvious ones. Each statistic reported has been determined from the simulation of 1 million games against an opponent using the optimal strategy. We first test the optimal solution itself. Since  $P_{0,0} = 0.5300$ , theoretically, a subject who plays optimally and goes first will win 53% of the time. Simulation results of an optimally performing player are in line with predicted results; the optimally playing individual won 52.89% of the games played.

Recall from our previous analysis that 5 and 6 dice both maximize the expected value of points. A player aware of this fact might adopt the following strategy: roll 6 dice if behind or tied with the opponent and five dice if ahead of the opponent. As we touched upon before, the insight for this decision strategy stems from the idea that a player who is behind or tied with his opponent would want to narrow the gap or take the lead by taking advantage of the larger variance of scores provided by six dice. On the other hand, a player who is ahead of his opponent would wish to play the more conservative choice of five dice in order to maintain that lead. Simulated results implementing this strategy yield a win percentage of 47.16%, indicating that the heuristic performs remarkably well.

Without any kind of knowledge about the probabilities involved, a player might adopt a very simple heuristic which we name the trial and adjustment strategy. In this strategy a player will randomly choose a number of dice for his opening roll. On each subsequent roll he will increase the number of dice by one if he received a nonzero score during his immediately previous turn or decrease his choice of dice by one if he obtained zero points during that turn. Of course this choice is bounded by a minimum of 1 die and maximum of 25 dice. Simulation results for each possible starting roll are provided in Figure 5.9. The horizontal axis denotes the choice of the starting roll while the vertical axis reports the percentage of games won using the heuristic. We draw a dashed horizontal line at 0.53 for easy comparison with the optimal strategy. It is clear that the effectiveness of the strategy is dependent upon the starting quantity of dice. However, even the

worst initial choice provides a win 17.53% of the time, and the appropriate starting choice can lead to a win 43.69% of the time. The steady decline in the percentage of wins beyond eight initial dice is to be expected. Each additional die increases the probability of rolling a one and will likely be a waste of a trial.

Figure 5.9: Trial and Adjustment Heuristic Performance in 1M Games




The other straightforward heuristic that a player might employ, and which was indeed the most common strategy reported, is what we call the consistent roll strategy. When a player uses this strategy he chooses some quantity of dice and repeatedly uses that quantity, turn after turn, regardless of past outcomes. However, we assume that he is forward thinking and does not use the heuristic blindly. That is to say that as he approaches the goal threshold, his decision becomes  $d = \min\{\text{heuristic choice}, \lceil \frac{100-s}{2} \rceil\}$ , where  $s$  is the player’s score. So a person using the “always roll 10” strategy will choose 2 dice, not 10, when he finds himself at 96 points.

Simulation results for the consistent roll strategy are provided in Figure 5.10 where the horizontal axis marks the possible choices of dice and the vertical axis relates the percentage of wins yielded by each choice. We again indicate the optimal strategy by a dashed horizontal line at 0.53. This graph is not as smooth as the graph for the trial and adjustment heuristic. However, the

jaggedness at 7, 9, and 13 dice may be attributed to the fact that these are the transitions to an average fewer number of required successful rolls. For example, 6 dice on average yield 24 points and would require 5 successful rolls, whereas 7 dice on average yield 28 points and would require 4 successful rolls. The local minimum at 17 dice is expected since it requires the greatest number of successful dice over the course of the entire game. Another aspect of this data which may seem surprising is that rolling 25 dice can lead to a win 11.88% of the time. This is not as unrealistic as it may initially seem. Given that the expected nonzero score of one die is 4 points, if a player rolls 25 dice, on average he will win the game in one turn. Since an optimally playing opponent would on average need 5 *successful* rolls, the total number of turns in 1 million games will likely exceed 10 million. If we multiply this by the probability of rolling a nonzero score with twenty five dice, i.e. 0.0104 (see Table 3.2), the simulated result is not unreasonable.

Figure 5.10: Consistent Roll Heuristic Performance in 1M Games



Although the trial and adjustment heuristic is the better strategy overall, the consistent roll heuristic outperforms the trial and adjustment heuristic by as much as 7% between 4 and 8 dice inclusive. Consistently rolling 6 dice provides a win 49.24% of the time, making it the best heuristic strategy available of those analyzed here. Thus, overall we see that the use of an appropriate heuristic can be very effective.

## 6 Discussion

Our investigation of dynamic decision making with respect to the dice game Hog has led to many interesting results. One of the most important of these is that player performance is largely stagnant despite repeated play, and even when improvement is detected, it is minimal and somewhat delayed. In this section we provide potential explanations for this occurrence. Of equal importance is understanding and accounting for the similarities and differences in performance between players in the game of Pig and players in the game of Hog. These discussions will reveal aspects of the experiment open for potential improvements and further exploration.

In Treatment 1, we found slight movement towards the optimal strategy over the course of several games. Performance was actually static during the first three, after which there was a sudden improvement that lasted through the fifth. This phenomenon is in agreement with existing literature and research. According to the Praeger Handbook of Learning and the Brain, “learning requires repetition” [11]. Howe elaborates that practice is a large contributor of success in both physical and intellectual skills [17]. Without “adequate time to integrate, experiment, and reflect,” learning may not flourish [11]. Therefore, 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 don’t think there were enough # of games for me to practice and learn and improve upon.” This is considerably disheartening; as Thaler points out, some of “the most important of life’s decisions, such as choosing a career or spouse, offer only a few chances for learning!” [40].

Determining why we only found learning in one treatment requires careful examination of the metric used to measure improvement. Notice that average  $\alpha$  in Treatment 1, game 1, was a mere 0.0132. This means that on average, the difference in the probability of winning between the optimal strategy and players’ actual decisions was only 0.0132. This exposes how even a large difference in the number of dice between the optimal and actual quantity chosen at a particular decision may have a minute difference in terms of the change in the probability of winning. It is possible that such a small difference is not perceivable to the subjects playing the game. If this is the case, the game of Hog suffers from a “flat maximum problem” which refers to a situation in

which “even wide discrepancies from the optimal strategy have very little effect on the outcomes of the decisions” [7]. Then feedback is far too poor for decision makers to learn from experience. It is not clear that this is true of Hog, however, because it is not consistent with a sizeable amount of the survey data. Several participants who reported an original strategy of “rolling a lot” or “risking a lot” mentioned that they changed their strategies to “roll fewer.” One specific respondent who had an original strategy of “2-3 dice” came to the conclusion that “5 is a lucky number.” As seen from the heuristic simulations, five is indeed an excellent heuristic. Therefore, even if the change in number of dice rolled does not largely affect the probability of winning, the fact that some players were at least improving in terms of heuristic strategies implies that feedback was meaningful after all.

Another reason for why performance might fail to improve is player attitude toward the nature of the game and/or a lack of knowledge about the probability concepts relevant to Hog. Particularly, survey results identified a fair number of participants who felt they had no influence over the direction of the game; over 12% of respondents referred to luck, randomness, and chance as the main culprit for their lack of strategy or improvement. This is a type of hindsight bias. Given biased memories, players “will find it very difficult to distinguish between a bad decision and a bad outcome” [40]. Specific responses such as “I have no control over the outcome” and “as games went on I realized it was more about luck and stopped taking notes” support this idea and suggest that players may not have made any conscious effort to increase their probability of winning. This could have been due to a lack of understanding of probability theory. Admittedly, the regression results suggest that knowledge of probability does not have an effect on player performance. However, recalling that there was no correlation between the number of classes taken and the number of probability questions answered correctly, this might have been a poor proxy for players’ knowledge of mathematical concepts applicable to the game. Another reason for lack of effort, particularly for Treatment 2, could have been that players felt the outcome of the game was predetermined due to the large initial disparity in points. Far fewer individuals used their scratch paper in the second treatment, supporting a more prevalent atmosphere of helplessness in changing the outcome of the game.

Finally, the survey data reveal that another cause for poor performance could be lack of mo-

tivation. Even if the game of Hog is suitable for learning, improvement will not take place if participants do not make any attempt to be systematic in their choices. Reaction to the experiment was mixed. One individual believed that Hog “could become addictive if available as a game where you play the computer (like some of the solitaire versions).” However, not everyone shared the same enthusiasm. Many participants admitted to being bored with the game (e.g. “the program was boringly[sic] slow so didn’t care much”) or that they were not at all serious about their performance (e.g. “randomly rolling for laughs”). It is unclear what fraction of players belong to this latter group. If substantial, we must be cautious of our conclusions since they may be based on decisions of players who were not properly motivated.

The most frequent complaint about the game was the pace. The Hog game 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 one 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 Table 5.1 and Table 5.2 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 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 shouldn’t kepted[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 appropriate change, if the study were repeated, may 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 the game of 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 players were not in fact playing optimally, it is not clear how poorly players really performed. That is, if a player knows his opponent is playing suboptimally, his

optimal decision strategy may not be the same as when he knows his opponent is playing optimally. As an extreme situation, consider a version of the game without a maximum number of dice in which a player is certain that his opponent will roll 1000 dice at each turn. In this situation since we do not care how long it takes to win but merely want to guarantee a win, it is possible that rolling a small number of dice might actually be the optimal choice. However, if instead he knows that his opponent will always roll a more reasonable number of dice, like 15, the answer is not very clear. To assess this problem we would have to calculate the probability of winning for a player and his opponent separately, unlike our current solution which applies to either player. 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.

While there are many viable reasons for why player performance in Hog had negligible improvement, explaining the differences in behavior between players in Pig and players in Hog is a more complicated task. Our hypothesis of more aggressive play in Hog did hold, in general. The overall ratio of conservative to aggressive rolls was smaller in Hog than in Pig (3 versus 5). Nevertheless, players on average were still far more likely to play conservatively than aggressively. Therefore, framing differences between the two games were not as salient as we initially believed. Even in Hog, rolling fewer dice increases the likelihood of a nonzero score, so the certainty effect is still very strong.

Perhaps a surprising result is that Hog players outperform Pig players despite the fact that Pig players have the advantage of greater feedback. A Pig player who might decide to roll  $x$  number of dice is able to see each die one at a time; he has the opportunity to revise his decision at each stage until he rolls a 1 or holds. A Hog player does not have the luxury of stopping before it is too late; he must make his choice and deal with the consequences. Actually, this result is not peculiar after careful consideration. We liken the situation to playing the stock market. Malkiel emphasizes how a simple buy-and-hold approach usually outperforms more technical strategies [23]. Similarly, standing by a decision quantity outperforms assessing the situation at the individual dice stage. Pig players are thus presented with too much information and have too many opportunities to react. By taking that option away from the player in the game of Hog, the player performs more



optimally. Because they do not have to focus on individual dice, Hog players are more aware of other important pieces of information such as the actual game state. We conclude that feedback needs to be not only meaningful and timely, but should not exceed a critical level. Otherwise players have more information than they can process, and performance declines.

In this study we addressed criticism directed at the complicated and often haphazard nature of most dynamic decision making experimental tasks. We introduce the game of Hog as an ideal study for decision making under uncertainty for its simplicity, timely feedback, and clear objectives. The game of Hog has a straightforward, analytical solution, allowing it to overcome one of the biggest challenges in developing an experimental task—having an ideal decision maker to which we can compare an observed decision maker [7]. Yet despite the simplistic nature of the game and immediate feedback, player performance was largely stagnant even with repeated play. Likely causes include insufficient time for learning, lack of meaningful feedback, and deficient player motivation. Addressing these concerns will require further research in experimental design.

## References

- [1] Art, L. & Mauch, H. (2007). Dynamic Programming: A Computational Tool. In Series J. Kacprzyk (Ed.), *Studies in Computational Intelligence*, 38. New York: Springer.
- [2] Atkins, P. W. B., Wood, R. E. & Rutgers, P. J. (2002). The Effects of Feedback Format on Dynamic Decision Making. *Organizational Behavior and Human Decision Processes*, 88, 587-604.
- [3] Bakken, B. E. (2008). On Improving Dynamic Decision-Making: Implications From Multiple-Process Cognitive Theory. *Systems Research and Behavioral Science*, 25, 493-501.
- [4] Becker, G. S. (1962). Irrational Behavior and Economic Theory. *The Journal of Political Economy*, 70(1), 1-13.
- [5] Bohan, J. F. & Shultz, J. L. (1996). Revisiting and Extending the Hog Game. *The Mathematics Teacher*, 89(9), 728-733.
- [6] Bradley, S. P., Hax, A. C. & Magnanti, T. L. (1977). *Applied Mathematical Programming*. Reading, MA: Addison-Wesley Publishing Co.
- [7] Brehmer, B. (1992). Dynamic Decision Making: Human Control of Complex Systems. *Acta Psychologica*, 81, 211-241.
- [8] Cheney, W. & Kincaid, D. (2008). *Numerical Mathematics and Computing*. Belmont, CA: Thomson Brooks/Cole.
- [9] Cormen, T. H., Leiserson, C. E., Rivest, R. L. & Stein, C. (2009). *Introduction to Algorithms*. Cambridge, Massachusetts: MIT Press.
- [10] Diehl, E. & Sterman, J. D. (1995). Effects of Feedback Complexity on Dynamic Decision Making. *Organizational Behavior and Human Decision Processes*, 62(2), 198-215.
- [11] Doubllass, C. K. (2006) Multimedia Technology. In S. Feinstein (Ed.), *The Praeger Handbook of Learning and the Brain*, 2 (pp. 315-320). Westport, CT: Praeger.
- [12] Edwards, W. (1962). Dynamic Decision Theory and Probabilistic Information Processing. *Human Factors*, 4, 59-73. In J. W. Weiss & D. J. Weiss (Eds.), *A Science of Decision Making: The Legacy of Ward Edwards* (pp. 62-75). New York: Oxford University Press.
- [13] Feldman, L. & Morgan, F. (2003). The Pedagogy and Probability of the Dice Game Hog. *Journal of Statistics Education*, 11(2). Retrieved October 10, 2010, from <http://www.amstat.org/publications/jse/v11n2/feldman.html>.
- [14] Friedman, M. (1966). The Methodology of Positive Economics. In, *Essays In Positive Economics* (pp. 3-46). Chicago: Univ. of Chicago Press.

- [15] Gonzalez, C., Vanyukov, P. & Martin, M. K. (2005). The Use of Microworlds to Study Dynamic Decision Making. *Computers in Human Behavior*, 21, 273-286.
- [16] Hogarth, R. M. (1981). Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgmental Heuristics. *Psychological Bulletin*, 90(2), 197-217.
- [17] Howe, M. J. (1998). *Principles of Abilities and Human Learning*. Hove, East Sussex: Psychology Press.
- [18] Jolls, C., Sunstein, C. R., & Thaler, R. (1998). A Behavioral Approach to Law and Economics. *Stanford Law Review*, 50(5), 1471-1550.
- [19] Jungermann, H. (1983). Two Camps on Rationality. In R. W. Scholz (Ed.), *Decision Making Under Uncertainty* (pp. 63-86). New York: North Holland.
- [20] Kahneman, D. (2003). Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, 93(5), 1449-1475.
- [21] Kleinmuntz, D. N. (1985). Cognitive Heuristics and Feedback in a Dynamic Decision Environment. *Management Science*, 31(6), 680-702.
- [22] Madrian, B. C. and Shea, D. F. (2000). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *NBER Working Paper Series*.
- [23] Malkiel, B. G. (1999). *A Random Walk Down Wall Street: Including a Life-Cycle Guide to Personal Investing*. New York: Norton.
- [24] Mankiw, N. G. (2007). *Principles of Microeconomics*. Mason, OH: Thomson South-Western.
- [25] Mathematical Sciences Education Board. (1993). The Hog Game. In, *Measuring Up: Prototypes for Mathematical Assessment*, 141-155. Washington: National Academy Press.
- [26] Mullainathan, S. & Thaler, R. H. (2000). Behavioral Economics. *NBER Working Paper Series*.
- [27] Neller, T. W. & Presser, C. G. M. (2004). Optimal Play of the Dice Game Pig. *The UMAP Journal*, 25(1), 25-47.
- [28] Neller, T. W. & Presser, C. G. M. (2005). Pigtail: A Pig Addendum. *The UMAP Journal*, 26(4), 443-458.
- [29] Neller, T. W., Russell, I. & Markov, Z. (2005). Solving the Dice Game Pig: An Introduction to Dynamic Programming and Value Iteration. Retrieved April 10, 2010, from <http://cs.gettysburg.edu/~tneller/nsf/pig/index.html>.
- [30] Neller, T. W., Presser, C. G. M., Russell, I. & Markov, Z. (2006). Pedagogical Possibilities for the Dice Game Pig. *Journal of Computing Sciences in Colleges*, 21(6), 149-161.
- [31] Nunnally, J. C. (1978). *Psychometric Theory*. New York: McGraw Hill Publishers.

- [32] Office of Institutional Analysis & Planning. (2009). Undergraduate Student Profile - Fall 2009. Retrieved February 24, 2010 from, [http://ir.unlv.edu/IAP/Reports/Content/UndergraduateStudentProfile\\_Fall2009.aspx](http://ir.unlv.edu/IAP/Reports/Content/UndergraduateStudentProfile_Fall2009.aspx).
- [33] Posner, R. A. (1998). Rational Choice, Behavioral Economics, and the Law. *Stanford Law Review*, 50(5), 1551-1575.
- [34] Printfection, LLC. (2010). Homo Economicus. [graphic]. *Shirts-4-Thought: Change Your Wardrobe, Change the World*. Retrieved January 2, 2010, from [http://www.shirts-4-thought.com/Homo-Economicus/\\_s\\_299479](http://www.shirts-4-thought.com/Homo-Economicus/_s_299479).
- [35] Rabin, M. (1998). Psychology and Economics. *Journal of Economic Literature*, 36(1), 11-46.
- [36] Seale, D. A. (2008). Research Proposal for the Study of Dynamic Decision Making in Jeopardy Race Games. *University of Nevada Las Vegas*, 1-10. Limited Distribution.
- [37] Seale, D. A., Stein, W. E. & Rapoport, A. (2011) Hold or Roll: Reaching the Goal in Jeopardy Race Games. Working Paper, Submitted to *Management Science*.
- [38] Silberberg, E. (1990). *The Structure of Economics, A Mathematical Analysis*. New York: McGraw-Hill, Inc.
- [39] Sterman, J. D. (1989). Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*, 35(3), 321-339.
- [40] Thaler, R. H. (2000). From Homo Economicus to Homo Sapiens. *Journal of Economic Perspectives*, 14(1), 133-141.
- [41] Tversky, A. & Kahneman, D. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2), 263-292.
- [42] Walpole, R. E., Meyers, R. H., Myers, S. L. & Ye, K. (2007). *Probability and Statistics for Engineers and Scientists*. Upper Saddle River, NJ: Pearson Prentice Hall.

## A Solving the Game Hog

Obtaining the optimal solution for the game of Hog requires a highly computationally intensive procedure. In order to solve

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

for each game state, we must first calculate  $\pi(d,k)$ , which may be determined using dynamic programming. We then partition the game states, solving each successive subset using an iterative method.

Although Neller and Presser have already solved the game of Hog, we provide a more thorough solution here. We offer a slightly different but equivalent calculation of  $\pi(d,k)$  and present a proof of its correctness, an exercise that Neller and Presser omit. Furthermore, our use of  $d^*$  in the calculation of  $P_{i,j}$  is an improvement over the Neller and Presser solution that uses  $d_{max}$  since this saves a considerable amount of computation. Finally, while Neller and Presser simply present the solution, we specify two different methods for its determination.

### A.1 A Select Review of Probability Theory

The following is a list of various definitions and theorems in probability theory that are essential to the understanding and proof of the optimal solution to the game of Hog. For further reading, we recommend [42], from which these theorems have been reproduced, as they are.

**Definition A.1.** *The probability of an event  $A$  is the sum of the weights of all sample points in  $A$ . Therefore,*

$$0 \leq P(A) \leq 1, \quad P(\emptyset) = 0, \quad \text{and} \quad P(S) = 1.$$

*Furthermore, if  $A_1, A_2, A_3, \dots$  is a sequence of mutually exclusive events, then*

$$P(A_1 \cup A_2 \cup A_3 \cup \dots) = P(A_1) + P(A_2) + P(A_3) + \dots.$$

**Theorem A.2.** *If an experiment can result in any of  $N$  different equally likely outcomes, and if exactly  $n$  of these outcomes correspond to event  $A$ , then the probability of event  $A$  is*

$$P(A) = \frac{n}{N}.$$

**Definition A.3.** The probability of an event  $B$  occurring when it is known that some event  $A$  has occurred is called a **conditional probability** and is denoted by  $P(B|A)$ .

**Theorem A.4.** If the events  $B_1, B_2, \dots, B_k$  constitute a partition of the sample space  $S$  such that  $P(B_i) \neq 0$  for  $i = 1, 2, \dots, k$ , then for any event  $A$  of  $S$ ,

$$P(A) = \sum_{i=1}^k P(B_i)P(A|B_i).$$

## A.2 Obtaining a Score $k$ , with $d$ Dice

**Theorem A.5.** 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:

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

*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, by Theorem A.2,

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

**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 non-zero score is two. Therefore, the smallest number that could possibly be rolled with  $d$  dice and result in a non-zero 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 1's 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. By Definition A.1 this gives us,

$$\pi(d - 1, 0) + \pi(d - 1, \tau) = 1 \quad (\text{A.2})$$

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 by Theorem A.4,

$$\pi(d, 0) = \pi(d - 1, 0)\pi(d, 0|d - 1, 0) + \pi(d - 1, \tau)\pi(d, 0|d - 1, \tau). \quad (\text{A.3})$$

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  $\frac{1}{6}$ ), using equation A.2 and substitution in A.3 we obtain

$$\pi(d, 0) = \pi(d - 1, 0) + \frac{1}{6}(1 - \pi(d - 1, 0)) = \frac{5\pi(d - 1, 0) + 1}{6}. \quad (\text{A.4})$$

**Case 4.** All other  $k, d$

We use similar logic to that of Case 4. 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, by Theorem A.4,

$$\pi(d, k) = \sum_{r=2}^{\min(6, k-2)} \pi(d - 1, k - r)\pi(d, k|d - 1, k - r). \quad (\text{A.5})$$

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

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

□

### A.3 A Brief Introduction to Dynamic Programming

Dynamic programming<sup>13</sup> is a type of algorithm, i.e. a well-defined list of steps used to transform some given input into a desired output [9]. It is a method for solving complex problems by exploiting the solutions of simpler subproblems. Each of these subproblems is only solved once; the value is stored for later reference, thereby preventing repetitive computations [9]. This technique is mostly applied to problems involving optimization [6].

Cormen, et. al., provide the following procedure for developing a dynamic-decision making algorithm: “(1) Characterize the structure of an optimal solution. (2) Recursively define the value of an optimal solution. (3) Compute the value of an optimal solution. (4) Construct an optimal solution from computed information” [9]. A problem whose optimal solution consists of the optimal solutions to its subproblems is said to have an optimal substructure [9].

When introducing dynamic programming, it is helpful to review a simple application such as the computation of the fibonacci numbers. Although not an optimization problem, we can introduce a dummy decision,  $d$ , which allows it to serve as a nice illustration of how the method works [1]. The fibonacci sequence, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144,..., is defined as

$$f_n = f_{n-1} + f_{n-2}, \quad \text{for } n \geq 3 \quad \text{and} \quad f_1 = f_2 = 1.$$

So we have the dynamic programming functional equation,

$$f(i) = \text{opt}_{d \in D} \{f(i-1) + f(i-2)\}$$

with base case  $f(1) = f(2) = 1$  [1]. Since the initial cases are given, we may move forward calculating one state at a time, a process known as forward induction [6]. The solution of each state is stored so that we have a memoized function, i.e. a procedure that remembers the results that have already been computed [9].

Often in dynamic programming we also make use of backward induction where the first stage calculated is actually the last stage in the problem, and we move backward solving one state at a time [6]. An example would be a shortest path question. In such a problem we move through a

---

<sup>13</sup>In this case by programming we mean a “tabular method” and not the act of writing computer code [9].



network, starting at some node 0 and ending at a final destination node  $N$ . Each edge connecting intermediate nodes, which we will consider as states, may have different weights or lengths. We determine the shortest path by finding which edge has the minimum value between state  $n-1$  and the final state (node  $N$ ), and work our way backwards one state at a time until we reach the initial state (node 0). It is extremely important to note that this process only works if our network is acyclic, which is to say that we cannot begin from any node in any direction and be able to return to it [6].

We can now see how dynamic programming may be used to determine the solution to  $\pi(d, k)$ . Following the steps from Cormen, et. al., we first conclude that the answer to any specific  $\pi(x, y)$  may be constructed from  $\pi(x', y')$  where  $x' \leq x$  and  $y' \leq y$ . Second, we have already defined this structure recursively in subsection A.2. Third, we may easily compute the solution starting with the base case until we have solved and stored each  $\pi(d, k)$  for  $0 < d \leq d^*$  and  $0 \leq k \leq 6d^*$  in some matrix  $dice_{1:d^* \times 0:6d^*}$ . Fourth, we may determine any  $\pi(x, y)$  with  $0 < x \leq d^*$  and  $0 \leq y \leq 6d^*$  by looking up element  $(x, y)$  in the *dice* matrix. An algorithm (presented in a pseudocode format) to do just this is available in subsection A.5, while actual C++ code may be found in Appendix B.

#### A.4 Solving $P_{i,j}, 0 \leq i, j < 100$

With the solution to  $\pi(d, k)$  under our belt we are now ready to turn to Equation 3.1, reproduced below.


$$P_{i,j} = \max_{0 < d \leq d^*} \sum_{k=0}^{6d} \pi(d, k)(1 - P_{j,i+k}). \quad (\text{A.7})$$

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 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 game 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 [27]. 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 game states into independent

groups and work backwards. That is, we calculate the probability of winning for all states with  $i + j = 198$  then  $i + j = 197$ ,  $i + j = 196$ , ... ,  $i + j = 2$ ,  $i + j = 0$ . Without loss of generality, we may replace 0 with 1 in each of the games states. This now corresponds to figure A.1, 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  $i + j = s$ ,  $s \in \{2, 3, 4, \dots, 197, 198\}$ .

Figure A.1: Visual Partitioning of Game States by the Sum of Players' Scores



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 the next diagonal partition. Since Equation A.7 is also defined recursively, we might initially be tempted to say that the problem can be solved with dynamic programming. Unfortunately, this would be a false conclusion; the network representing this problem is not acyclic. 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  $(i, j)$ , a cycle. 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 Equation A.7 as the following:

$$P_{i,j} = \max_{0 < d \leq d^*} \left\{ \left[ \sum_{k=2}^{6d} \pi(d, k)(1 - P_{j,i+k}) \right] + \pi(d, 0)(1 - P_{j,i}) \right\}. \quad (\text{A.8})$$

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

$$P_{i,j} = \max_{0 < d \leq d^*} \{f(i, j, d) + \pi(d, 0) - \pi(d, 0)P_{j,i}\}. \quad (\text{A.9})$$

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 Equation A.10 below.

$$P_{i,i} = \max_{0 < d \leq d^*} \frac{f(i, i, d) + \pi(d, 0)}{1 + \pi(d, 0)}. \quad (\text{A.10})$$

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  $i \neq j$ :

$$P_{i,j} = \max_{0 < d \leq d^*} \{f(i, j, d) + \pi(d, 0) - \pi(d, 0)P_{j,i}\} \quad (\text{A.11})$$

$$P_{j,i} = \max_{0 < \delta \leq \delta^*} \{f(j, i, \delta) + \pi(\delta, 0) - \pi(\delta, 0)P_{i,j}\}. \quad (\text{A.12})$$

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

$$P_{i,j} = \frac{f(i, j, \bar{d}) + \pi(\bar{d}, 0) - \pi(\bar{d}, 0)[f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0)]}{1 - \pi(\bar{d}, 0)\pi(\bar{\delta}, 0)} \quad (\text{A.13})$$

$$P_{j,i} = \frac{f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0) - \pi(\bar{\delta}, 0)[f(i, j, \bar{d}) + \pi(\bar{d}, 0)]}{1 - \pi(\bar{\delta}, 0)\pi(\bar{d}, 0)}. \quad (\text{A.14})$$

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  $\delta$  dice, Player 1 will choose  $d$  such that he maximizes  $P_{i,j}$ , i.e.

$$P_{i,j} = \max_{0 < d \leq d^*} \frac{f(i, j, d) + \pi(d, 0) - \pi(d, 0)[f(j, i, \bar{\delta}) + \pi(\bar{\delta}, 0)]}{1 - \pi(d, 0)\pi(\bar{\delta}, 0)}. \quad (\text{A.15})$$

As we vary  $\bar{\delta}$  from 1 to  $\delta^*$ , 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

$$P_{j,i} = \max_{0 < \delta \leq \delta^*} \frac{f(j, i, \delta) + \pi(\delta, 0) - \pi(\delta, 0)[f(i, j, \bar{d}) + \pi(\bar{d}, 0)]}{1 - \pi(\delta, 0)\pi(\bar{d}, 0)} \quad (\text{A.16})$$

for each possible  $\bar{d}$ ,  $0 < \bar{d} \leq d^*$ . This will lead us to find all combinations of  $\bar{d} - \delta$  for which neither player would choose to alter his choice of dice quantity, a Nash equilibrium<sup>14</sup>. 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 A.1 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  $\delta = 7$ .

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

<sup>14</sup>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.

Since we let  $d_{max} = 25$ , for all games states in which  $i, j \leq 50$ ,  $d^*$  and  $\delta^*$  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 Jacobi 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<sup>15</sup> [8]. We continue generating these  $x^{(i)}$ ’s until a predetermined level of precision is achieved. This means that we repeatedly solve

$$x_i^{(k)} = \left[ - \sum_{j=1, j \neq i}^n (a_{i,j}/a_{i,i})x_j^{(k-1)} + (b_i/a_{i,i}) \right] \quad (1 \leq i \leq n) \quad (\text{A.17})$$

(assuming nonzero diagonal elements) until  $\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| < \epsilon$ , where  $\epsilon$  is our error tolerance [8].

A Slightly modified version of Jacobi iteration which typically allows for faster convergence is Gauss-Seidel iteration. Since the elements of  $\mathbf{x}$  are determined serially and not in parallel, we solve

$$x_i^{(k)} = \left[ - \sum_{j=1, j < i}^n (a_{i,j}/a_{i,i})x_j^{(k)} - \sum_{j=1, j > i}^n (a_{i,j}/a_{i,i})x_j^{(k-1)} + (b_i/a_{i,i}) \right] \quad (\text{A.18})$$

since newly calculated values may be used immediately [8].

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 [27]. Thus, such a system could not usually be solved iteratively using either of the two methods mentioned above. 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  $x_i = 1.0 \forall i$ . We therefore recommend its use since it saves a considerable amount of computation time. Algorithms for both the game-theoretical solution and the iterative solution may be found in subsection A.5 below; C++ code is available in Appendix B.

---

<sup>15</sup>Since this is merely a cursory glance at iterative solutions, we refer interested readers to [8] for appropriate convergence criteria.

## A.5 Solution Algorithms

The following algorithms are presented in a pseudocode that should be interpretable to anyone with basic programming knowledge. It is presented as a broad overview. That is to say that summations, maximums, and other computations may need to be further broken down into other loops; we avoid this for simplicity and ease of readability. However, an example of C++ code adhering to these algorithms (using the iterative approach) may be found in Appendix B. Comments are preceded by the % symbol, as in MATLAB code, and boxed off sections within the algorithms pertain to alternative choices of implementation. These algorithms, though correct, are not necessarily the most efficient.

### An Algorithm for solving $\pi(d, k)$

```
procedure Score( $d_{max}$ )  
integer  $d_{max}, d, k, r$   
real array  $dice_{1:d_{max} \times 0:6d_{max}}$  (initialized to 0.0)  
for  $k = 0$  to  $6$  do  
     $dice_{1,k} \leftarrow 1/6$   
end for  
 $d_{1,1} \leftarrow 0.0$   
for  $d = 2$  to  $d_{max}$  do  
     $dice_{d,0} \leftarrow (5 * dice_{d-1,0} + 1)/6$   
    for  $k = 2d$  to  $6d$  do  
         $dice_{d,k} \leftarrow \frac{1}{6} \sum_{r=2}^{\min(6,k-2)} dice_{d-1,k-r}$   
    end for  
end for  
return  $dice$   
end Score
```

**An Algorithm for solving  $P_{i,j}$ ,  $0 \leq i < 100$  and  $0 \leq j < 100$**

```

procedure  $Pwin(g, d_{max}, dice, \epsilon, m)$ 
integer  $g, d_{max}, m, c, i, j, k, d_x^*, d_y^*, \bar{d}, d'$ 
real  $\epsilon, x', y'$ 
integer array  $roll_{1:g-1 \times 1:g-1}$  (initialized to 0)
real array  $dice_{1:d_{max} \times 0:6d_{max}}$  (initialized to 0.0),  $opt_{1:g-1 \times 1:g-1}$  (initialized to 1.0),
     $matrix_{1:d_y^* \times 1:d_x^* \times 1:2}$  (initialized to 0.0)
for  $k = g - 1$  to 1 do
     $i \leftarrow k$ 
     $j \leftarrow g - 1$ 
    while  $j - i > 0$ 
         $d_x^* \leftarrow \min\{d_{max}, \lceil \frac{100-i}{2} \rceil\}$ 
         $d_y^* \leftarrow \min\{d_{max}, \lceil \frac{100-j}{2} \rceil\}$ 

```

*% Game Theoretical Solution*

**for**  $d' = 1$  **to**  $d_y^*$  **do**

**for**  $\bar{d} = 1$  **to**  $d_x^*$  **do**

$$matrix_{d', \bar{d}, 1} \leftarrow \frac{f(i, j, \bar{d}) + dice_{\bar{d}, 0} - dice_{\bar{d}, 0} [f(j, i, d') + dice_{d', 0}]}{1 - dice_{\bar{d}, 0} dice_{d', 0}}$$

$$matrix_{d', \bar{d}, 2} \leftarrow \frac{f(j, i, d') + dice_{d', 0} - dice_{d', 0} [f(i, j, \bar{d}) + dice_{\bar{d}, 0}]}{1 - dice_{d', 0} dice_{\bar{d}, 0}}$$

**end for**

**end for**

**call** **Equilibrium**( $matrix, opt, roll, i, j, d_x^*, d_y^*$ )

*%—or—*

```
% Iterative Solution
```

```
 $c, x', y' \leftarrow 0$ 
```

```
while  $\|opt_{i,j} - x', opt_{j,i} - y'\| > \epsilon$  and  $c < m$ 
```

```
     $x' \leftarrow opt_{i,j}$ 
```

```
     $y' \leftarrow opt_{j,i}$ 
```

```
     $opt_{i,j} \leftarrow \max_{0 < d \leq d_x^*} \left\{ \sum_{k=2}^{6d} dice_{d,k}(1 - opt_{j,i+k}) + dice_{d,0}(1 - opt_{j,i}) \right\}$ 
```

```
     $roll_{i,j} \leftarrow d$  (corresponding to  $opt_{i,j}$  in previous line)
```

```
     $opt_{j,i} \leftarrow \max_{0 < d \leq d_y^*} \left\{ \sum_{k=2}^{6d} dice_{d,k}(1 - opt_{i,j+k}) + dice_{d,0}(1 - opt_{i,j}) \right\}$ 
```

```
     $roll_{j,i} \leftarrow d$  (corresponding to  $opt_{j,i}$  in previous line)
```

```
     $c \leftarrow c + 1$ 
```

```
end while
```

```
     $i \leftarrow i + 1$ 
```

```
     $j \leftarrow j - 1$ 
```

```
end while
```

```
if
```

```
     $d_x^* \leftarrow \min\{d_{max}, \lceil \frac{100-i}{2} \rceil\}$ 
```

```
     $opt_{i,i} = \max_{0 < d \leq d_x^*} \left\{ \frac{[\sum_{k=2}^{6d} dice_{d,k}(1 - opt_{i,i+k})] + dice_{d,0}}{1 + dice_{d,0}} \right\}$ 
```

```
     $roll_{i,i} \leftarrow d$  (corresponding to  $opt_{i,i}$  in previous line)
```

```
end if
```

```
end for
```

```
for  $k = g - 2$  to 1 do
```

```
     $j \leftarrow k$ 
```

```
     $i \leftarrow 1$ 
```

```
    while  $j - i > 0$ 
```

```
        % Insert code from outermost while loop in previous for loop
```

```
    end while
```

```
end for
```



```

if
    % Insert code from previous if statement
end if
return roll, opt
end Pwin

procedure Equilibrium(matrix, opt, roll, i, j,  $d_x^*$ ,  $d_y^*$ )
integer i, j,  $d_x^*$ ,  $d_y^*$ ,  $\bar{d}$ ,  $d'$ , x, y
real  $x_{max}$ ,  $y_{max}$ 
integer array roll1:g-1×1:g-1 (initialized to 0), nash1:dy*×1:dx*×1:2 (initialized to 0.0)
real array opt1:g-1×1:g-1 (initialized to 1.0), matrix1:dy*×1:dx*×1:2 (initialized to 0.0)
for  $d' = 1$  to  $d_y^*$  do
     $x_{max} \leftarrow \max_{0 < \bar{d} \leq d_x^*} matrix_{d', \bar{d}, 1}$ 
     $x \leftarrow \bar{d}$  (corresponding to  $x_{max}$  in previous line)
     $nash_{d', x, 1} \leftarrow 1$ 
end for
for  $\bar{d} = 1$  to  $d_x^*$  do
     $y_{max} \leftarrow \max_{0 < d' \leq d_y^*} matrix_{d', \bar{d}, 2}$ 
     $y \leftarrow \bar{d}$  (corresponding to  $y_{max}$  in previous line)
     $nash_{d', x, 2} \leftarrow 1$ 
end for
for  $d' = 1$  to  $d_y^*$  do
    for  $\bar{d} = 1$  to  $d_x^*$  do
        if  $nash_{d', \bar{d}, 1} + nash_{d', \bar{d}, 2} = 2$ 
             $roll_{i, j} \leftarrow \bar{d}$ 
             $opt_{i, j} \leftarrow matrix_{d', \bar{d}, 1}$ 
             $roll_{j, i} \leftarrow d'$ 
             $opt_{j, i} \leftarrow matrix_{d', \bar{d}, 2}$ 
        return opt, roll
    end for
end for

```

```
    end if
  end for
end for
end Equilibrium
```

## B Hog Solution Code, C++

The following code, written in C++<sup>16</sup>, first determines  $\pi(d, k)$  for  $0 < d \leq 25$  (and hence,  $0 \leq k \leq 150$ ), and then iteratively calculates the optimal solution for the two-player game of hog with goal threshold of 100 points. The output is a MATLAB script file (.m) which will produce Figure 3.2.

```
#include<iostream>
#include<fstream>
#include<iomanip>
#include<vector>
#include<cmath>

using namespace std;

// GLOBAL CONSTANTS
const int DMAX = 25; // Maximum number of dice
const int GOAL = 100; // Point total to be reached
const int EQUATIONS = (GOAL-1)*(GOAL-1); // Number of states to be determined
const int SIDE = 6; // Number of sides on the dice
const int ITMAX = 100; // Maximum number of iterations
const double EPSILON = 0.000000001; // Error tolerance

// FUNCTION PROTOTYPES
double mag(double,double,double,double);
double known(const vector< vector<double> >&, const vector< vector<double> >&,int,int,int);
int maximum(const vector< vector<double> >&,int,int);

int main ()
{
    // VARIABLES
    int count = 0; // To store number of iterations
    int m = 0; // To store summation limit in calculation of \pi(d,k)
    int dstar_x = 0; // maximum number of dice for P_{i,j}
    int dstar_y = 0; // maximum number of dice for P_{j,i}
    double x_old = 0; // To store iteration t-1 for P_{i,j}
    double y_old = 0; // To store iteration t-1 for P_{j,i}
    int i = 0; // Index
    int j = 0; // Index
    ofstream out; // Output variable
    // \pi(d,k) matrix
    vector< vector<double> > dice(DMAX+1, vector<double>(DMAX*SIDE+1, 0.0));
    // Optimal Solution: P(Win)
    vector< vector<double> > optimal(GOAL, vector<double>(GOAL, 1.0));
    // Optimal Solution: Number of dice
    vector< vector<int> > optroll(GOAL, vector<int>(GOAL, 0));
    // P_{i,j,d} matrix
    vector< vector<double> > grid(EQUATIONS+1, vector<double>(DMAX+1, 0.0));
```

---

<sup>16</sup>Because C++ indexes its arrays and vectors starting from 0, we declare our array sizes to be one greater than necessary and ignore the first entry; this allows us to nicely match the array indices to the logic of the problem itself.

```

// DETERMINE \pi(d,k)
// d = 1; k \in {0, 2, 3, 4, 5, 6}
for(int k = 0; k <= SIDE; k++)
    dice[1][k] = 1./SIDE;
dice[1][1]=0.0;
for(int d = 2; d <= DMAX; d++)
{
// d > 1 and k = 0;
    dice[d][0] = ((SIDE-1.)*dice[d-1][0] + 1.)/(double)(SIDE);
// k < 2d or // k > 6d
    // already initialized to zero
// otherwise
    for(int k = 2*d; k <= SIDE*d; k++)
    {
        m = min(SIDE, k-2);
        for(int r = 2; r <= m; r++)
            dice[d][k]=dice[d][k] + dice[d-1][k-r]/(double)(SIDE);
    }
}

// FIND OPTIMAL SOLUTION
for(int k = GOAL-1; k > 0; k--) // i + j = 198 to i + j = 100
{
    i = k;
    for(j = 99; (j-i)>0; j--)
    {
        count = 0;
        x_old = 0;
        y_old = 0;
        dstar_x = min((int)ceil((GOAL-i)/2.), DMAX);
        dstar_y = min((int)ceil((GOAL-j)/2.), DMAX);
        while(mag(optimal[i][j], x_old, optimal[j][i], y_old)>EPSILON && count < ITMAX)
        {
            x_old = optimal[i][j];
            y_old = optimal[j][i];
            for(int d = 1; d <= dstar_x; d++)
                grid[(i-1)*(GOAL-1)+j][d] = known(dice,optimal, i, j, d) + dice[d][0]
                - dice[d][0]*optimal[j][i];
            optroll[i][j] = maximum(grid, (i-1)*(GOAL-1)+j, dstar_x);
            optimal[i][j] = grid[(i-1)*(GOAL-1)+j][optroll[i][j]];
            for(int d = 1; d <= dstar_y; d++)
                grid[(j-1)*99+i][d] = known(dice, optimal, j, i, d) + dice[d][0]
                - dice[d][0]*optimal[i][j];
            optroll[j][i] = maximum(grid, (j-1)*(GOAL-1)+i, dstar_y);
            optimal[j][i] = grid[(j-1)*(GOAL-1)+i][optroll[j][i]];
            count++;
        }
        i++;
    }
}
if (i == j)
{
    dstar_x = min((int)ceil((GOAL-i)/2.), DMAX);
    for(int d = 1; d <= dstar_x; d++)
        grid[(i-1)*(GOAL-1)+j][d] = ((known(dice, optimal, i, j, d)

```

```

        + dice[d][0]/(1 + dice[d][0]));
    optroll[i][j] = maximum(grid, (i-1)*(GOAL-1)+j, dstar_x);
    optimal[i][j] = grid[(i-1)*(GOAL-1)+j][optroll[i][j]];
}
}

for(int k = GOAL-2; k > 0; k--) // i + j = 99 to i + j = 2
{
    j = k;
    for(i = 1; (j-i)>0; i++)
    {
        count = 0;
        x_old = 0;
        y_old = 0;
        dstar_x = min((int)ceil((GOAL-i)/2.), DMAX);
        dstar_y = min((int)ceil((GOAL-j)/2.), DMAX);
        while(mag(optimal[i][j], x_old, optimal[j][i], y_old)>EPSILON && count < ITMAX)
        {
            x_old = optimal[i][j];
            y_old = optimal[j][i];
            for(int d = 1; d <= dstar_x; d++)
                grid[(i-1)*(GOAL-1)+j][d] = known(dice, optimal, i, j, d) + dice[d][0]
                - dice[d][0]*optimal[j][i];
            optroll[i][j] = maximum(grid, (i-1)*(GOAL-1)+j, dstar_x);
            optimal[i][j] = grid[(i-1)*(GOAL-1)+j][optroll[i][j]];
            for(int d = 1; d <= dstar_y; d++)
                grid[(j-1)*(GOAL-1)+i][d] = known(dice, optimal, j, i, d) + dice[d][0]
                - dice[d][0]*optimal[i][j];
            optroll[j][i] = maximum(grid, (j-1)*(GOAL-1)+i, dstar_y);
            optimal[j][i] = grid[(j-1)*(GOAL-1)+i][optroll[j][i]];
            count++;
        }
        j--;
    }
    if(i == j)
    {
        dstar_x = min((int)ceil((GOAL-i)/2.), DMAX);
        for(int d = 1; d <= dstar_x; d++)
            grid[(i-1)*(GOAL-1)+j][d] = ((known(dice, optimal, i, j, d)
            + dice[d][0]/(1 + dice[d][0]));
        optroll[i][j] = maximum(grid, (i-1)*(GOAL-1)+j, dstar_x);
        optimal[i][j] = grid[(i-1)*(GOAL-1)+j][optroll[i][j]];
    }
}

// PRINT OPTIMAL NUMBER OF DICE TO BE GRAPHED IN MATLAB
out.open("optimum_dice.m");
out << "x=[1:99]; \ny = x; \nz = [ \n";
for(int j = 1; j < GOAL; j++)
{
    for(int i = 1; i < GOAL; i++)
        out << optroll[i][j] << " ";
    out << '\n';
}

```

```

    out << "]" ; \nsurf(x,y,z);\n";
    out.close();

    return 0;
}

// FUNCTION DECLARATIONS
double mag(double x1, double x2, double y1, double y2)
// Returns the vector norm of the difference of two 2x1 vectors
{
    return pow(pow(x2 - x1, 2.0) + pow(y2-y1, 2.0), 0.5);
}

double known(const vector< vector<double> >& dice, const vector< vector<double> >& optimal,
int i, int j, int d)
// Calculates f(i,j,d)
{
    double num = 0.0;
    for(int k = 2; k <= SIDE*d; k++) // for each score total
        if(i+k < GOAL)
            num = num + dice[d][k]*(1-optimal[j][i+k]);
        else
            num = num + dice[d][k];
    return num;
}

int maximum(const vector< vector<double> >& matrix, int row, int dstar)
// Returns the largest value in the vector
{
    int max;
    max = 1;
    for(int n = 2; n <= dstar; n++)
        if(matrix[row][n] > matrix[row][max])
            max = n;
    return max;
}

```

## C Subject Recruitment Flier



You can sign up for one of the following sessions by sending an email to Shipra De at [des2@unlv.nevada.edu](mailto:des2@unlv.nevada.edu). Please indicate your session preference(s) in the email.

Session #	Date and Time	Location
1	W Oct 20 7-8:15pm	BEH-240
2	T Oct 26 2:30- 3:45pm	BEH-240
3	M Nov 1 4:00-5:15 pm	BEH-240
4	M Nov 1 7:00-8:15pm	BEH-240

## D Instructions



University of Nevada Las Vegas  
Honors College

### Instructions for Race Game

---

Welcome to the experiment on Race Games! This experiment is open to any students enrolled in MGT-301. The instructions for the experiment are simple. If you follow them carefully and make good decisions, you could earn up to \$20 by playing five rounds of the game. In addition, you will earn a \$5 "show-up" fee and possible extra credit (please confirm with your instructor).

#### Description of the Task

In this Race Game, you will play five rounds (games) against a randomly determined opponent. The identity of your opponent will be anonymous. Each game is identical. The object of the game is to be the first person to reach 100 points. You earn points by rolling a number of electronic dice on the computer screen. It is your choice at the beginning of each turn to determine how many dice you would like to roll (maximum 25). Note that this decision is independent of previous turns, i.e. if you choose to roll 4 dice during a given turn, you may choose to roll fewer dice, say 2, or a greater number of dice, perhaps 5, during your next turn. If the number 1 appears on any of the dice, you earn zero points. Otherwise the sum of the numbers appearing on all dice will be added to your game total. In either case it then becomes your opponent's turn to play. Your opponent plays by the same rules.<sup>17</sup>

In summary:<sup>18</sup>

- Each player determines the number of dice he or she would like to roll at the start of each turn.
- If a 1 appears on any of the dice rolled, the player earns zero points for that turn.
- If a 1 does not appear on any dice, the sum of all numbers rolled is added to the player's game total.
- The first player to accumulate 100 points wins the round.

---

<sup>17</sup>Instructions for Treatment 2 included the following paragraph appearing before the itemized summary: *The player who rolls first will be determined by chance at the beginning of the first game. In subsequent games, the player who begins play will be alternated. In each game the player who rolls first will start with 0 points, while the player's opponent, who is second to play, starts the game with 50 points.*

<sup>18</sup>Treatment 2 instructions had the following additional bullet appearing between bullets 3 and 4: *At the beginning of each game the first player to roll begins with 0 points, while the player's opponent begins with 50 points.*



The player who rolls first will be determined by chance at the beginning of the first game. In subsequent games, the player who begins play will be alternated.<sup>19</sup>

At the end of the fifth round, you will be asked to answer a brief survey. Note that payment for participation cannot be collected until the survey is complete.

Finally, at the start of play you will be supplied with a sheet of blank paper and a pen which may be used to record any information that you feel may help you play the game. Do not write your name or any other personal information on this paper as it will be collected with your exit survey and should contain no information that could identify you.

### **Receiving Extra Credit and Payment**

Your earnings will be paid to you in cash at the end of the experiment upon completing the exit survey. Remember that you will earn extra credit towards your MGT-301 grade and \$5 for participating in the study, regardless of how many rounds you play. In addition, you will earn \$4 for each game you win.

Notification of extra credit will be forwarded to your MGT-301 participating professor on a weekly basis. If you have any questions regarding this experiment, please contact Professor Seale at 702-895-3365 or [dseale@unlv.nevada.edu](mailto:dseale@unlv.nevada.edu), or email Shipra De at [des2@unlv.nevada.edu](mailto:des2@unlv.nevada.edu).

Thank you for your participation and good luck!

---

<sup>19</sup>This paragraph was omitted from the Treatment 2 instructions. Also, the second sentence was a misprint in the Treatment 1 instructions; in actuality the winner of each game had to allow his opponent to go first in the following game.

# E Survey

<b>RACE GAMES EXIT SURVEY</b>					
<b>Player Number:</b> _____		<b>Age:</b> _____		<b>Gender:</b> <input type="checkbox"/> Female <input type="checkbox"/> Male	
<b>Ethnicity (please check <input checked="" type="checkbox"/>):</b>					
<input type="checkbox"/> Hispanic or Latino <input type="checkbox"/> Not Hispanic or Latino <input type="checkbox"/> Unknown					
<b>Race (please check <input checked="" type="checkbox"/>):</b>					
<input type="checkbox"/> American Indian		<input type="checkbox"/> East Asian		<input type="checkbox"/> South Asian	
<input type="checkbox"/> White		<input type="checkbox"/> Hawaiian or Pacific Islander		<input type="checkbox"/> Black or African American	
<input type="checkbox"/> More than One Race		<input type="checkbox"/> Other		<input type="checkbox"/> Unknown	
<p><i><b>Directions:</b></i> In the following section, there are phrases describing people's behaviors. Please indicate how accurately each statement describes <b>you</b>. Describe yourself as you generally are now, not as you wish to be in the future. Describe yourself as you honestly see yourself, in relation to other people you know of the same sex as you are, and roughly your same age. So that you can describe yourself in an honest manner, your responses will be kept in absolute confidence. Please read each statement carefully, and then check the appropriate box <input checked="" type="checkbox"/>.</p>					
<b>Enjoy being reckless.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Would never make a high risk investment.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Stick to the rules.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Seek danger.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Am willing to try anything once.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Know how to get around the rules.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Seek adventure.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Avoid dangerous situations.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Take risks.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	
<b>Would never go hang-gliding or bungee-jumping.</b>					
<input type="checkbox"/> Very Inaccurate	<input type="checkbox"/> Moderately Inaccurate	<input type="checkbox"/> Neither Inaccurate nor Accurate	<input type="checkbox"/> Moderately Accurate	<input type="checkbox"/> Very Accurate	

**Directions:** Please answer the following questions to the best of your knowledge and ability. All dice referred to in the questions should be considered to be **fair, six-sided dice** (Please check ).

**How many classes (post high school) have you taken in which you have learned some statistics and/or probability theory?**

- None       1       2       3       4 or more

**What is the probability of *not* rolling a 1 using one die?**

- 1/2       2/3       3/4       4/5       5/6

**What is the probability of *not* rolling any 1's using two dice?**

- 1/2       7/12       2/3       25/36       5/6

**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'

**Do you gamble?**

- Never       Sometimes       Often

**In general, do you feel lucky?**

- No       Yes

**Directions:** Please answer the following questions in a legible manner. Take a moment to consider each question before writing down your response. Answer all parts of the question.

Games Played: \_\_\_\_\_

If not 5, please state your reason for quitting: \_\_\_\_\_

\_\_\_\_\_

Games Won: \_\_\_\_\_

What was your original strategy? \_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

Did your strategy change from game to game? If yes, how so? \_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

Did your strategy ever change based upon the behavior of your opponent? \_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

Do you feel your performance improved from game to game? If not, why do think this was the case? \_\_\_\_\_

---

---

After playing the game once did you want to play again? Why or Why not? \_\_\_\_\_

---

---

Did you enjoy playing the game? Why or why not? \_\_\_\_\_

---

---

Thank you for your participation!  
Don't forget to pick up your cash award as you exit the computer lab.

FIN