WHAT ARE THE ODDS?:

A UTILISATION OF SETTLERS OF CATAN TO PROVIDE A FRAMEWORK FOR MONTE CARLO BASED GAME THEORETIC ORIENTED DECISION-MAKING.

A Thesis

Presented to:

The Faculty of the School of Engineering and Applied Sciences

University of Virginia

In Partial Fulfilment

Of the Requirements for the Degree

Master of Science (Systems and Information Engineering)

By

A. P. D. G. Everett

B.A.; University of Washington, 2009

B.Sc.; University of Washington, 2009

15 July 2016

Approval Sheet

The thesis is submitted in partial fulfilment for the degree:

Master of Science

A. P. D. G. Everett

The thesis has been read and approved by the examining committee:

K. Preston White, Jr. (Adviser)

Michael C. Smith (Chair)

James W. Lark, III

Mark R. Floryan

Accepted for the School of Engineering and Applied Science:

Craig H. Benson, Dean, School of Engineering and Applied Science

August 2016

Proverbs 24:6:

For it is by wise guidance that you wage your war, and the victory is due to a wealth of counsellors.

Aeneid, Book 1:

Arma virumque cano, Troiae qui primus ab oris (My song is of war, of men who first came from Troy)

Aeneid, Book 6:

Bella! Horrida Bella! (Wars! Horrible Wars!)

The Hon Herbert Asquith, The Fallen Subaltern:

As goes the Sun-god in his chariot glorious, When all his golden banners are unfurled, So goes the soldier, fallen but victorious, And leaves behind a twilight in the world.

1) Executive Summary

The importance of making decisions well is paramount in all fora, especially when making long range plans (what would generally be called strategy). Random events in life, as well as unpredictable decisions made by others, can impair or eliminate years of time and work. Accurate planning is needed in order to respond to changing conditions brought on by those events and decisions.

Settlers of Catan is a Eurogame popular across many spheres, from tech executives to NFL players to friends playing at home. Settlers previously has been the object of study and modelling in the artificial intelligence realm. In this research, Settlers of Catan is analysed for its suitability for providing a framework to study and modelling of strategic decision-making. We demonstrate that Settlers has strong value modelling the strategic/operational level of decision-making. Play involves random events determined by the roll of a pair of dice and decisions made by other players regarding resource trades. Defeat and victory are strictly defined and judging what works and what doesn't is a direct exercise with no ambiguity.

The complexity of Settlers of Catan rivals or even exceeds that of the board game Go. As explained in this study, Settlers of Catan is an *n*-person, finite, extensive-form game, which means that it possesses one or more subgame perfect equilibria. Regular backwards induction only has value for finite games of perfect information. Analysis of subgames allows extension into fora that are of imperfect information, which is to say most spheres of human activity.

In order to assess issues of strategy and trade for this game, roughly 180 people were found to complete a general survey that covered these issues, as well as measured demographic and personality information thought relevant to the game. Roughly 40 games of two types were also conducted, upon which observations of strategy and trade were made.

In addition to devising as a protocol to record game play, this study develops and applies a means to map personality characteristics to a quantitative strategic thinking quotient. This is method generalizable and what has been demonstrated in through use of Settlers of Catan, can be used in many other one decision-making contexts. It is a first foray into this sphere using this approach, so the results are tentative, but do demonstrate promise for future study and application.

2) Introduction

Why study games?

To quote the computer scientist Bruce Abramson describing the foundations of the study of games: "*The mathematical study of games predates digital computers by several decades. With the coming of the computer age, the marriage of games, as decision making models, to computers, as decision making machines, was swift and natural. In 1950, with computer science in its infancy and the term "artificial intelligence" as yet unborn, Claude Shannon's Programming a Computer for Playing Chess begat the computer game. In this classic work, Shannon justified chess programming as a valid scientific pursuit by claiming that aside from being an interesting problem in its own right, chess bears a close resemblance to a wide variety of more significant problems, including translation, logical deduction, symbolic computation, military decision making, and musical composition; in any of these fields, skilful performance requires thought, and satisfactory solutions, although generally attainable, are rarely trivial.*"[1]

Games, by their nature, are constrained things with their exacting rules and minimal total universe. In games, defeat and victory are strictly defined, and judging what works and what doesn't is a direct exercise with no ambiguity. So, following in the vein of Claude Shannon and many others, I am using a game to help elucidate issues outside of the sheer confines of the game directly studied, because games are a good test of artificial intelligence.

Why study strategy?

Strategy, the word and concept, as important as it is, has no single definition that is universally agreed upon. Edward Luttwak in his book *Strategy* has an entire appendix dedicated to listing various definitions of the word. In a game theoretic context, a strategy is the series of all actions that fully define the behaviour of a player. In a military context, strategy is the highest level in a hierarchy of several interrelated levels of activity in warfare.[2] According to the US Marine Corps, "At its most basic, strategy is a matter of figuring out what we need to achieve, determining the best way to use the resources at our disposal to achieve it, and then executing the plan.".[3] My own definition of strategy is an overarching plan to attain one or more goals under conditions of uncertainty.[4]

As it says in Marine Corps Doctrinal Publication (MCDP) 1, *Warfighting*, "All actions in war take place in an atmosphere of uncertainty, or the "fog of war."…While we try to reduce these

---

[1] Abramson, The Expected Outcome Model of Two-Player Games,
[2] MCDP 1 (*Warfighting*), Section 1
[3] MCDP 1-1 (*Strategy*), Section 1
[4] Given my background in meteorology/atmospheric sciences, I originally used the word "synoptic" (as it is meant in a meteorological context) instead of "overarching" in my original version of this formulation, which I changed for purposes of clarity for the reader.

unknowns by gathering information, we must realize that we cannot eliminate them—or even come close…War is intrinsically unpredictable. At best, we can hope to determine possibilities and probabilities…By judging probability, we make an estimate of our enemy's designs and act accordingly."[5] This text indicates both the high value attached to understanding and planning what your opponent does and given the probabilistic nature to which it refers, the need to use a Monte Carlo based method to effectively model decisions in this (or really any) context.

Strategic thinking and planning in the business world is derived from military strategy. This is why, for instance, classic military texts such as Sun-Tzu's *Art of War*, *The Book of the Five Rings* by Miyamoto Musashi, Niccolo Machiavelli's *Art of War* and *The Prince*, and many others are regularly read by those in business and assigned as reading by business school curricula.[6] This is also why there are more than a few papers in business-oriented journals such as the *International Journal of Business and Management* and several others as well as military journals that discuss the relationship of military strategy to business strategy, and the applications of the former to the latter.[7] The relationship is so pronounced that US Naval War College Professor of Strategy James Holmes wrote in *The National Interest* earlier this year (April 2016) an article entitled "*Business Strategy Isn't Military Strategy"*, in which he states:

"In short, the logic of business competition maps vaguely, at best, to the realm of violent interaction between antagonists. It's possible for military officialdom to grow too accustomed to operating by peacetime logic. Commanders and their political masters can concentrate so much intellectual energy and so many resources on efficiency, industrial processes and other linear undertakings that they lose sight of the battlefield—and forget that topsy-turvy rules reign there. War—a grim, remorseless teacher—may furnish a reminder."

However this does mean that something that is developed to improve strategic decision-making has a wide variety of applications, and whilst the mapping isn't one-to-one, there is a correlation

[5] MCDPs are higher order publications that contain the fundamental and enduring beliefs of the USMC with respect to armed conflict (MCDPs 1, 1-1, 1-2, and 1-3 are foundational publications) and serve as the guiding doctrine for the conduct of major warfighting activities (MCDPs 2, 3, 4, 5, and 6).

[6] Yes, The Prince is a book of political philosophy, but as Clausewitz noted in *On War*, war is a continuation of politics with other means (the common version of the aphorism is a mistranslation of the original German "Der Krieg ist eine bloße Fortsetzung der Politik mit anderen Mitteln", as James Holmes notes in The Diplomat in a 2014 article "Everything You Know about Clausewitz is Wrong" and even I, with my limited German proficiency, can confirm. Furthermore, as Christopher Lynch notes in the introduction of his translation of Machiavelli's *Art of War*, his three major works, *The Art of War* (which was the only major writing published in his lifetime) along with the more famous *Discourses on Livy* and *The Prince* are all interrelated, and it is impossible to understand Machiavelli's views on politics without also understanding his views on warfare and its conduct.

[7] One of the better recent surveys of strategic theory and practice applied to business was published last year (2015) by the Turkish Army War College, which covered a lot of the recent literature on the topic: http://www.harpak.edu.tr/Bilimsel_Faaliyetler/Kara_Harp_Akademisi/ICMSS_2015/2.2.pdf

between the two. So yes, they are related fora, but they are separate and it should be said that even in this context, too much of a good thing isn't.

A Brief History of Games and their application to Strategic thought

Games have existed for thousands of years. It cannot be denied that games have a recreational element, however even many years ago, they have been seen for their potential to be more than diversion for the idle. Games have existed since at least 3000 BCE and have been found in all cultures. Mancala, Senet and the Royal Game of Ur all date to this period. Go has been extant since 2000 BCE (more on Go later since it is a major strategy game that is relevant to this study). Dice have existed since 700 BCE. Chess has existed since roughly 700 CE, etc.[8] Let it be said that humanity has a long relationship with games, which ironically enough, that fact doesn't make games as a whole any easier to define.

There are many classes of strategy-oriented games, and the most classic are what are known as abstract strategy games, which are undeniably the oldest and most familiar to the widest variety of people. An abstract strategy game is only loosely tied to a thematic concept, if it is tied to a thematic concept at all. These rules as they exist in abstract strategy games do not attempt to simulate reality, but rather serve the internal logic of the game. J. Mark Thompson wrote in his article "Defining the Abstract":[9]

"There is an intimate relationship between such games and puzzles: every board position presents the player with the puzzle, What is the best move?, which in theory could be solved by logic alone. A good abstract game can therefore be thought of as a "family" of potentially interesting logic puzzles, and the play consists of each player posing such a puzzle to the other. Good players are the ones who find the most difficult puzzles to present to their opponents."

Chess, Go, Crosses and Noughts, and Draughts all fit into the category of abstract games.

Simulations are the reverse. This type of game is an attempt to imitate the decisions and processes inherent to some real-world situation. Abstract games cannot be completely divorced from simulations, thus games can be thought of as existing on a continuum of almost pure abstraction to almost pure simulation.[10]

---

[8] The time line of games in history was found at MacGregor Historical Games, http://www.historicgames.com/gamestimeline.html
[9] Thompson, J. Mark. (2000, July) Defining the Abstract. *The Games Journal.*
[10] Section 4 of Sebastian Martin Möring's PhD Dissertation, *Games and Metaphor – A critical analysis of the metaphor discourse in game studies*

Settlers of Catan (the primary object of this study, hereinafter generally shortened to SoC or Settlers) is a Eurogame, so fits roughly equidistant from both ends of the continuum, more abstract than simulations and less abstract than Chess and other abstract strategy games. Settlers paved the way for the Eurogame genre in the United States and outside Europe. It wasn't the first Eurogame to find an audience outside Germany, but it became much more popular than any of its predecessors. It quickly sold millions of copies in Germany, and in the process brought money and attention to the genre as a whole.[11,12]

Models of Conflict

Modelling conflicts, and elements of the same, is a practice that goes back over two centuries. The professionalisation of the armed forces, something that is described in Morris Janowitz's *The Professional Soldier*, and Samuel P. Huntington's *The Soldier and the State*, both seminal works in the field of civil-military relations, was required by the increase in complexity of planning, training, and equipping armies and navies. Navies became professional endeavours before armies (training and manning personnel for naval vessels is an activity that takes years, as is building the ships themselves, which is why the US Constitution in Article 1, Section 8, amongst the enumerated powers of the federal government states: "To raise and support Armies, but no Appropriation of Money to that Use shall be for a longer Term than two Years; To provide and maintain a Navy"), and in armies, technical branches such as logistics, engineering, and artillery became professional endeavours before the traditional combat arms of infantry and cavalry.[13] Warfare is not purely formulaic; no set of equations can fully describe the particulars of a military engagement, however that hasn't prevented some extremely bright people from attempting to do so.[14] Lewis Fry Richardson, an applied mathematician and meteorologist who was a pacifist for religious reasons made one of the first serious attempts to quantitatively study war in the early to mid-20th century, and given the nature of this study, a discussion of him and his work will occur later in this study.

It should be noted that war games are a class of simulations with specific application to the military,

[11] Harford, Tim (17 July 2010). "Why we still love board games". *FT Magazine*.
[12] Curry, Andrew (23 March 2009), "Monopoly Killer: Perfect German Board Game Redefines Genre", *Wired Magazine*
[13] It is for this reason that RMA Woolwich was founded in 1741, with the intent to produce: "good officers of Artillery and perfect Engineers", whereas Sandhurst, which was originally founded to train officers for the infantry and cavalry, whose leaders traditionally hailed from the aristocracy, wasn't founded until 1801. The US Military Academy was originally subordinate to the Chief of Engineers of the US Army, and wasn't removed from his purview until after the end of the Civil War for a similar reason (although it trained officers for the infantry and cavalry as well as engineers and artillery).
[14] This was arguably the biggest failing of Robert McNamara's approach to the Vietnam War, in that he and his subordinates tried to treat victory as a formula such that if the US supplied sufficient inputs, they'd get the desired output. Amongst the many problems in this is the fact that McNamara *et al* chose the wrong metrics (such as body count), which could be gamed.

simulating battles, campaigns, and even whole wars. Tactical simulators used by military and police forces for what is termed "close quarters battle" are in the same sphere as war games.

Why Settlers of Catan and What is its value applied to strategic games

SoC is a non-deterministic game that can be played by 3 to 4 people in the standard version (5 to 6 using the expansion) that given the fact of imperfect information, and game structure (to be addressed more later in section 4) has a large strategy element and multiple paths to victory, especially given the fact that game rules allow for random board set-up, so conditions are not identical from one game to the next. Furthermore, strategy games are a direct link between classic games (those that are two-person and of perfect information) and of video games (which all use AI engines), so there is a definite interest to AI researchers. In the words of a friend of mine who is an Intelligence Officer in the US Marine Corps:[15]

"*Settlers is the best model for the operational level of war in that the player who best decides how they are going to win and sticks with it wins. The moment when you realise that you will never get longest road or largest army and decide to hybridise your strategy is probably the closest point you can get to being that command who has to decide that the adversary has shifted strategy and the current plan must be abandoned.*"

Outside of a military context, it has other, relevant relationships to aspects of life. Settlers is a game that relies on economic concepts such as supply and demand, management of resources, and trade. It is possible to play Settlers without trade, but that assumes that a player is able to acquire all the resources they need by their own production. Often times, this is not the case, so players must trade with each other to get what they need. As was noted in a Wall Street Journal article discussing Settlers and its impact on the high technology crowd:[16]

"*LinkedIn's Mr. Hoffman, who estimates he has inducted nearly 40 Silicon Valley executives into the game, says tech entrepreneurs are drawn to Settlers because it "most closely approximates entrepreneurial strategy." The title pushes players to collaborate and swap resources to get points, while the random rolls of the dice force people to constantly revamp their strategies for winning. That's much like running a start-up, Mr. Hoffman says.*"

[15] (Capt.) Brian Strom, and specifically, he's a MAGTF Intelligence Officer with a Ground Intelligence Background. Prior to commissioning in the US Marine Corps, he took a degree in physics from Princeton and served in the US Navy. His professional expertise as well as his experience with Settlers has helped improved this study immensely, and his assistance, both providing input regarding applications of strategic thought as well as his participation in taking my Settlers Strategy Survey is greatly appreciated.
[16] Phillip Tam; *An Old-School Board Game Goes Viral Among Silicon Valley's Techie Crowd*, 17 Dec 2009

In short, Settlers of Catan is valuable because it models real-life interpersonal dynamics. And because of the wide range of people who play this game, from tech executives to NFL players to nerds playing with their friends at home, there have been a fair amount of resources expended towards studying Settlers and its applications, far more than have been used in studying the validity of certain simulations that are primarily or solely meant for military or law enforcement use.[17]

Research Motivation and Questions I Sought to Address

Two things primarily prompted my interest to perform this study, one mostly personal and the other from professional intellectual curiosity.

First and most personal, I have played Settlers of Catan since the 2000-2001 timeframe, when I was introduced to it by German friends. I found the game intellectually engaging and fascinating, the issues of trade and competition I found as stimulating in many respects as poker, and with strategic elements I found even more engaging than chess. So much so, that in the roughly 15 years since, I've played roughly 1300 games of Settlers. It is to be noted that Germany has a long-standing tradition of complex board games, and are famously obsessed with quality board games, so much so, in Germany, there is an award that has been around since 1978 for the best game to premier any given year (Spiel des Jahres, which Settlers of Catan won in 1995).

Second and more professionally, given my personal and professional interest in probability and applications of it, such as decision modelling, I was introduced to Monte Carlo Tree Search. As I state elsewhere, Monte Carlo based methods date back to the 1940s, and their application to games dates back to the 1980s (Bruce Abramson explored the idea in his 1987 PhD dissertation). It was first employed to Go in the early 1990s. Monte Carlo Tree Search was formalised in 2006 by Rémi Coulom. Given this, I became interested in its potential applications towards non-deterministic games. What I found was that this system had been used in modelling skat, poker, Magic: The Gathering, and Settlers of Catan, all in the last decade. Settlers of Catan is a game that is even more complex than Go, given both the aleatory elements brought on by both dice and the randomness of the game board set-ups, as well as the behavioural component brought on by game elements, trade, strategic concerns, etc. These traits make it ideal I think for attempting to model strategic decision making by humans.

[17] Kevin Clark wrote last year in the Wall Street Journal, Green Bay Packers players have a tradition of regularly playing Settlers after practices (*Green Bay's Board-Game Obsession*, 15 January 2015)

This prompted me to contemplate the following questions:

i) Why use Monte Carlo Tree Search (MCTS) in this context?
An MCTS based model allows for sampling of the game space, whereas other methods would require an exhaustive search that will use far too much computational power and effort.

ii) How can we use domain knowledge to improve MCTS?
MC methods at their most basic are random in nature, using knowledge, which can be refined to a feasible (within game rules) set.

iii) How can we arrive at a balance between exploration and exploitation of the game space using MCTS?
At each decision point, all moves should be investigated, but weighted to the most value-additive.

Foundational Issues

In order to create a viable MCTS model that appropriately accounts for actual human behaviour, first human behaviour had to in fact be observed and recorded systematically such that it would be useful. As Barry Silverman notes with respect to modelling human behaviour in strategic games and training models:

"A common challenge running throughout these applications is to increase the realism of the synthetic agents' behaviour and cognition. This is not an idle fancy, but a serious objective that directly affects the bottom line of commercial concerns, mission achievement in non-commercial organizations, and the safety and health of individuals who need to transfer skill sets from virtual to real worlds."[18]

As he also notes, there are the following challenges with respect to accurately modelling human behaviour:[19]

1) Developers possessing insufficient behavioural knowledge.
2) Artificial life has focused on low-level cognitive functioning.[20]
3) AI focuses on high-level cognitive functioning.
4) Behavioural and cognitive researchers tend to ignore implementation.

[18] Chapter 9, *The Science and Simulation of Human Performance, Volume 5* (2004)
[19] Ibid.
[20] Artificial life here is defined as virtual life that appears to act in a real way, however anything they do is reactive, and does not involve any deliberation or cognitive processing.

5) A dearth of interchange standards between game makers and the AI community.
6) A need to validate "useful" models of human behaviour.

This work was developed in an attempt to address in part some of these enumerated issues, specifically 1, 3, 4, and 6, with the intent to allow future researchers to take the data and tools I've generated to help further the realm of human behaviour modelling and its utilisation in strategic games and training models.

As a part of this study, in order to provide a check against the initial heuristics I developed to describe trading and general strategic practice, I crafted a survey for Settlers of Catan (in which trading practices and strategic forms are evaluated) as well as observed a group of Settlers games and had players fill out end-of-game surveys in order to ascertain strategic thought in that context. The original intent of the surveys conducted and observation of Settlers games was to quickly provide a check against those initial heuristics I developed. However, because of the numerous issues involved in these parts, including lack of documentation of Settlers games, a standard means to record them, lack of strategic literature, etc, this study changed focus to cover those issues, and lay the foundation for further work.

This study is the first work that uses Settlers of Catan as a framework for modelling that wasn't developed under a computer science faculty, so there are differing foci than what might be encountered in other studies. A greater focus is given to the strategic elements behind the game, human factors surrounding gameplay, especially given dealing with the random game elements, complex issues such as trade, etc. as well as the strategic behavioural issues that are a part of Settlers of Catan. I had started off this study thinking of strategy algorithms for Settlers play and it occurred to me in media res that before such could be implemented that people needed to understand more about human play and that improving this understanding would in itself be a worthwhile contribution. As such, this study attempts to bridge the gap between the computational and the humane.

3) What is MCTS?

Monte Carlo Tree Search (MCTS) is a method for making optimal decisions in artificial intelligence (AI) problems, typically move planning in combinatorial games. It combines the generality of random simulation with the precision of tree search.

Research interest in MCTS has risen sharply due to its spectacular success with computer Go (most recently with Alpha-Go and its beating two of the best current masters of the game in the last year, including the world champion in March 2016) and potential application to a number of other difficult problems.[21] Its application extends beyond games, and MCTS can theoretically be applied to any domain that can be described in terms of {*state*, *action*} pairs and simulation used to forecast outcomes.

The basic MCTS algorithm is simple: a search tree is built, node by node, according to the outcomes of simulated playouts. The process can be broken down into the following steps.

- **Selection**

  This is about choosing the "best" next action, this phase will decide how to build the next level of the tree. Starting at a root node *R*, recursively select optimal child nodes until a leaf node *L* is reached.[22]

- **Expansion**

  If *L* is a not a terminal node (i.e. it does not end the game) then create one or more child nodes C and select one. Simply, if the MCTS algorithm finds a new state (all states are nodes in this context), it adds it to the tree's current states.

- **Simulation**

  Run a simulated playout from *C* until a result is achieved. Here is where the necessity of good domain knowledge is vital. Since probabilities are worked with, a high quality heuristic will give a better probability for the nodes that would lead to the solution, and low probability for the nodes that looks far from the goal.

---

[21] The paper that the AlphaGo team wrote was published in Nature, volume 509 (28 January 2016) covered AlphaGo beating Fan Hui, the current European champion. There were two articles about this event that well covered the more social impacts of this discovery as well as describing it, in The Atlantic (*How Google's AlphaGo Beat a Go World Champion*, 28 March 2016, Christopher Moyer), and Wired (What the AI Behind AlphaGo can teach us about being human, 19 May 2016, Cade Metz), both of which take into account the matches against the world champion, Lee Sedol, and the exposition after AlphaGo had beaten him. It should be noted that AlphaGo's algorithm specifically uses MCTS to find knowledge previously gained from an artificial neural network by extensive training from both computer and human play.

[22] A root node is a node distinguished from the rest of the tree nodes, normally it is either shown as to the left or on top.
A leaf node is a form of child node, it is one in which has no children itself.

- **Backpropagation**

  Update the current move sequence with the simulation result.

This process can be seen as an analogue to what Professors Dixit and Nalebuff refer to in their primer on game theory as the first rule of strategic behaviour: "Look forward and reason back".[23] The process they describe can be broken down into three steps. First, looking at the ultimate decision and assuming that if it comes to that point, the deciding player will choose their optimal outcome (however that may be defined). Second, advancing one step back to the penultimate decision, and assume the next player would choose their best outcome, treating the following decision as fixed (given that we have already reasoned what that player will pick if it should come to that point.). Third, continuing this process back until all decisions have been fixed.

The primary application of MCTS has been towards games. Those games include classic board games such as Go and Crosses and Noughts (Crosses and Noughts is known in the US as tic-tac-toe), real-time video games such as Ms Pac-Man and non-deterministic games such as skat, poker, and Settlers of Catan.[24] It has also found application in solving transportation problems.[25]

As Bruno Bouzy proved, the evaluation of moves in Monte Carlo tree search converges to minimax.[26] He notes, that the basic version of MCTS, whilst it does converge, does so very slowly. Despite the fact of the slow convergence rate, it has many advantages that merit its use over alpha-beta pruning and other algorithms that minimise the search space. Specifically, MCTS does not need an explicit evaluation function (which are static in nature, given they're imposed prior to modelling and do not change during the model run).[27] Simply by implementing the mechanics of the game is all that is needed to explore the search space (that is to say the generation of allowed moves in a given position and the end of game conditions). Given that trait, MCTS can be employed in games without a developed theory of the game or in general game playing.

[23] They wrote a fantastic, non-technical primer on game theory and its applications, *Thinking Strategically*. This rule is on pg 34.
[24] A non-deterministic game is defined as a game with one or more random elements.
[25] Trunda and Bartak, "Using Monte Carlo Tree Search to Solve Planning Problems in Transportation Domains", MICAI 2013
[26] Bouzy, "Old-fashioned Computer Go vs Monte-Carlo Go", 2008.
[27] Evaluation functions were developed by Claude Shannon in his study of computers playing chess in 1950.

4) The game: Settlers of Catan

What is Settlers of Catan?

Settlers of Catan (SoC), created by Klaus Tauber in 1995, is a multiplayer (3-4 players in the standard game, 5-6 using the expansion), non-deterministic, sequential turn game. The goal of SoC is to settle an island, named Catan. In the standard form of the game, three or four players can play, each starts with two settlements and two road segments, and the goal is to acquire 10 Victory points (the two settlements are worth two).

Rules and Game Board

The standard form of the game board is formed from 19 hexagonal tiles in a 3-4-5-4-3 configuration, consisting of one desert hex (which produces no resource), and 18 hexes of five types that produce resources, that can be described as a set *Re* where:

Re = {brick, brick, brick, ore, ore, ore, sheep, sheep, sheep, sheep, wheat, wheat, wheat, wheat, wood, wood, wood, wood}

randomly distributed within the game space.

Each of the tiles in set *Re* is paired with an enumerated disc that corresponds to the probability distribution of the sum of two U~(1,6) discrete random variables (gameplay requires use of two 6-sided die), which I will refer to hereafter as the "dice distribution", which forms the set of numbers *N*, where: *N* = {2,3,3,4,4,5,5,6,6,8,8,9,9,10,10,11,11,12}.[28]

The 7 is not part of the set *N*, as rolling a 7 (probabilistically the most common outcome from the dice distribution) activates the robber (which is an essential element of gameplay) that initially resides in the desert hex.

[28] Thanks to my thesis adviser, Preston White, I tend to refer to the probability distribution of two U~(1,6) discrete random variables as "the dice distribution", and whilst formal literature doesn't express this under this specific terminology (even though it's a common enough concept), I believe that calling this distribution by that appellation enhances clarity in discussions on this topic.

As a reminder to the reader, the dice distribution is depicted in the following figure:

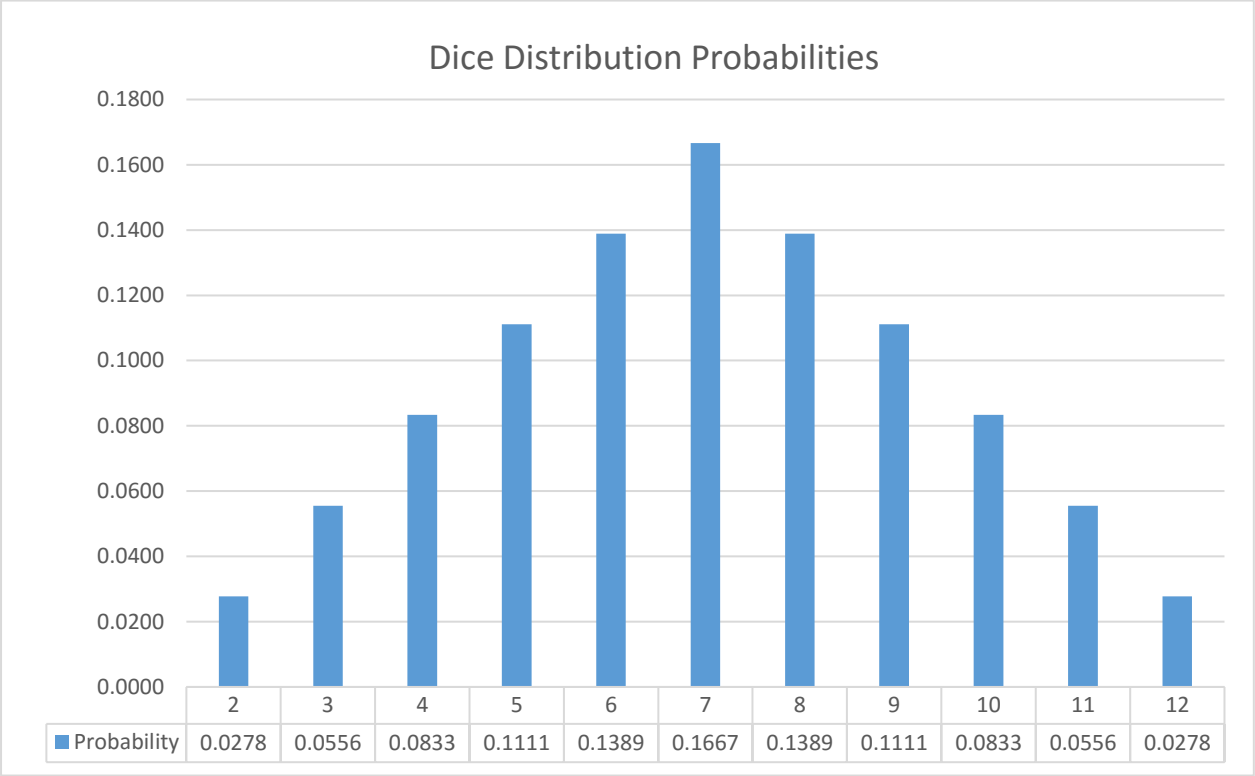

Figure 1 - Dice Distribution Probabilities

As you can see in Figure 1, 1/6th of the total probability mass is in the range 2 – 4, by reflection, this is also true of 10 – 12, which means that 3/6th of the probability mass is in the range 5 – 9 (7 excepted of course, for which 1/6th of the probability mass lay there). It should be noted that because of this distribution of *N*, the central tendency of the dice distribution is under-represented by the discs, as can be seen in Figure 2.[29]

[29] In the actual probabilities portion of the table, given that 7 discs do not exist (since 7s activate the robber), probabilities were re-scaled out of 30, since it is with that population that the discs are created.

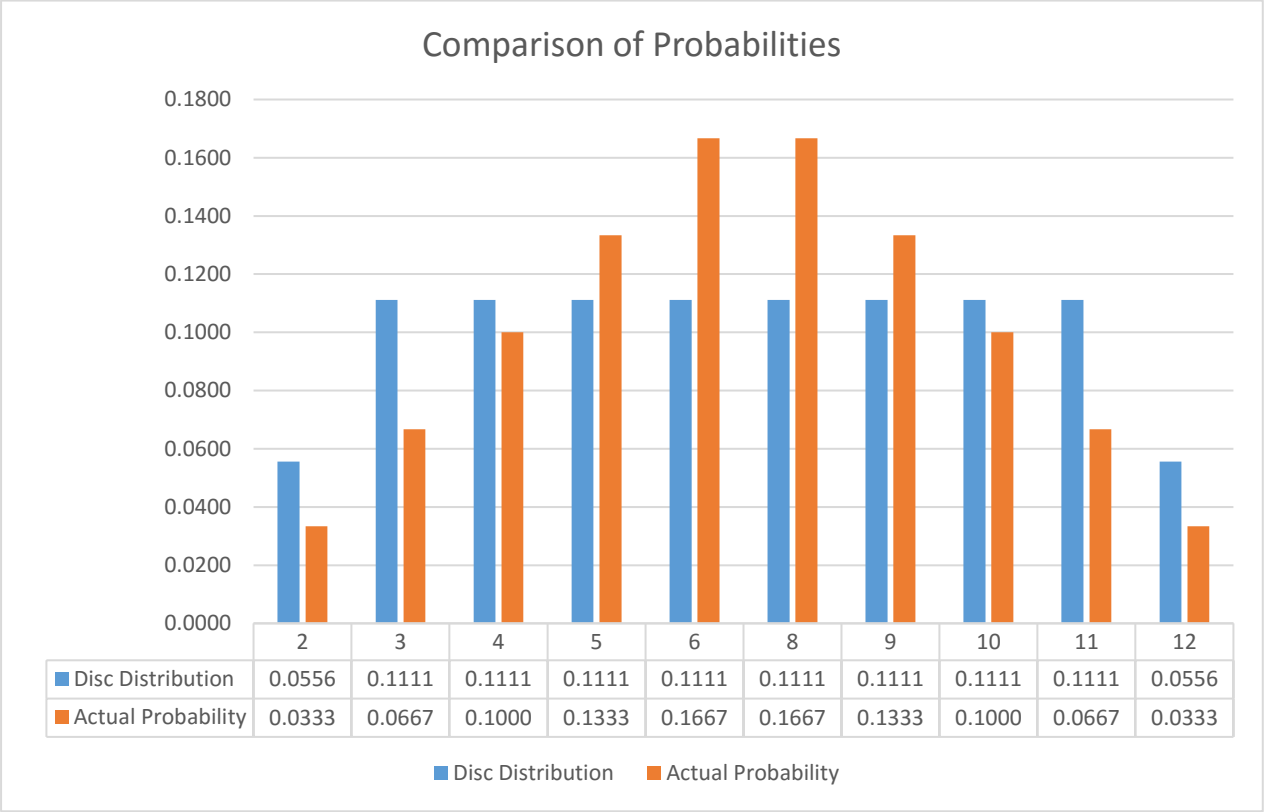

	2	3	4	5	6	8	9	10	11	12
Disc Distribution	0.0556	0.1111	0.1111	0.1111	0.1111	0.1111	0.1111	0.1111	0.1111	0.0556
Actual Probability	0.0333	0.0667	0.1000	0.1333	0.1667	0.1667	0.1333	0.1000	0.0667	0.0333

Figure 2 - Comparison of Probabilities between the disc distribution and actual probabilities

As you can see, the actual probability of rolling a 2 – 4 (out of 30) is 20%, which by reflection means the same about rolling from 10 – 12, which means that rolling a 5, 6, 8, or 9, happens 60% of the time that a 7 is not rolled, however is only represented by 44.44% of the discs. I assert that it is a significant flaw in the design of the game (further discussion of this is in section 6, where I propose a solution to this problem).

Ringing the island are water tiles, several of which are home to a port. Ports come in two types: 1) General, which allow a player to trade any three of a resource for one of another with the bank, and 2) Specialised, which allow a player to trade any two of a specific resource type for another with the bank. Otherwise, players have to trade at the ratio 4:1.

Note: the massive number of potential game boards that exist in Settlers of Catan. There are six orientations of the ports/water tiles, six orientations of the enumerated discs, 19 tiles of six types, all leading to a total of roughly 8.8*10^12 game board configurations:

$$\frac{6*6*19!}{3!3!4!4!4!1!} = 8.799558768 * 10^{12}$$

Of course this means that a player will, almost certainly, never play the exact game board twice

in their lifetimes, given the fact that assuming an individual played a game an hour (which is doable, but marginally fast given my experience), and never stopped, that player would need to play 367 billion days (or just over 1 billion years), assuming that the board is iterated after the end of every game.[30] Sufficed to say, it won't ever happen. There are similarities in boards though, which can be utilised to figure out most effective strategies for board types.[31]

During the game, the players can build settlements and cities (which are upgrades of the settlements) on the vertices of the hexes, and roads on the edges. Settlements can be placed on any free vertex that respects the distance rule: no settlement may be located on a vertex that is connected to another settlement or city by exactly one edge, or in other words, you need to have at least two edges between buildings.

As previously stated, each player starts with two settlements, as well as two road segments, which they are allowed to place in order of their position in mirror fashion (e.g. for a 3 player game, player one places his first settlement and road segment, then player two, then player three, then player three a second time, player two, and lastly player one – with similar mechanics for a 4-player game). The placement of the second settlement determines what each player's initial resource draw is, and this fact has strategic implications.

Production

On the beginning of each player's turn, they roll the two dice. The number generated (which as previously stated corresponds to one or more resource tiles save 7) determines which resources are generated on that turn for any player that as a settlement or city that is adjacent to a tile so enumerated. As previously stated, rolling a 7 activates the robber. Game rules require that upon rolling that 7, that the robber ***must*** be moved to a different hex from the one it is currently on. Whatever hex the robber is moved to is now incapable of production as long as the robber is on that hex. Furthermore, the player moving the robber is required to take one resource from one of the players whose settlements/cities is on the edge of that hex (if any). Figure 3 provides an overview of Settlers play on a single turn.

It should be noted that in the early parts of the game, most players play "nice" so will move the

[30] Assuming a year = 365.2422 days.

[31] Assuming you fixed the ports and discs, that would reduce the total number of board options to 2.44*10^11 (and given how most people play, they set up the ports and discs the same way every time). Given that, and the same assumptions previously stated, a player could play every board state in roughly 27.885 million years.

robber to a hex that has no development attached, or in the rare case in which every tile has a settlement attached to it, move the robber to a tile with the 2 or 12 disk, given the probabilistic infrequency those numbers are rolled. When starting this project, I foresaw modelling that behaviour as a challenge, but also something that needs to be done, given so far, none of the literature I've seen has mentioned that behaviour. However in later stages of the game, especially with four players, it is much harder to be "nice" and you end up in the position of what in chess is known as the zugzwang, wherein the compulsion to move (the literal meaning of the word) places you in a worse position than the one you were in, as you end up adversely impacting one or more players, which causes them to be willing to adversely impact you in the same way.

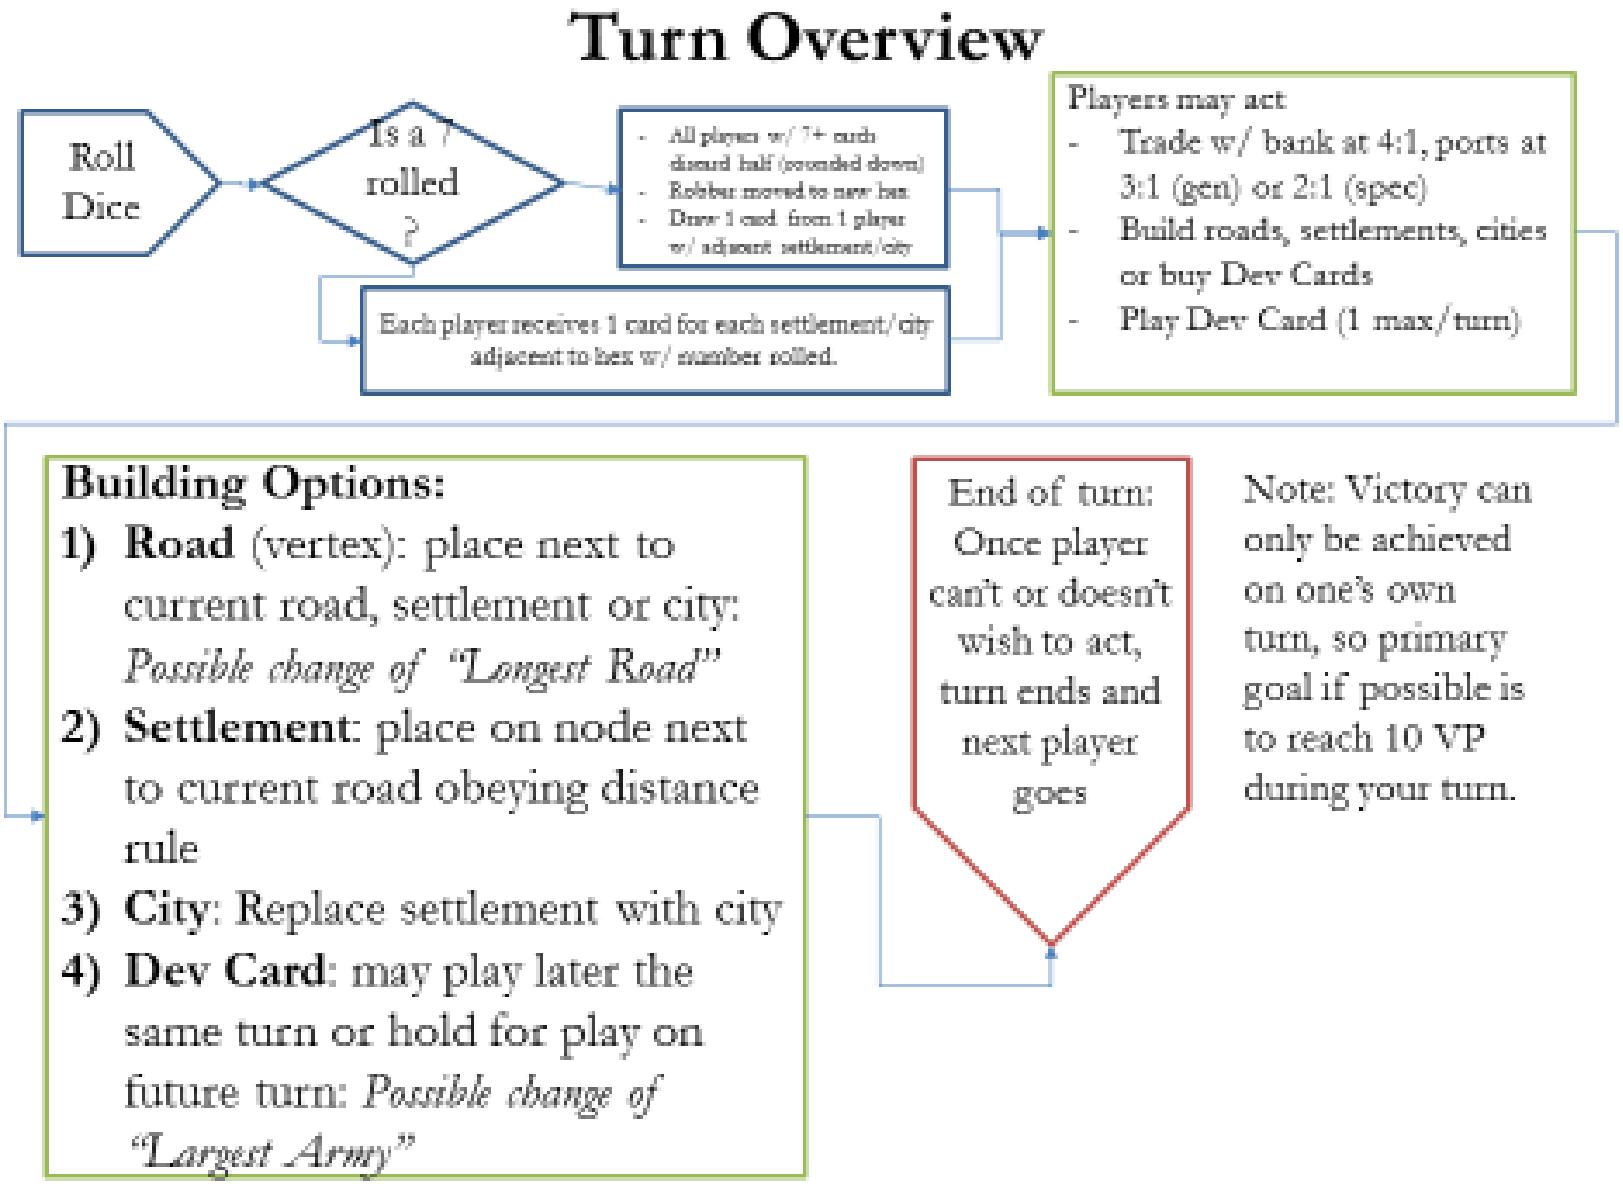

Figure 3 - Settlers of Catan Turn Overview

Structures and development cards

After the production phase, the active player may commit to zero or more actions. Said player can choose to spend resource cards on:

i) A road – which costs one brick and one wood, and must be attached to the player's current road network.

ii) A settlement – which costs one each: brick, wood, sheep and wheat, and must be attached to the player's road network, following the restrictions placed by the distance rule. The settlement allows for access to resources the settlement is at the vertex of, including port access if a player builds at one of the vertices of the port hex.

iii) A city – which is an upgrade from a settlement – which costs two ore and three wheat. A city allows the player to acquire two of whatever resource it is on the vertex of, assuming the number of the hex adjacent to that vertex is rolled.

iv) A development card – which costs one each wheat, ore and sheep. There are 25 total development cards of the following types:

- 14 Knights (which allow for taking immediate control of the robber).[32]

- 5 Victory point cards (representing various structures your civilisation has on Catan.

- 2 Monopoly cards (which allow for possessors to take control of all of a single declared resource upon use, more on this later).

- 2 Road Building cards (which allow you to gain two road segments immediately to be added to your road network).[33]

- 2 Year of Plenty cards (which allow the player to draw two resources of any type from the bank immediately).[34]

A player may use only ONE development card per turn.

[32] One must play at least 3 knights in order to attain the Largest Army Trophy (worth 2 VP, more on this later).

[33] As two road segments cost 2 wood and 2 brick, there is a fixed value. In receiving this card, resource net equivalent of 4 is gained for 3 expended, however, potentially, given that it is wood and brick gained, it may not be worth the 1 ore, 1 wheat, and 1 sheep expended

[34] The value here is much more variable, the flexibility is available that it isn't with the monopoly or road building cards, however, you do get one less net resource than is expended.

A summary table of development costs is as follows:

Development	Development Cost
Settlement	1 Wood, 1 Brick, 1 Sheep, 1 Wheat
City	3 Ore, 2 Wheat
Road	1 Wood, 1 Brick
Development Card	1 Ore, 1 Wheat, 1 Sheep

Game playing notes from my experience that I tested in this study:

1) Some players will, if they have a lot of one resource and a monopoly card, trade those resource cards away to other players, and then use a monopoly card to call them all back. However, if a player ever does this, as might be expected, no player is likely to trade with them again, not just for the remainder of that game, but also for future games (this is a game-theoretic element I've seen none of the papers mention with respect to their models).
2) If you see a development card face down for more than two or three terms, most players (and certainly experienced players) tend to assume that said card(s) are VP cards.

Trading

As mentioned above, players are allowed to trade with the bank at a ratio of 4:1 for any resource, and if they have access to a port, either general or specific, they may trade by the rules previously mentioned. A player may also trade with any other player, assuming that all players consent to the trade.

Game playing note: a common heuristic that many experienced SoC players will use is that if they see a player with 7 or more VPs showing (assuming face-down development cards held for several turns as VP cards) is that they won't trade with that player (for fear that in so doing, they'll make it possible for that player to win). This is also a characteristic that is absent from any of the papers I've seen on modelling SoC. I assumed this as a given when drafting the strategy survey I administered as a part of data accumulation portion of this thesis. In my observations of Settlers games, as well as from the administered survey, this does not in fact always happen (some more experienced players traded with an opponent at 7 VP, however, it was usually a trait of the less experienced).

Victory Points and Awards

Various developments are worth VPs. Settlements are worth 1 VP, cities are worth 2. The first player to build 5 or more road segments gets the award of the "Longest Road" trophy, which is worth 2 VPs (which can only be stripped from that player if another player builds a longer road). The first player to play 3 Knights gets the award of the "Largest Army" trophy, which is also worth 2 VPs (which of course can only be stripped from that player if another player plays more knights). Along with the previously stated development cards that are worth 1 VP each, these are the only ways to gain VPs. The first player to reach 10 VPs is the victor of the game.

Strategy Types

Coming into this study, based on my previous experience in Settlers of Catan, I distilled six strategies that players utilised in the game. Those strategies are below.

Primary/Basic Strategies

i) Brick-wood: This strategy is based on the premise of building roads and settlements. As there are only three brick hexes (same as ore), unlike the rest that have four, and as both roads and settlements are necessary for early expansion, it's often effective. This strategy also allows for acquisition of the "Longest Road" award and the two victory points that it yields.

ii) Ore-wheat: This is the other primary strategy. This strategy allows you to convert your settlements into cities, and to purchase development cards. Development cards lead to both the "Largest Army" award, as well as the VP cards.

Mixed Strategies and more advanced concepts

iii) Blended: this is trying to thread the needle between the two primary strategies listed above.

iv) Rare resource Monopoly: Find the rarest resource on the board, and monopolise it. Of course, this will maximise the player's trading ability. Brick and ore are the most likely candidates for this, given that there are fewer of those hexes than of the other resource types.

v) Abundant Resource/Port: Find what will be the most prevalent resource on the board, get access to it, and to the specific port that specialises in it. This will allow a player, for instance, to trade a wood for two sheep with another player, knowing full well that they can, in turn, transform those two sheep into whatever resource they actually need at the time.

vi) Number diversity: Given the game rule that puts a player at risk if they hold more than seven cards at a time, it is often wise to normalise their resource flow throughout the game by having access to the highest number of numbered hexes. Ideally, the player should try to play at the vertices that allow for access to hexes enumerated 4, 5, 6, 8, 9, and10. This leads to maximisation of the probability that a player will gain at least some resource each turn.

Ideally, if the number discs allow for it, a player should try to get both wood and brick at the same time, or ore and wheat at the same time, since they are complementary resources in the basic strategies.

It should be noted that you could win a game without wood or brick (if you upgrade your two settlements into cities, and with a development card focused strategy, you acquire the Largest Army trophy and 4 VP cards for 10 total VP), however, you cannot win a game without ore. (5 settlements and longest road, for 7 total VP). This also means that you cannot win the game without sheep, insofar as a strategy that ignores building roads and settlements in favour of a heavy development card focus would also require that you have at least 7 sheep resources, and given the statistical unlikelihood of a single player drawing 4 of 5 VP cards is so low as to be realistically non-existent. It should be noted that in the 19 primary games observed for this study (more on that in section 6), no one person performed such a feat. As you can tell from Figure 4, the number of VP cards drawn in gameplay is roughly speaking, exponentially decreasing.

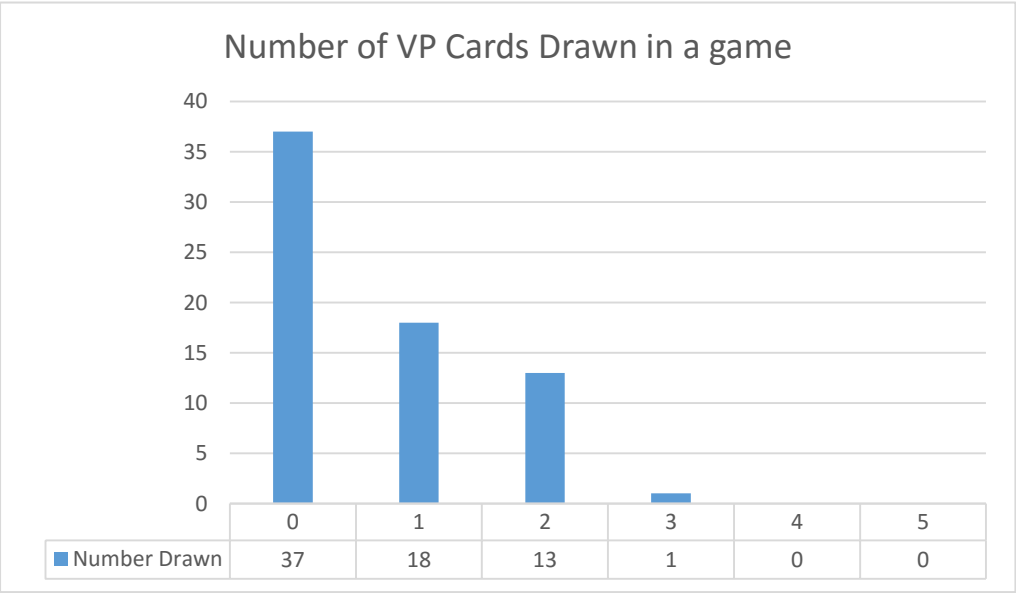

Figure 4 - Number of VP Cards Drawn in a Settlers game

Most players don't even acquire one VP card (37), less than half of that number (18) even acquire one VP card, and only one player in all 19 games acquired 3. So a strategy based on avoiding all wood and brick is theoretically feasible, however, it is apparently realistically infeasible, and so should be excluded unless and until data exists to indicate its viability.

It also should be noted that a player can, in fact, acquire the 10 VPs necessary for victory without building any additional settlements or upgrading their settlements to cities at all, the player attempting to employ such a strategy would have to acquire both trophies (Longest Road and Largest Army), as well as acquire at least 4 of the 5 VP cards. Given the reasons stated in the previously explaining the no wood-brick strategy, this strategy is seen as just as infeasible, practically speaking, as that one is, however, it is a strategy that relies on all the resource types, so is more feasible in that respect, even if it is no more apparently feasible in outcome.

Of the six strategies enumerated above, it is easy to assert that Wood-Brick and Ore-Wheat are what in a game-theoretic sense are known as pure strategies, and that Blended is a mixed strategy of the previous two (as it is defined as such).[35] Number diversity is arguably the reverse side of the coin to blended's obverse (as it too relies on getting a mix of resources, however, it comes with a different focus than blended) since its objective is to maximise resource gain, probabilistically, which is also a balancing act.

Expansion

The expansion, which allows for either five or six players to play the game, expands the game board from 19 to 30 hexes in a 3-4-5-6-5-4-3 configuration. Those 11 additional hexes include: 1 desert, 2 wheat, 2 wood, 2 sheep, 2 brick and 2 ore. As would be expected, there are additional player piece sets for two additional players, 25 additional resource cards (five of each type), 9 additional development cards (6 knights, one each of monopoly, road building, and year of plenty) and 28 numbered discs like the ones described in the basic game. Given that this expansion will not be discussed, I will not expend any more time describing the rules and differences between the base game and the expansion.

[35] It should be noted though that one can regard a pure strategy as a degenerate case of a mixed strategy, in which that one particular pure strategy is selected with probability 1 and every other strategy with probability 0, however as mixed strategies are viewed as the children of pure strategies, not vice versa, it is not recommended to do so.

Final Thoughts

Given that Settlers of Catan is a turn-based sequential game, it can be modelled by a game tree. This fact allows for, and I think invites, the utilisation of the Monte Carlo Tree Search method, which is why it has been attempted a few times so far. This game is worthy of study, it may have fewer game board states than Go, but because of its aleatoric elements, as well as the issues surrounding trade, they bring a level of complexity that that rivals if not exceeds Go's.

5) Literature review

The current literature, such as it exists given the recentness of the development of MCTS, provides a framework with which to analyse questions, but for which there is a lot left as of yet undefined or unexplored as well as disparate elements that can be better integrated. Given the catholic nature of this study, it should be noted that sources in the specialities of economics, political philosophy, psychology, and other social and behavioural sciences were checked along with sources from computer science, mathematics, and the physical sciences.

One of the more consistent features of the literature reviewed is that negotiating and trading between players was usually ignored. The general reason is that it was seen as too complex of a problem, and there was too little data upon which to create algorithms to use for this purpose.

Lack of rationality by players

As Guhe and Lascarides in 2014 noted, given human deviance from utilising analytical solutions (if they existed) as well as the fact that the game tree for SoC is not surveyable (a given example is the fact that the options for negotiating trades in natural language are not bounded), utilisation of algorithms such as backwards induction for computing expected utilities do not work on their own.

The response they propose to address this problem is the development of a symbolic model that consists of heuristic strategies for playing the game.

Pseudo-code to describe games

Méhat and Cazenave in 2010 describe the use of Game Descriptive Language (GDL), which is used to describe games that is based on first-order logic, which can be used to describe finite games with an arbitrary number of players. They do note however that GDL cannot be used to describe games

which contain aleatory (that is to say random) elements.[36] To deal with the aleatory element, GDL-II extends GDL to allow for the description of random elements and incomplete information.[37]

MCTS in Settlers of Catan

Szita *et al* (2010) implemented an agent that is able to play against computer-controlled as well as human players. They demonstrated basic success in adapting MCTS to multi-agent environments and presented two approaches to providing the agent with limited domain knowledge. One part of their analysis was the effect of starting position, in which they demonstrate that in a 4-player game, with random agents, seat order has an effect on final ranking in a game, with position 1 having the strongest position, winning 30.42% of games in their simulations, with an average score over all games of 5.98. Position 4 was second best, winning 23.44% of games, and with an average score of 5.63 over all games.

AI method to decide in Chess

Lai (2015) created an AI engine called Giraffe that evaluates chess positions and narrows down the search tree to the most beneficial moves, using machine learning methods, thus pruning the tree to just a few branches. His work does not actually utilise MCTS, but using machine learning, his work parallels it, thus can be adapted utilising MCTS.

Richardson's Conflict Modelling

As Brian Hayes notes in the beginning of his précis into the life and work of Lewis Fry Richardson with respect to modelling wars, "Demographically, it hardly matters. War deaths amount to something like 1 per cent of all deaths; in many places, more die by suicide, and still more in accidents. If saving human lives is the great desideratum, then there is more to be gained by prevention of drowning and auto wrecks than by the abolition of war." After noting that war as an activity is one that engages and inflames human passions, he then states "Stepping outside the bounds of one's own culture and ideology is also a challenge, not to mention the bounds of one's time and place. We tend to see all wars through the lens of the current conflict, and we mine history for lessons convenient to the present purpose. One defense against such distortions is the statistical method of gathering data about many wars from many sources, in the hope that at least some of the biases will balance out and true patterns will emerge. It's a dumb, brute-force approach and not fool

[36] Méhat and Cazenave, *Combining UCT and Nested Monte Carlo Search for Single-Player General Game Playing.*,
[37] Thielscher, *A General Game Description Language for Incomplete Information Games*,

proof, but nothing else looks more promising."[38]

Richardson was amongst the first to attempt to model war quantitatively, and he did it because he was a pacifist for religious reasons (his family were Quakers) and he thought that if he could quantify war, that he could help bring an end to war. His idealism must be applauded, but as Plato has been attributed as saying: "Only the dead have seen the end of war".[39] Lewis Richardson was no crank, as noted British meteorologist Julian Hunt noted in his survey of the life and career of Lewis Richardson: "His name can be added to the list of inventors of computational mathematics (with von Neumann, Courant, and Turing), of modern meteorology and fluid mechanics (with Bjerknes, Taylor, and Prandtl), of quantitative techniques in psychology and social sciences (with James), and of analysis and modelling of complex systems (with Norbert Wiener). In each of these fields his work is still being cited. Between 1980 and 1984, in one citation index, there were over 200 references to his work."[40] It was Richardson's work with meteorology that lead to his development of mathematically based conflict models, an activity that began in 1919 after the end of World War I and continued unto his death.

Richardson developed two things with respect to his conflict modelling that I think are relevant to mention in this study. The first is his creation of a base-10 logarithmic model of conflict, which he used to chart the period of 1820 to 1950, of which he identified 315 conflicts. What he found was that the two magnitude-7 wars (WWI and WWII) were the cause of 60% of all deaths in warfare during that period. Richardson identified seven wars that were magnitude-6 (causing the death of 500,000 to 2,000,000), etc.[41]

Richardson's model is depicted in Figure 5 (orange bars depict number of casualties, green bars the number of conflicts of that magnitude):[42]

---

[38] Hayes, *Computing Science: Statistics of Deadly Quarrels*, *American Scientist.* This article is one of two that surveyed the life and work of LFR, the other was published in 1998 in the Annual Review of Fluid Mechanics and written by noted British meteorologist Julian Hunt. Both are referenced in this section.

[39] There is debate about if Plato actually said it, the earliest known attribution was made by George Santayana in the early 20th century. Anyone who's interested in reading about that is encouraged to read the discussion on the Plato Dialogues website (http://plato-dialogues.org/faq/faq008.htm).

[40] Hunt, "*Lewis Fry Richardson and His Contributions to Mathematics, Meteorology, and Models of Conflict*,", *Annual Review of Fluid Mechanics*

[41] The seven massive death conflicts listed by Richardson are, in chronological order, using the names he assigned them: the Taiping Rebellion (1851-1864), the North American Civil War (1861-1865), the Great War in La Plata (1865-1870), the sequel to the Bolshevik Revolution (1918-1920), the first Chinese-Communist War (1927-1936), the Spanish Civil War (1936-1939) and the communal riots in the Indian Peninsula (1946-1948).

[42] A maginitude-0 event on this scale is a quarrel in which 1 to 3 people died, a magnitude-1 event is one in which 10-32 die, a magnitude-2 event is one in which 100-320 die, etc.

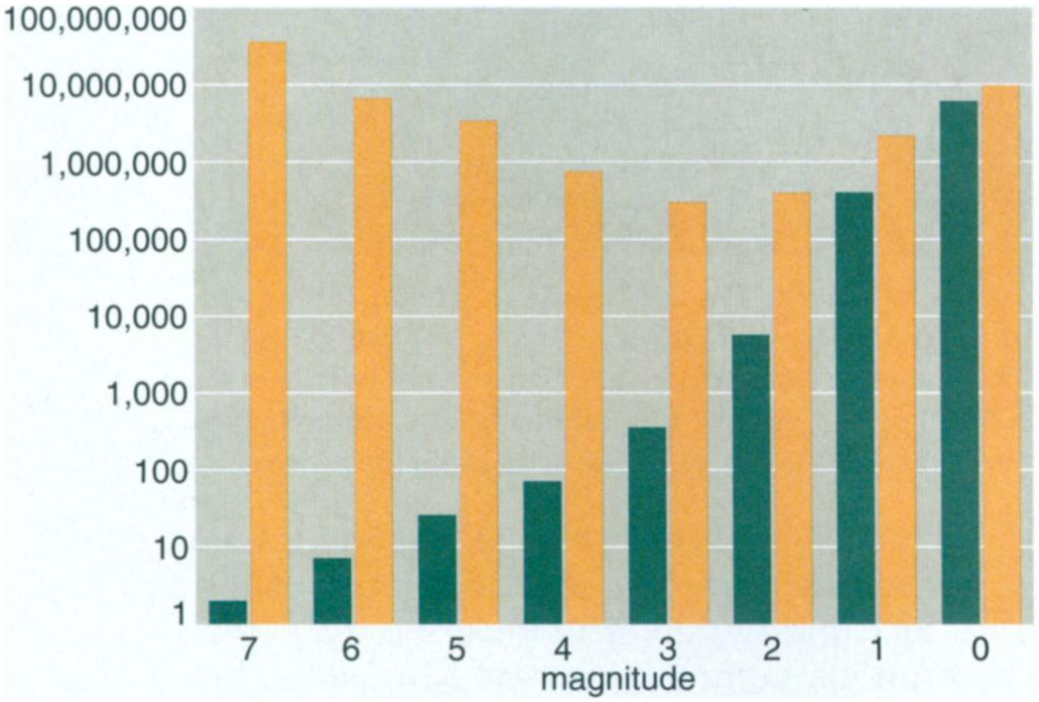

Figure 5 - Richardson's Magnitude of Wars Model

I think this model is rather revealing in what it conveys.[43] In order to deal with the issue of casualty ranges, Richardson utilises the logarithm based approach to make the process of analysis consistent. I used a comparable approach in order to deal with the issue of ranged values for games played in order to develop a unitary experience score for survey participants (which you can read about in section 6).

Richardson also found out that the frequency of the outbreaks of war was very closely modelled by the Poisson distribution, as is depicted in Figure 6:[44]

[43] Graphic from Hayes, *op. cit.*, pg 12
[44] *Ibid*

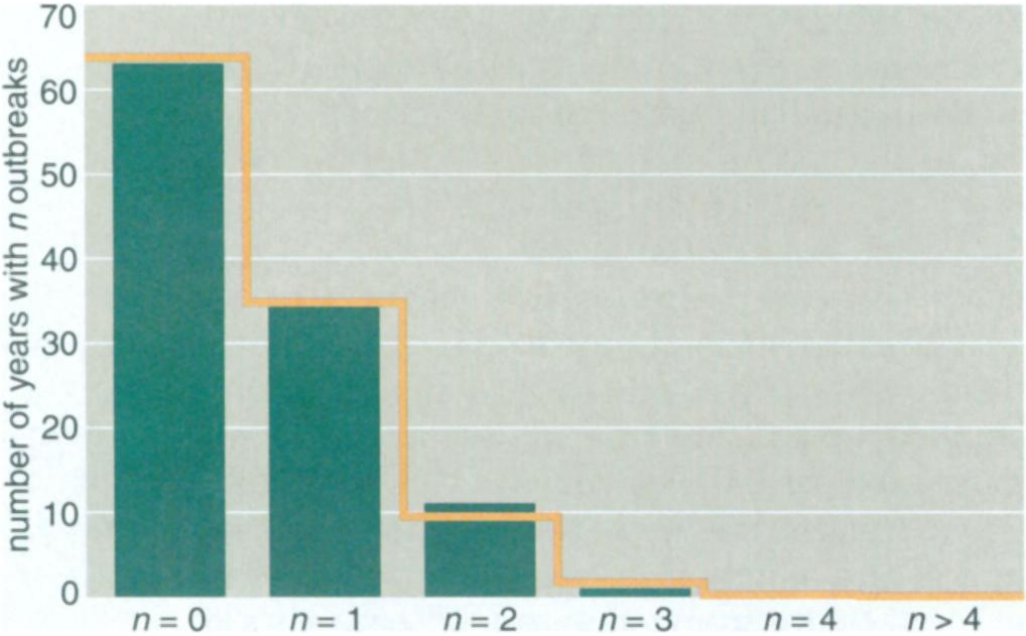

Figure 6 - Richardson's Demonstration of War as a Poisson Process

This of course strongly suggests that the not just are events in war often times random and unpredictable, that the outbreak of war itself is fundamentally a random process and that wars themselves are randomly distributed accidents.

Game theory

Even though game theory is old (von Neumann and Morgenstern published *Theory of Games and Economic Behaviour* in 1944), it has a lot of tools that are useful and valuable in this study. Game trees, expected utility, subjective probability, the need for randomising strategies, and so on.

The fact that Settlers is an extensive-form game with aleatory elements is relevant here, since chance events, which Settlers has many of because of it being a dice-driven game, as well as the randomness of game boards, can be encoded as "moves by nature". How do we know that Settlers is an extensive-form game? As it happens, Sergiu Hart wrote a study entitled *Games in extensive and strategic forms* in a book he co-edited with Robert Aumann in 1992 in which he presented the following list of characteristics of an n-player extensive-form game:

- A finite set of *n* (rational) players
- A rooted tree, called the *game tree*
- Each terminal (leaf) node of the game tree has an *n*-tuple of *payoffs*, meaning there is one payoff for each player at the end of every possible play
- A partition of the non-terminal nodes of the game tree in *n*+1 subsets, one for each (rational) player, and with a special subset for a fictitious player called Chance (or

Nature). Each player's subset of nodes is referred to as the "nodes of the player". (A game of complete information thus has an empty set of Chance nodes.)

- Each node of the Chance player has a probability distribution over its outgoing edges.
- Each set of nodes of a rational player is further partitioned in information sets, which make certain choices indistinguishable for the player when making a move, in the sense that:
    - there is a one-to-one correspondence between outgoing edges of any two nodes of the same information set—thus the set of all outgoing edges of an information set is partitioned in equivalence classes, each class representing a possible choice for a player's move at some point—, and
    - every (directed) path in the tree from the root to a terminal node can cross each information set at most once
- the complete description of the game specified by the above parameters is common knowledge among the players

Time could be spent on how each of these criterion are met by Settlers of Catan here, however, it is self-evident given the game description previously provided. The fact that Settlers is a finite extensive form game also means, as Martin Osborne pointed out in his book *An Introduction to Game Theory*, it has a subgame perfect equilibrium.[45] This fact runs against the assertion of Guhe and Lascarides, who cite the non-surveyability of the issue of trade as the reason why Settlers can't be solved. I would assert that, consistent with game theory, the fact that even if the language of trade can't be entirely surveyed (and that is true in my experience with this study) that the action of trade is surveyable, thus solvable utilising backwards induction. Trades happen when both players believe it to be in their strategic interest to trade. When evaluating games by backwards induction, you don't need to know what language was used, just that a trade happened, and the terms by which it occurred.

The expansion of game theory into the computer science realm has started to occur over the two decades, Noam Nisan *et al.* put out a book in 2007, called *Algorithmic Game Theory* that helped formalise the subject. In 2008, computer scientist Yoav Shoham wrote an article entitled

[45] A strategy profile is a subgame perfect equilibrium if it represents a Nash equilibrium of every subgame of the original game. Less formally, this means that if 1) the players played any smaller game that consisted of only one part of the larger game and 2) their behaviour represents a Nash equilibrium of that smaller game, then their behaviour is a subgame perfect equilibrium of the larger game. (cf Osborne, *An Introduction to Game Theory*).

Computer Science and Game Theory that was a review of the interplay between these two fields.[46] It is from algorithmic game theory that the tool of multi-agent systems originates, and that in turn is vital to managing any realistic simulation of a game as complex as Setters of Catan. My search of the literature found some examples in which an agent-based approach was attempted.

Multi-agent method (1)

Branca and Johansson in 2007 describe the utilisation of a multi-agent system [types: 1) settlement, 2) evaluator, 3) trade, 4) mediator], which they combine with JSettlers to create an AI bot to play SoC that utilises a weighted utility model to evaluate AI player action.[47] A stated limitation of their work was that AI bot was only able to trade with the bank, and not with other players. Although not reaching the goal, they demonstrated that it was possible for an MAS-based solution to get close to the performance of monolithic solutions.

Multi-agent method (2)

Saleem and Raees Natiq in 2008 also implemented a multi-agent system approach in JSettlers for their MSc thesis in attempting to solve the problem. [48],[49] One of the unique features of this study was they actually attempted to solve the issue of modelling trading practice in Settlers. They propose a general architecture that I believe to be a decent framework upon which to organise such an effort. That architecture is as follows:[50]

[46] This review was published in the Communications of the ACM, which is of course the monthly magazine for the members of the Association for Computing Machinery (ACM), one of the two umbrella organisations in the US for academic and scholarly interests in computing.

[47] JSettlers is an open-source computer model of SoC, written in Java, that has been commonly used in SoC oriented research.

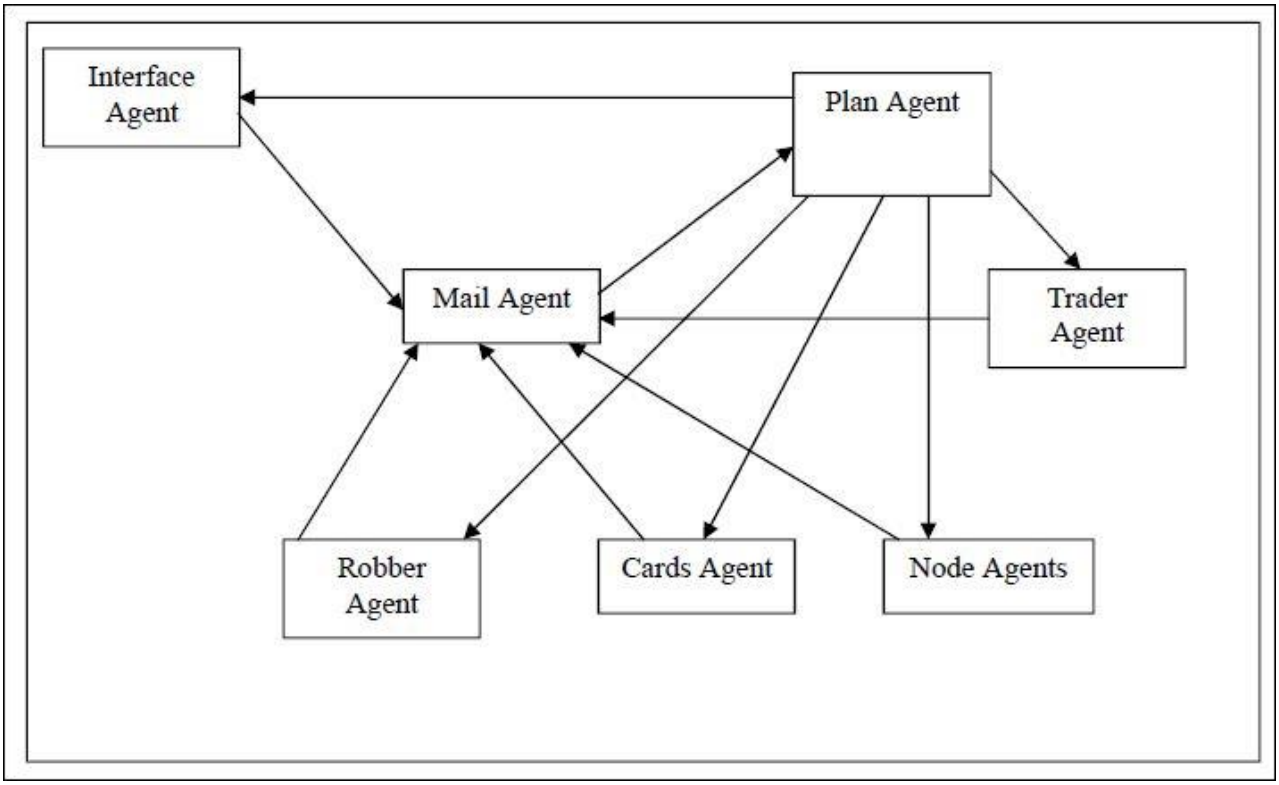

Figure 7- Multi-Agent System General Architecture

The constituent parts of the multi-agent system are as follows (the descriptions are theirs, although in some cases the phrasing has been altered for clarification):

- Interface Agent: responsible for communicating with the Plan Agent which in turn communicates with the other agents, and sends the response back to the Interface Agent, which then executes the proposed action.
- Plan Agent: responsible for planning out the course of action. The Plan Agent sources outs and calls the appropriate Agent to execute the task, gets the result from that agent and sends back the response to the Interface Agent.
- Node Agents: a collection of 52 agents which are resident on the board locations where settlements can be possibly built. Each node agent is resident on a particular hex on the board. This agent verifies if a settlement can be built on any of the locations adjacent to the hex on which this agent resides. It sends back the response to the Plan Agent.
- Cards Agent: As the name implies, this agent is responsible for playing development cards. It checks at a particular move, which card would be beneficial to play and sends the response to the Plan Agent. It also manages the bot's resource cards.

[48] Saleem and Raees Natiq submitted a single thesis for their MSc degrees (I assumed it was because it was a joint project).
[49] Both of these works reference Thomas's PhD dissertation: *Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic*, which I also checked in order to ascertain their use and derivations.
[50] *A Multi-agent player for Settlers of Catan - A MAS Player in game playing*, section 3.1

- Robber Agent: responsible for moving the Robber on the board. It basically analyses the possibility to gain maximum advantage by moving the Robber to a hex, where maximum damage can be given to other players.[51]
- Trader Agent: responsible for making compelling trade offers to other players in exchange of resources needed by the player to build settlements that could give the player an advantage during the course of play.
- Mail Agent: responsible for dispatching messages amongst all the agents.

This system was a noble attempt to solve the problem, but it too failed with respect to trading and the strategic elements of the same. They ordered the iteration as follows: trade with other players, trade with a specific or general port, and no trade at all. This trading agent ignores trading with the bank at the 4:1 rate entirely, which I think is a severe limitation of this system. Their multi-agent system breaks down with respect to trade and makes unrealistic assumptions, but they are to be commended for making the attempt.

More on the Research Questions and what the Literature says.

As I discussed in the introduction, there are three primary research questions this study attempts to address: 1) Why use an MCTS-based model? 2) How can we use knowledge to improve MCTS? 3) How can we arrive at a balance between exploration and exploitation using MCTS?

Regarding the first question, why use an MCTS-based model at all, when you can solve a problem through an exhaustive method. A reader might be curious as to why a Markov Decision Process (MDP) based model isn't employed in this context. As the reader may recall, a Markov decision process is defined as the 5-tuple: $(\boldsymbol{S}, \boldsymbol{A}, \boldsymbol{P}(*,*), \boldsymbol{R}\ (*,*), \boldsymbol{\gamma})$, where:

- $\boldsymbol{S}$ is a finite set of states
- $\boldsymbol{A}$ is a finite set of actions
- $\boldsymbol{P_a(s, s') = \Pr(s_{t+1} = s' \mid s_t = s, a_t = a)}$ is the probability that action $\boldsymbol{a}$ in state $\boldsymbol{s}$ at time $\boldsymbol{t}$ will lead to state $\boldsymbol{s'}$ at time $\boldsymbol{t+1}$.
- $\boldsymbol{R_a(s, s')}$ is the expected or actual immediate reward after transitioning to state s' from s.
- $\boldsymbol{\gamma \in [0, 1]}$ is the discount factor between current and future rewards. Normally γ is set near one.

[51] I would not choose to use this focus, since most people tend to attempt to minimise damage to any one player (either by spreading out their use of where to lay the robber or by playing the robber on unoccupied hexes).

What a Markov decision process does is finds a "policy" for the decision-maker (denoted as π), that is a distribution over given action states $\pi(a|s) = \mathbf{Pr}(\boldsymbol{a_t = a \mid s_t = s,})$. A policy in this sense dictates all behaviours of the agent, and once an MDP is merged with a policy in this fashion, this fixes the action for each state and the resulting arrangement behaves like a Markov chain, and can be solved by dynamic programming.

Settlers has a finite set of states at every turn, a finite set of actions, a series of probabilities of this type, rewards that can be classified in such a way, and the discounting factor is a well-trod tool in dealing with financial and economic benefits across time. So, obviously, an MDP-based method could be used here to derive an almost optimal sampling for the domain of the problem, the biggest issue here is that the amount of work needed is exhaustive, ergo when the domain of the problem is huge, the MDP method won't succeed in generating quality results within a reasonable amount of time, which is why the need of an alternative is urgent.

Regarding the second question, on utilising knowledge to improve modelling efforts, whilst rational utility theory would lead people to reasonably assume that all actors will engage in behaviour to maximise their expected utility at all times. Dan Ariely notes in his book *Predictably Irrational* that the vast majority of people consistently act in ways that are systematically and predictably non-optimal. In the vast majority of cases, the irrationality is something within a bounded range that is foreseeable, which assists modelling efforts. The second objective of this study is to refine the utilisation of MCTS to account for this.

Regarding the third question, on balancing exploration and exploitation in models, SoC does not have an analytic solution under the formulation of Guhe and Lascarides. [52] However as Hart noted the conditions of an n-player, finite, extensive form game (which Settlers fits), and all extensive-form games have a one or more subgame perfect equilibria, an optimal solution can be simulated, against which real world play can be measured. As Szita *et al* noted in 2010, if all options are weighted equally, the strategy played is fundamentally weak, and the MC simulation suffers as a result, so assigning larger initial weights to options that heuristic knowledge informs are more promising.

Conversely, Szita *et al* also note, that if the selection strategy is too deterministic, the exploration

[52] Guhe and Lascarides, Game Strategies for The Settlers of Catan, *2014 IEEE Conference on Computational Intelligence and Games*, 2014

of the search space becomes too affected, ergo the MC simulation also suffers.[53] So a balance between exploring and exploiting must be found in the simulation strategy. Finding that balance is the second objective of this study.

The first formula for balancing exploitation and exploration in games, called UCT (*Upper Confidence Bound* 1 *applied to trees*), was introduced by Kocsis and Szepesvári (2006). The idea underpinning this algorithm is to sample actions selectively; therefore choosing sub optimal actions early in sampling will improve the performance. They recommended choosing in each node of the game tree the move, for which the expression:

$$\frac{wi}{ni} + c\sqrt{\frac{\ln t}{ni}}$$

has the highest value, that is to say that in the end, it chooses the arm with the best upper confidence bound. In this formula:

- w_i = number of wins after the i^{th} move
- n_i = number of simulations after i^{th} move
- c = exploration parameter, theoretically root(2), in practice usually chosen empirically
- t = the total number of simulations, equal to the sum of n_i.

The first component of the formula above corresponds to exploitation; it is high for moves with high average win ratio. The second component corresponds to exploration; it is high for moves with few simulations. Most contemporary implementations of MCTS are based on some variant of UCT.

Practically speaking, the value of UCT is that it has a smaller error probability when terminated properly and convergence to the almost optimal result if enough time is considered, and is much more efficient than minimax in this model, and the difference in performance grows as the tree deepens. The reason for that effect is that randomly sampling nodes of the tree leads to a reliable estimate of the true minimax value of the tree, even with error in the heuristic. Good guidance from the sampling is a primary factor to UCT's success.[54]

[53] Szita *et al*, Monte-Carlo Tree Search in Settlers of Catan., *Proc. Adv. Comput. Games, 21–32*, 2010
[54] Raghuram *et al*, "On the Behavior of UCT in Synthetic Search Spaces", 2011

6) **<u>What Data Was Collected</u>**

Two primary methods were used to collect data for this study. The first was the use of a general strategy survey, which had a number of questions on strategy and trading preferences within Settlers games.[55] The second was observational data drawn from observing 39 Settlers games of two classes, 19 primary games, and 20 secondary games (more details will be given later in this section).

General Surveys

The first form of data collection for this study was a general strategic survey, which covered participants' knowledge of strategy within the game Settlers of Catan, as well as personality and demographic information of the participants that I presumed to be relevant. Two main forms of personality information that were evaluated were Machiavellianism level (as measured by MACH-IV test scores) and questions to evaluate risk tolerance.

Demographic data included gender, years of formal education, degree type, veteran's status, along with game specific questions about years played and games played. It was assumed that a person's race/ethnicity would have no impact on results, so, it was excluded. Gender was also assumed wouldn't bear out on strategic competence, but would potentially impact MACH-IV scores so what strategies might be employed by players, that is to say that it was assumed that women would be on the whole less Machiavellian than men amongst other things, so gender information was solicited to test that. Veteran's status was also asked about, with the assumption that military service would in fact bear on a person's strategic thinking competence as well as personality information such as their risk tolerance and MACH-IV score, however only nine survey participants who had military service responded, therefore that characteristic had to be disregarded from the analysis for this study, since nine was far too few to perform any respectable statistical analysis, however, that is a potential line of inquiry, insofar as the veterans population that did respond was more Machiavellian than the total population.[56,57]

---

[55] As is required by Federal law and University of Virginia policy, the strategy survey was reviewed by the Social and Behavioural Sciences Institutional Review Board of UVA, was approved, and was assigned the project number: 2016-0081-00

[56] The boxplot data of veterans v. non-veterans MACH-IV Scores

Veterans MACH-IV scores: Q1: 55, Q2: 60, Q3: 77, Lower Whisker: 50, Upper Whisker: 84

Non-Veterans MACH-IV scores: Q1: 52, Q2: 60, Q3: 67, Lower Whisker: 34, Upper Whisker: 87

[57] When discussing this theory with a friend of mine who is a submariner in the US Navy (LCDR Brandon Ray), he commented how there might be a differential between those with line/combat arms experience and those whose military experience was in a support role such as logistics, and staff jobs such as medical, legal, or chaplain (a theory that I agree is probably also accurate, but of course need data to test it). As in, I think military service provides an ontological shift, that the sheer act of undergoing initial entry training in the military will make someone more Machiavellian than not, however within that shift, operations people are likely to be more Machiavellian than logistics or staff people. Along with his years of experience as a submariner, Brandon is also a current graduate student at the University of Washington, having attained a Masters in Atmospheric Sciences. His baccalaureate degree is from Northwestern University.

Using various listservs both through the University of Virginia and personal contacts to spread word of this study, 174 people took the Settlers strategic survey for this study, of which, 167 valid (by which it is meant complete) entries were received, upon which all analysis was conducted.[58]

Of the 167 survey takers with completed surveys, 30.5% (n = 51) were female, 69.5% (116) were male. The mean years of education was 16.7 and a standard deviation of 2.2, indicating that the sample averaged at least a bachelor's degree, and roughly 40% of my sample had post-graduate education (in my survey: 16 = BA/BSc, 17 = attended graduate school, no degree attained, 18 = MA/MSc, 19 = professional doctorate (JD, DPharm, etc) or post-coursework graduate student, 20 = PhD/ScD/MD).

Figure 8 is a summary report of the data

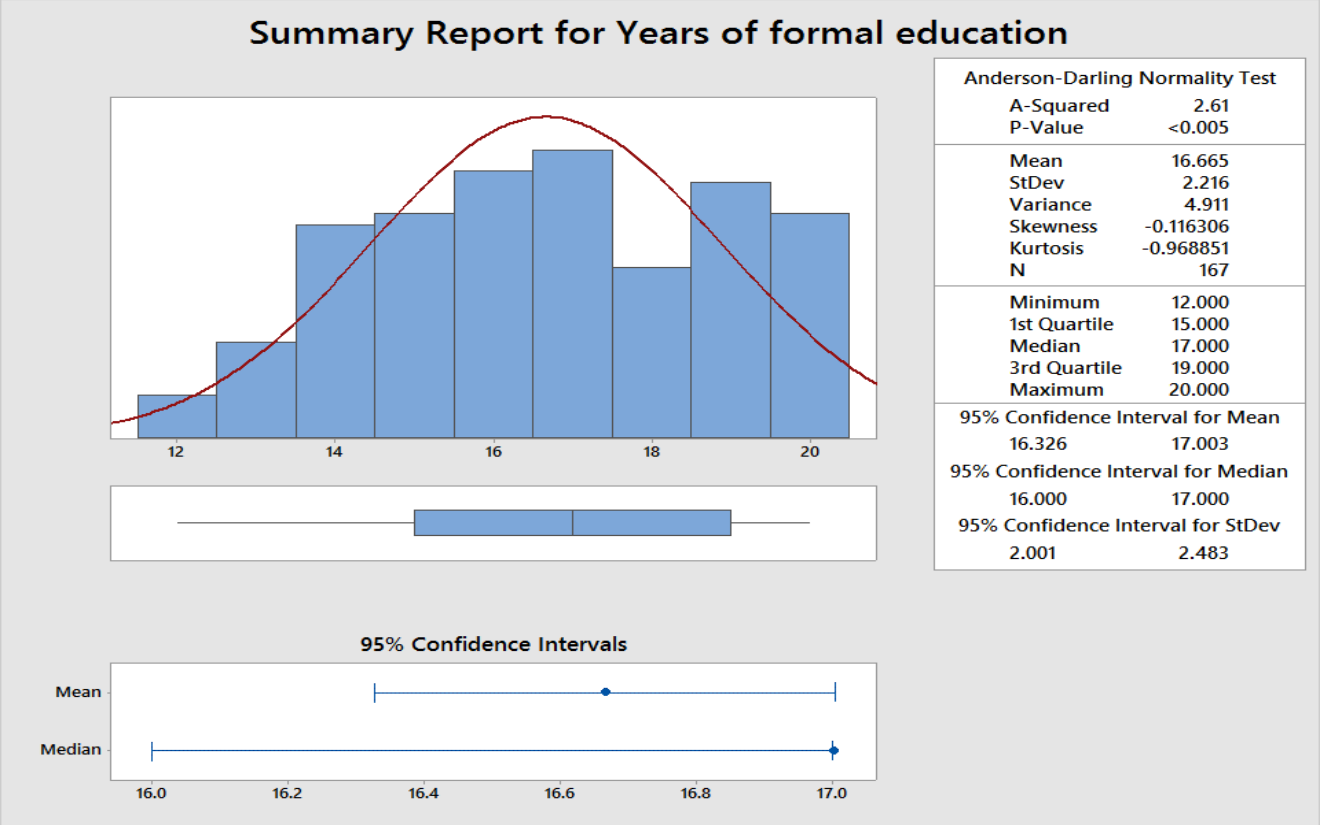

Figure 8 - Years of Formal Education Completed

Given that the survey participants were all volunteers, it is a self-selecting pool, however, what this summary report does allow is to confidently state that the pool is highly educated, given they

[58] Of the invalid entries, one I ruled out because the participant looked to be under age 18 on the basis of the information provided (and given IRB, I was not authorised to utilise the participation of minors), and the rest, their entries were incomplete, so I was able to use their data for some things.

averaged (in all senses of the word) at least a baccalaureate degree.

The basis of using the Machiavellianism as a gauge for explaining behaviours in games has existed over a decade. Looking for explanations for observed behaviour experimental games, Gunnthorsdottir *et al* found that particular individual choices which didn't correspond to standard Nash Equilibrium self-interest predictions, but instead that in trust games, those with high Mach-IV scores tended to follow the "homo economicus" equilibrium strategies, whereas those with low scores tended to deviate from equilibrium and acted with conventionally accepted social preferences and moral norms.[59]

As defined by the *Oxford English Dictionary*, Machiavellianism is: "the employment of cunning and duplicity in statecraft or in general conduct". In the 1960s, Richard Christie and Florence L. Geis developed a test for measuring a person's level of Machiavellianism, the Mach-IV test. Using their scale, Christie and Geis conducted multiple experimental tests that showed that the interpersonal strategies and behaviour of "High Machs" and "Low Machs" differ, and these basic results have been repeatedly replicated.[60]

I looked at the relationship between degree type and MACH-IV as well, suspecting that education in business/economics and/or law would correlate to higher Machiavellianism scores, and whilst the data I have would indicate there is such a relationship, too few responses of people with backgrounds in both categories meant that no statistically valid claim can be made in that regard.

The questions used in the survey to assess risk tolerance were taken from Tversky and Kahneman. Tversky and Kahneman have two different framings of a question with the same expected value but differing variance. They use it to demonstrate that people's preferences are not logically consistent. By asking both questions next to each other in the survey developed for this study, the intent was to induce people to think consistently about their risk tolerance.

Here is how the first question appeared in my survey: "Imagine that your country is preparing for the outbreak of an unusual disease, which is expected to kill 600 people. Two alternative

[59] Gunnthorsdottir, A., *et al*; "Using the Machiavellianism instrument to predict trustworthiness in a bargaining game". *Journal of Economic Psychology 23*, 49-66 (2002)
[60] Christie, R. andGeis, F. (1970) "Studies in Machiavellianism". Academic Press

options to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the options are as follows: If Option A is adopted, 200 people will be saved. If Option B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. Which of the two courses of action would you favour?"

Here is how the second question appeared in my survey: "Imagine that your country is preparing for the outbreak of an unusual disease, which is expected to kill 600 people. Two alternative options to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the options are as follows: If Option C is adopted, 400 people will die. If Option D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. Which of the two programs would you favour?"

For the most part, my placing the two risk tolerance questions together worked as expected, which is consistent with Tversky and Kahneman.[61] The responses are shown in Figure 9:

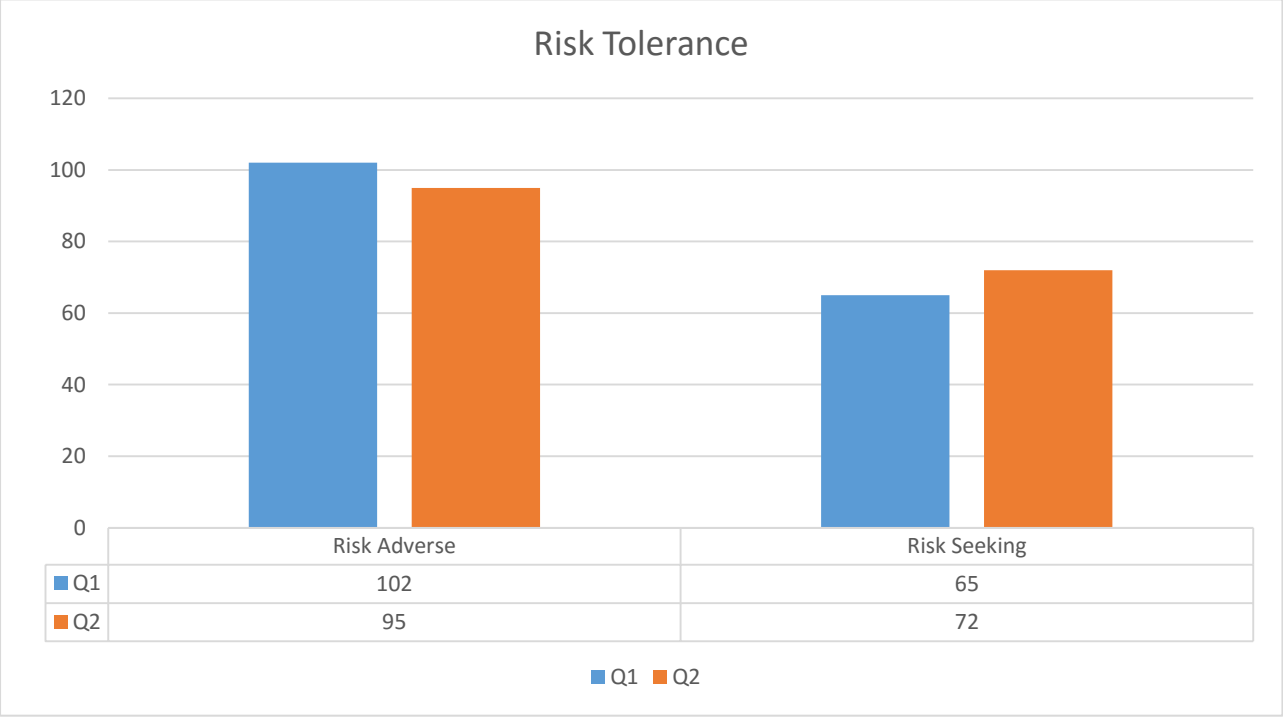

Figure 9 - Risk Tolerance

[61] Doing it in the described manner is a differing, and in many ways opposite result, than what Tversky and Kahneman were trying to achieve. They were trying to demonstrate inconsistency in thought due to risk aversion and how things are phrased. What I was trying to achieve was consistency of thought by placing the questions together.

As you can see in Figures 9 and 10, most people responded the same way to both questions. However, as you can see from the data, not all did. There was a net shift of seven from risk adverse to risk-seeking on the basis of the question phrasing. The breakdown was as follows:

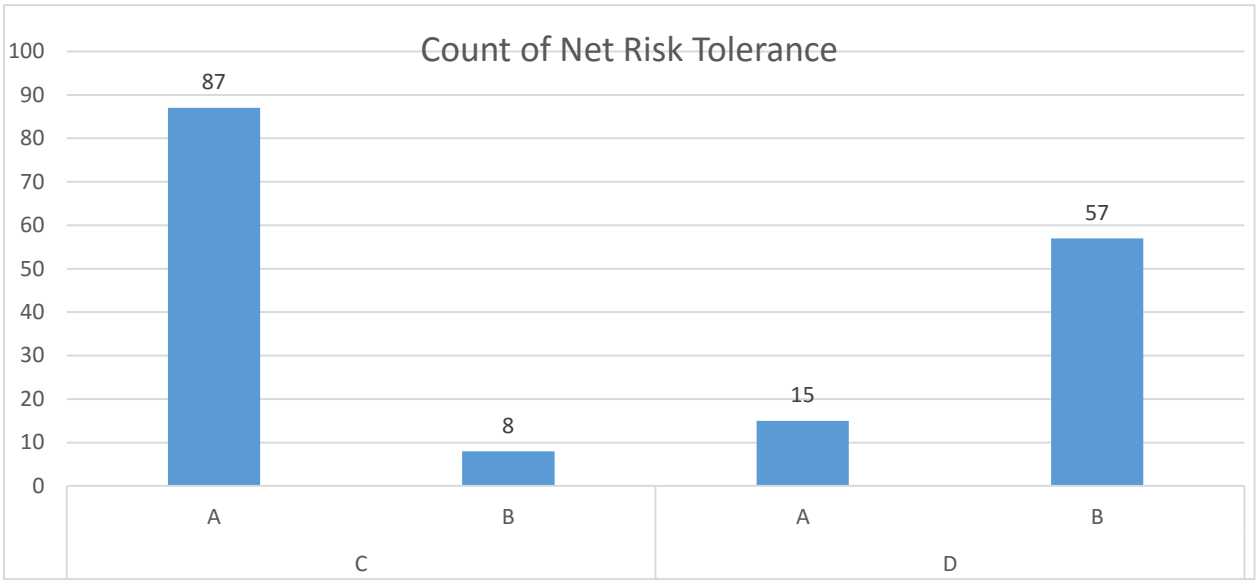

Figure 10 - Count of Net Risk Tolerance

87 people answered both A and C (which were the risk-adverse options), 57 people answered both B and D (the risk-seeking options), and 23 people were inconsistent. If you look at the Chi-Squared test for this (as depicted in Figure 11), you can see that it's not random.

	Risk Questions		
	C	D	
A	87 *58.02* (14.47)	15 *43.98* (19.09)	**102**
B	8 *36.98* (22.71)	57 *28.02* (29.96)	**65**
	95	**72**	**167**

χ^2 = 86.230, df = 1, χ^2/df = 86.23, P(χ^2 > 86.230) = 0.0000

expected values are displayed in *italics*
individual χ^2values are displayed in (parentheses)

Figure 11 - Chi-Square Analysis of Risk Tolerance Questions

It is fair to assert on the basis of this data that most people are basically risk-adverse (roughly 60%), however, that a sizeable minority are in fact basically risk seekers (roughly 40%).

Observed Games

In this study, there were two classes of games I recorded, primary games (of which I had 19), and secondary games (of which I had 20).[62]

A primary game in this study is one in which I recorded all data, rolls in their order in the game, all trades, took photographs of board states, etc. That is to say that other than recording the words that players used, I recorded everything necessary to recreate exactly what happened in each and every game. After each game, I had each player-participant fill out an end-of-game survey, which allowed me to inquire as to their strategic thinking during the game.[63]

A secondary game in this study utilised a board set-up from a primary game, as well as the player order and initial settlement placement, but in those games, only a net tally of rolls (how many 2s, 3s, etc) was kept, as well as who won each game was recorded.

This division was performed for two primary reasons: 1) Because I wanted to validate the results that I acquired as reasonable and truthful, which required a check against the data. 2) I could email out a picture of the board, and indicate player order, and get in return who won or lost, and the net tally of rolls. All primary games were administered by me, utilising a game recording method developed by me, secondary games were administered by a variety of people, which is why the data-recording requirements were kept at a minimum, since it obviated the need to train someone else in my recording method or in fact observe those games myself, which of course saved resources, primarily my time. It should be noted that in 13 of those 20 games (all 4-player), there was a different victor ludorum than occurred in the related primary games.[64] This was not even originally planned, but was implemented during the course of this study, once it became obvious to me that I was not going to get the 60-75 games I had planned on. Of the secondary games conducted for this study, the position winning percentages well matched what Szita *et al* found in their study with respect to random agents. Player 1 won 8 times, Player 4 won 6 times, and Players 2 and 3 won 3 times each. This is even a more pronounced result than

[62] I in fact did have 20 primary games conducted, but I've lost the data that I recorded for one game, as such, it is being ignored in this count.

[63] Like with the general strategy surveys, the end of game surveys were also reviewed by UVA's Social and Behavioural Sciences IRB and approved.

[64] Victor ludorum = Winner(s) of the game(s). It works in singular and plural given how the Latin is conjugated with this phrase.

Szita *et al* found.[65]

Furthermore, it was observed over the course of the primary games that the "pure" strategies of wood-brick and ore-wheat were only 75-80% pure as implemented by players in the game. As been stated elsewhere, strategies that entirely ignore certain types of resources are theoretically feasible, but are practically infeasible insofar as you need a lot of other things to go right, and some of those are probabilistically very rare indeed. For instance, in all primary games observed, only in one of those games did the Longest Road trophy not change hands between players at least once, in most games (13), it changed hands at least twice. This is because someone attempting to get the longest road in the game is often noticed very quickly and other players will try to break up the road if able, or at least build their own developments in such a way to attempt boxing in such a player. A player attempting a development card focused strategy will bring other players to buy development cards in the attempt to stop the player from acquiring VP cards. This happened in half of all games in which a development card based strategy was attempted by at least one player. However, in only two games observed did the Largest Army trophy change hands at all. It should be noted that the same model that Richardson created to model arms acquisition by rivals can also be used in this context.

It should be noted that there is one notable imprecision in the rules. There is no specification regarding if ports can be utilised in the turn that someone settles on them. For purposes of conducting the primary games, in accordance with the policy on development cards, which can't be used in the turn purchased, I held to the standard of no using ports on the turn settled on, however, unless/until that issue were addressed by Mr Tauber, that will hamper modelling efforts.

Things Observed from the Data

As I expected, VP derived from structures form the core of players' total VP scores. As can be seen in Figure 12, there is a reasonably strong correlation between VP derived from structures built, and their total VP scores ($R^2 = 0.55$).

[65] Primary game results weren't conclusive, as I didn't have enough 3-player games nor 4-player games to be able to draw any conclusions in this matter. Arguably the secondary games weren't either, given that only 20 of them were held, however, that is something that should be investigated in further studies.

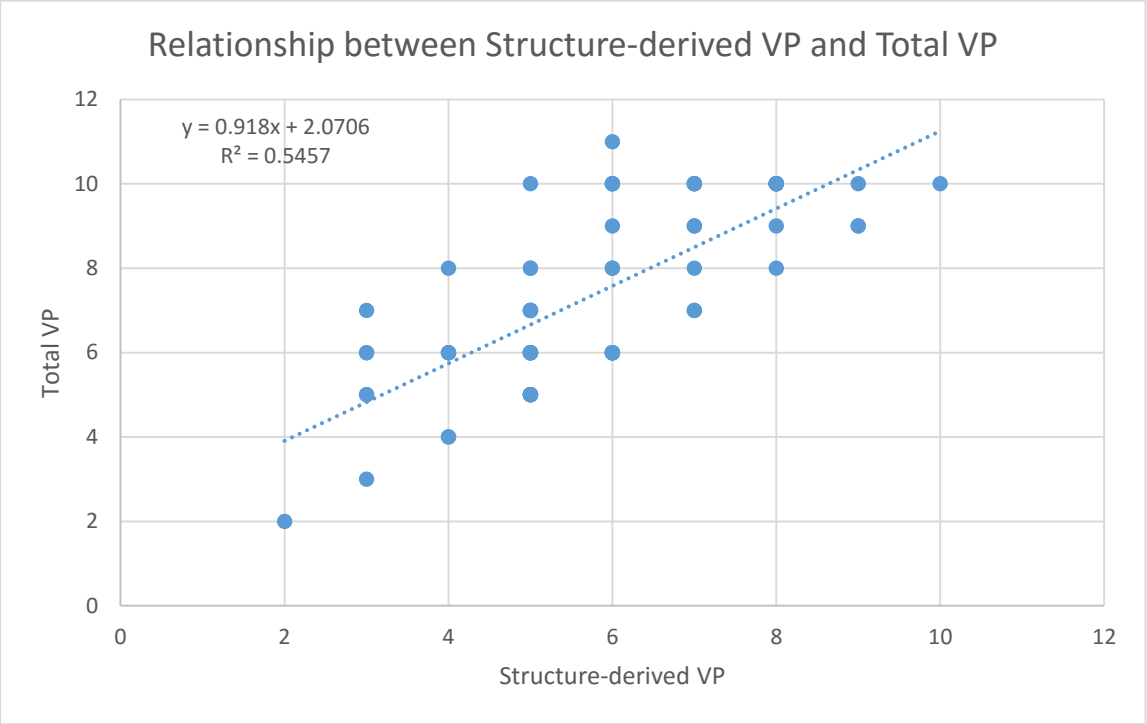

Figure 12 - Relationship between Structure-derived VP and Total VP

Given all the various elements of the game, the human factors, as well as the various ways to acquire VP, an R-squared at or near 0.6 is as realistically strong of a correlation as I think can be expected. What Figure 12 also demonstrates is that even though the entire triangle bound by (2,2), (2,11), and (10,10) is theoretically possible as game score results for players, that the actual shape of outcomes is a roughly triangular form that is a point at (2,2), with its "base" ranging from x = 5 to x = 10 where y = 10. As you can see from Figure 13, that is an even stronger correlation than the number of trophies to the number of total victory points, and note that every single time that both trophies were held by a single player, that player won (this occurred in 3 of the 19 primary games).

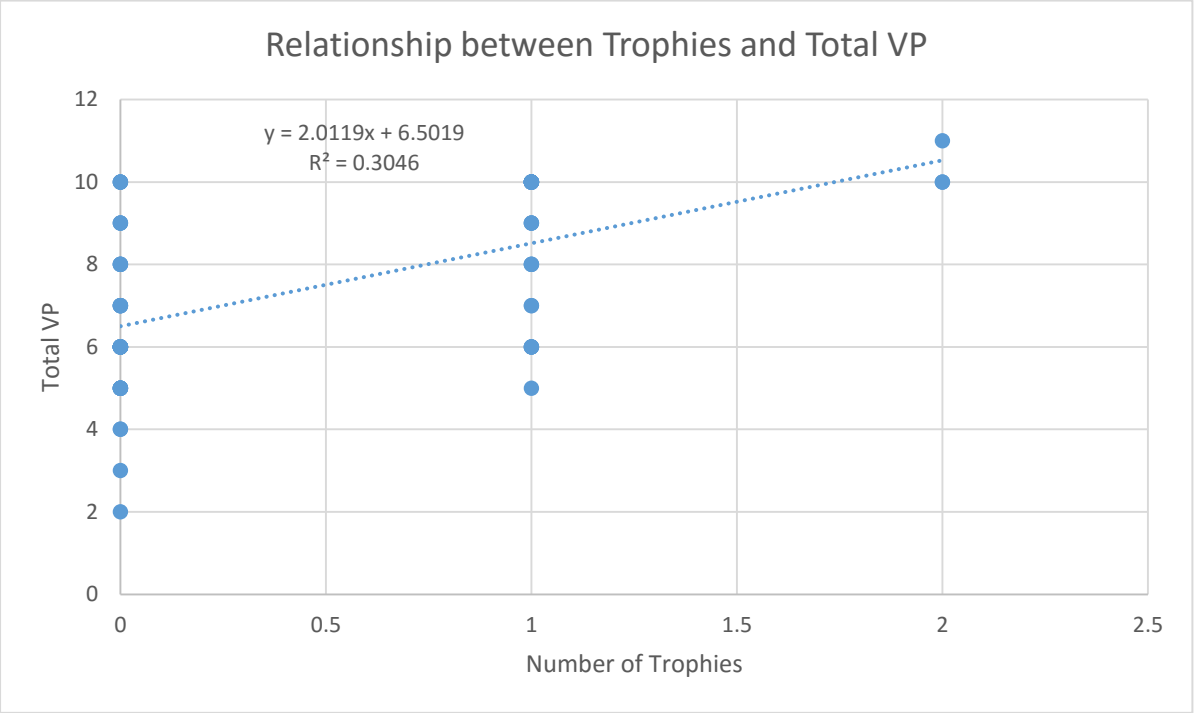

Figure 13 - Relationship between Trophies and Total VP

It should be noted that along with those three games in which one player had both trophies, in 12 of the 19 primary games, the victor ludorum had one trophy, and in 4 of the games, the victor had zero trophies.

No other two-variable linear correlation in the data collected from end-of-game surveys was of any value as the fits all were below $R^2 = 0.10$ save the correlation between a player's longest road and their total roads built (0.797), which is to be fair not very relevant to most discussions.

In order to assess how players attained victory, I created a table to see what winning players had, development-wise in the game. As should be expected, all 19 primary games each had a winner, but of course, they utilised various strategies to get there. Summary statistics are given in Table 2:

	City	Sett	Trophy	Knights	VP Cards	Action Cards	Dev Cards	Longest Road	Total Road
Min	1	0	0	0	0	0	0	2	4
25%	1.5	1	1	0.5	0	0	1.5	4	7
50%	3	2	1	3	1	0	4	7	9
75%	3.5	4	1	3	1.5	1	5.5	9	10
Max	4	5	2	5	2	3	9	11	12

Table 1 - Summary Statistics of Winning Players' Inventories

As you can tell, looking at the minima of all games, that is an infeasible pathway to victory, since that would only provide 2 VP with one structure, and there's no way in the game to only have one structure. In contrast, if you look at the maxima for all games, that is also an infeasible pathway to victory, given that 4 cities, 5 settlements, 2 trophies, and 2 VP cards would yield 19 VP, which is almost double the required number to win. The median case however is a highly feasible pathway to victory, even though it too would over-perform the requirement to attain 10 VP, as an inventory of 3 cities, 2 settlements, 1 trophy, and 1 VP card would yield 11 VP.[66] Even a longest road of 7 would be likely to take the longest road trophy most games, so again the median case over-performs requirements.

As I discuss elsewhere in this study, a strategy based on zero wood and brick, as well as one based on building or upgrading no settlements and cities are both realistically infeasible (although theoretically possible). As you can see in Figure 14, 18 of 19 primary game victors had at least 6 VP derived from structures, and all had at least 5 structural VP.

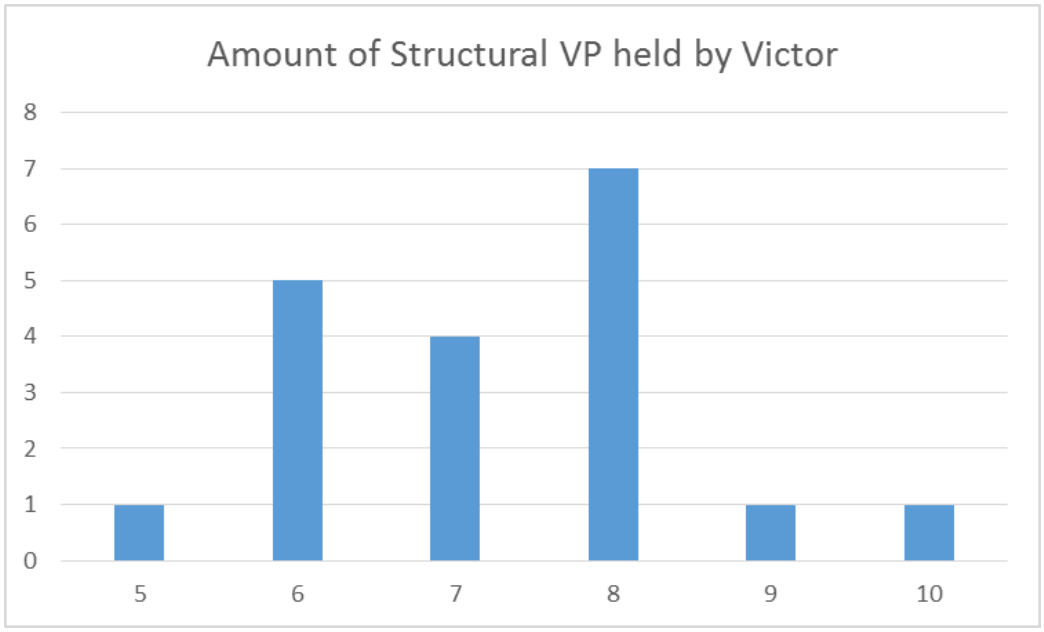

Figure 14 - Structural VP held by Victor

As a part of the end-of-game surveys, I asked all players to indicate which were their primary strategies during the game (and their secondary strategies if they had one). The results were certainly illuminating even as limited my observations are. So, utilising the strategies I specified

[66] It should be noted that the exact probability of getting 3 knights and1 VP card in exactly 4 development cards is 14.05%. This is of course a multinomial probability calculation, of this form: $\frac{4!}{3!1!0!}(\frac{14}{25})^3(\frac{5}{25})^1(\frac{6}{25})^0$. This is more likely than if you drew 4 cards, they'd all be knights (9.8%). However, it's less likely than if you drew 2 knights, 1 VP card, and 1 action card (18.06%). More than likely, in order to have 3 knights and a VP card, 5 or even 6 cards need to be bought. This would imply, that from a resource efficiency point of view, that the two each ore, wheat, and sheep expended on that fifth and sixth card would be better spent on either that fourth city (since that two sheep could more than likely be turned into an ore by trading with another player), or on another settlement (since only one wheat and one sheep go into a settlement, and the other wheat, sheep and 2 ore can be traded for brick and wood)

in Section 4, I numbered them in the order placed above: Wood-Brick = 1, Ore-Wheat = 2, …, Number Diversity = 6. As can be seen in Appendix 1, no strategy was strictly dominate, nor strictly dominated, all six strategies were used by victor ludorum. However, as you can see from Appendix 1, the primary strategy with the most wins was Ore-Wheat with 7, followed by Rare Resource with 5. Of the secondary strategies, the most dominate secondary strategy wasn't even using one with 8 not utilising one (4 of 7 of the ore-wheat users, as well as both of the Abundant Resource/Port users didn't utilise a secondary strategy), however, the most utilised secondary strategy was Abundant Resource/Port with 5.

Now it should be noted that an unlisted strategy in my original list was a Development Card focused strategy (which in my experience seemed derivative of the ore-wheat strategy, which is why I didn't list it amongst my list).[67] What I found was that 5 of 19 winners were utilising a development card focused strategy, including 3 of the 4 ore-wheat primary strategy adhering victors with no listed secondary strategy.

As a part of the survey, I attempted to ascertain when and how people trade in Settlers of Catan. So I asked a series of questions about trading preferences, and the data was revelatory to be sure.[68] As I expected, everyone, or effectively so, is willing to trade with their opponent(s) if both they and their opponent are under 7 VP (173/174 answered in the affirmative).[69] Most players answered in the negative to the question "Will you trade with your opponent if both you and your opponent are above 7 pts?" (108 no, 66 yes). One that was a bit surprising was how many are willing to trade if they and their opponent are ***at*** 7, almost 80% (139) answered in the affirmative, only 35 said no. Most people are willing to trade with their opponent if they're at 7 (96%/167), 8 (94.3%/164), and 9 (90.8%/158), however, as should be expected, there's more reticence to trade as their point total goes up. I then asked three questions attempting to ascertain if people are willing to trade if their opponent is at 7, 8, and 9 VP. This is where things get rather interesting. For the question: "Will you trade with your opponent if your opponent is at 7 pts?", which I expected to be answered mostly no, was in fact in total a strong yes, with almost 76% (132) answering in the affirmative. As I expected, most (73.6%/128) answered no to "Will you trade with your opponent if your opponent is at 8 pts?", and an even higher

[67] As it happens, one of the papers I reviewed for this study (Saleem and Natiq) did in fact have it listed amongst its strategy types, however, by the time I saw that, I had already sent out my Settlers strategic survey, so there was no way to change it.
[68] All 174 responses received were used for this portion.
[69] To be fair, I expected this to be answered by everyone in the affirmative, so the fact that even one person answered in the negative is something potentially to be explored.

percentage (90.2%/157) answered no to "Will you trade with your opponent if your opponent is at 9 pts?" On the basis of my previous experience, I was expecting that last question to be answered uniformly in the negative. That generosity may be against one's self-interest, but it makes sense up to a point. Where things start to defy logical preference is whilst better than 60% (105) won't trade with their opponent if the opponent is at 8 and they're at 7, and almost 76% (132) won't trade with their opponent if the opponent is at 8 at they're at 6 (so 42 will), one more person will trade with their opponent if the opponent is at 8 and they're at 5. This makes no sense at all, a player who is further behind should be less willing to trade the further they are behind. This behaviour is also extant when the opponent is at 9 VP. So, when the opponent is at 9 and they're at 8, 144 (82.8%) said they wouldn't trade and 30 said they would. When the opponent is at 9 and they're at 7, 154 (88.5%) said no and 20 (11.5%) said yes. And when the opponent is at 9 and they're at 6, 153 (87.9%) said no and 21 (12.1%) said yes. That is to say that players are taking the risk of trading when they feel they "know" they're going to lose, but aren't willing to trade when it would, in fact, permit a higher chance of winning. This behaviour is not just sub-optimal by acting against interest, it doesn't even adhere to being consistent and violate the postulates of subjective expected utility.[70]

This result goes directly against my own experience in this matter, in which players tend to be reticent about trading with others at/above 7 VP. This could be explained by the type of people I'm friends with, of the 167 surveys which I analysed for this purpose, 33 were taken by friends of mine, the median for my friends was 64 (which is 4 points into the High Mach category) and 55% of my friends were considered High Mach, whereas there was a median of 59 for all others (which is 2 points less than the High Mach threshold), with 45% belonging to the High Mach category. It seems reasonable to conclude that, caeteris paribus, my friends are more Machiavellian than the population at large, which would be a reasonable explanation as to why my expectations in some areas were broken by the data I collected in my surveys and observed watching primary games.

I also attempted to ascertain how many people pulled off what in game theory terms is I think the ultimate defection, that is how often players would trade away all of a resource and then use a Monopoly card to recall that which they just traded away. Now, according to game theory, someone who acts in such a way, you would expect many if not most people to never trade with

[70] That is to say such behaviour violates the Ellsberg Paradox. See *Risk, Ambiguity, and Decision* for more detail.

that person again, or at least wait a good long while before doing so. The data I collected indicates that people are way more forgiving than this. According to the survey data, 52.3% (n = 91) of participants said yes to the question: "Will you trade with your opponent if your opponent ever trades a particular resource away, and then uses a monopoly card to call that resource back?" and another 43.7% (76) said they would after a few turns. Only 4% said they'd never trade with that opponent again. Related to that, I asked participants if they'd ever done so themselves. 16.1% (28) said that they did often, 27% (47) said that they did occasionally, and 36.8% (64) said that they'd done so at least once in their playing career. Only 20.1% (35) said that they'd never done so in their Settlers playing careers.

According to the rules of Settlers of Catan, only the player whose turn it is may initiate trades. However, given human interaction, this isn't as simple as it seems. The language issue is a big part of that. There are many ways to express the desire to trade, both the person whose turn it is, as well as other players who are, or might be, willing to trade. This rule assumes that the only way that trades occur is that the player whose turn it is looks at one other player, and makes a specific offer. In my supervision of the primary games for this study, and what I observed, this was rarely what happened. Sometimes, the player whose turn it was would ask a question along the lines of "Who has <resource>?" Sometimes, it is along the lines of "I have <x>, does anyone want it?" There's also the strategic element of attempting to ascertain what other players have before deciding to offer a trade. Also, something that needs to be kept in mind is what a friend of mine (who in fact was one of the player-participants for this study) said: "That while there are rules this is still a leisure activity".[71] The fact that this game is seen as a leisure activity is potentially why also that results from the survey and the games deviate from theory.

Also, with respect to trading, between ¼ and $^2/_7$ths of the trades I observed in the primary games were "defensive" in nature, that is to say if a player was well over 7 cards, or they had a lot of the same type, they would initiate trades with other players or even the bank to drive down their resource card totals to less than or equal to 7.

As was expected, the aleatory element played a big factor in results. Whilst the distribution for the rolls of all primary games conducted was roughly equivalent to where it ideally should have

[71] Josiah S Byers. Along with being a player-participant, he holds a masters in international relations from the University of Washington and was valuable as a resource to check things against, both for his experience in Settlers of Catan as well as input regarding applications of strategic thought.

been, the operative word here is "roughly". Figure 15 is a histogram comparing the ideal, discrete ideal, and sums of rolls over all primary games conducted.

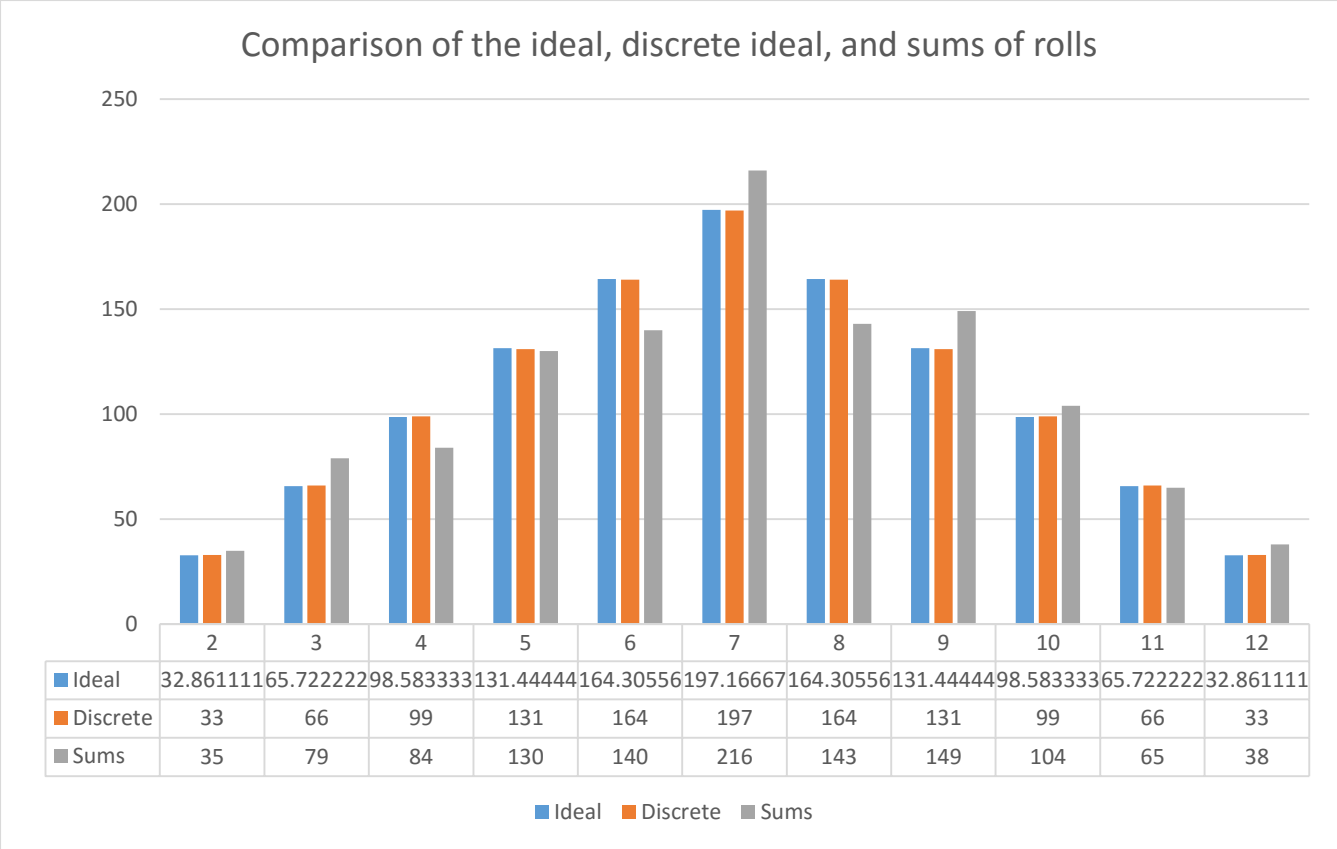

	2	3	4	5	6	7	8	9	10	11	12
Ideal	32.86111	65.72222	98.58333	131.4444	164.3056	197.1667	164.3056	131.4444	98.58333	65.72222	32.86111
Discrete	33	66	99	131	164	197	164	131	99	66	33
Sums	35	79	84	130	140	216	143	149	104	65	38

Figure 15 - Comparison of the ideal, discrete ideal, and the sums of rolls

The number of all dice rolls was 1183 of all primary games conducted for this study (see Appendix 1, where a table of dice rolls for primary games is located). Given the dice distribution is out of 36, a multiplication factor of 32.86111 needed to be applied to scale the dice distribution up to the same scale as the actual number of dice rolled. Summing the idealised dice rolls also equalled 1183, as did the rounding from rational to natural numbers. The sum of squared errors of the discrete ideal from the actual ideal was 1.149691358 (which is the theoretical optimal SSE in this case). The actual dice rolls were not nearly so precise in their estimates. The SSE of the actual dice rolls was 2159.483025, and the percentage off individual dice were off from their ideal ranged from -14.8% to 20.2% (Figure 16).

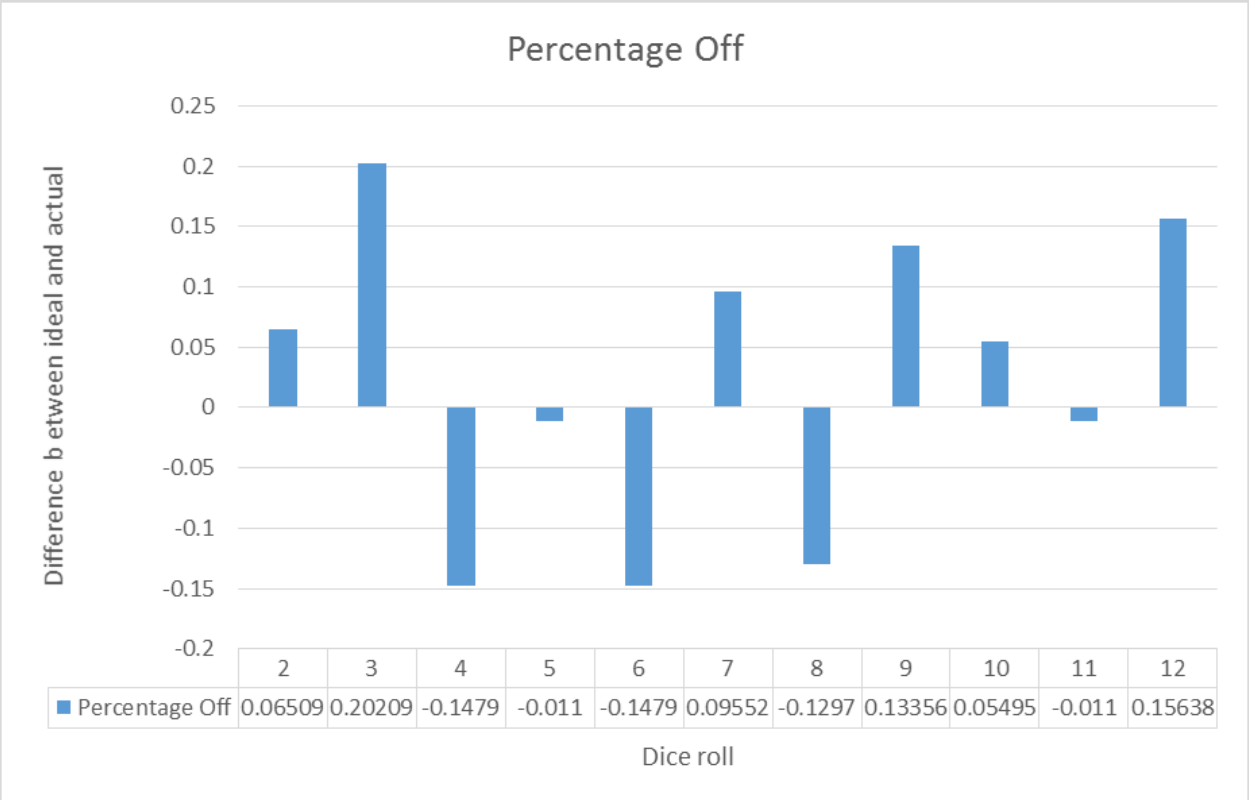

Figure 16 - Rolls and their percentage off the ideal

However, two things of note go against this apparently major discrepancy.

First, if you sum all the actual dice rolls and their values ($2 * 35 + \cdots + 12 * 38$), that total is 8341. For the discrete ideal, that total is 8281. The difference between those two numbers is just over 7 per mille, which is certainly reasonably close (as the difference is under 10 per mille (or 1 per cent)).[72]

Second, checking the game dice rolls against simulations run in @Risk, I have to say that given a total run of 1183 throws, that the actual dice rolls acted as probability theory would dictate that they should. And as rough as this histogram may appear, most of the dice roll histograms for individual games look nothing like they ideally should, all of which I was able to simulate using @Risk.

In short, the strong law of large numbers works, but it may take a while to get there. More importantly, in an individual game, the player may not be able to rely on what the dice distribution says they should get on average.

[72] To be precise, it is: 0.007245502

Practically speaking, of course, this means that you can do everything "right" strategically and still lose on the basis of how the dice are rolled in any particular game. This is the element of chance that planning and strategy can't entirely overcome.

For instance, in game 1 of this study, the initial layout is in Figure 17:

Figure 17 - Game 1 initial placement

It is fair to assert that the red player had the best initial layout of the three players. The player is on both a 6, and an 8, as well as a 5 and a 9. So, the entire centre mass of the dice distribution is covered. However, red didn't win this game, red in fact place last as can be seen in Figure 18:

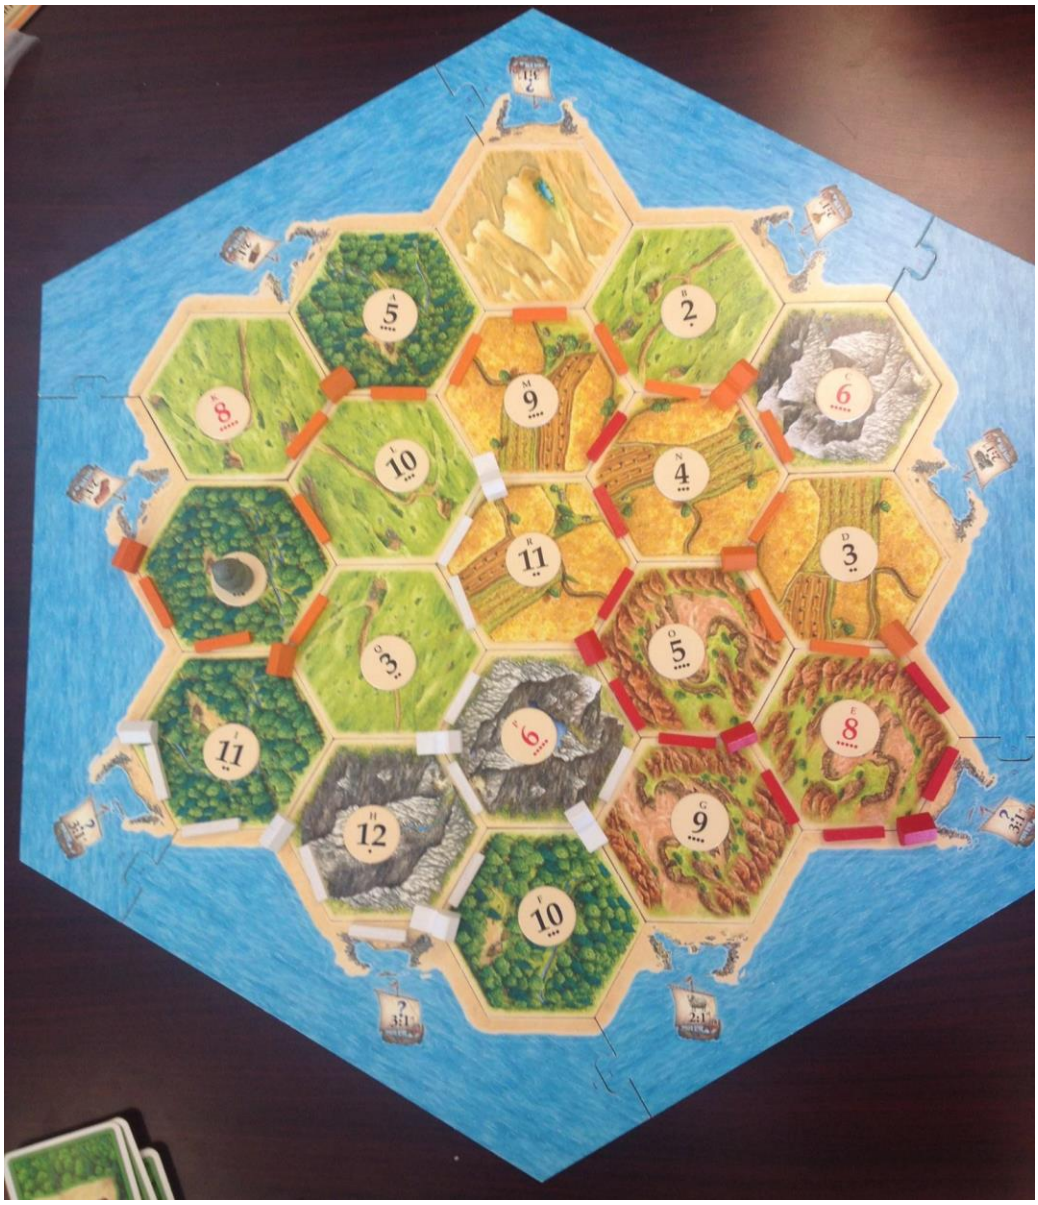

Figure 18 - Game 1 end of game state

Not only did the red player not win, he had the lowest score amongst the three players.[73] What happened was a game roll distribution that looked like this:

[73] It should be noted that the author of this study was the red player in this game, and only when someone specific is referenced will a gender specific pronoun be used, otherwise, the general neutral "they" would have been used over the male specific "he".

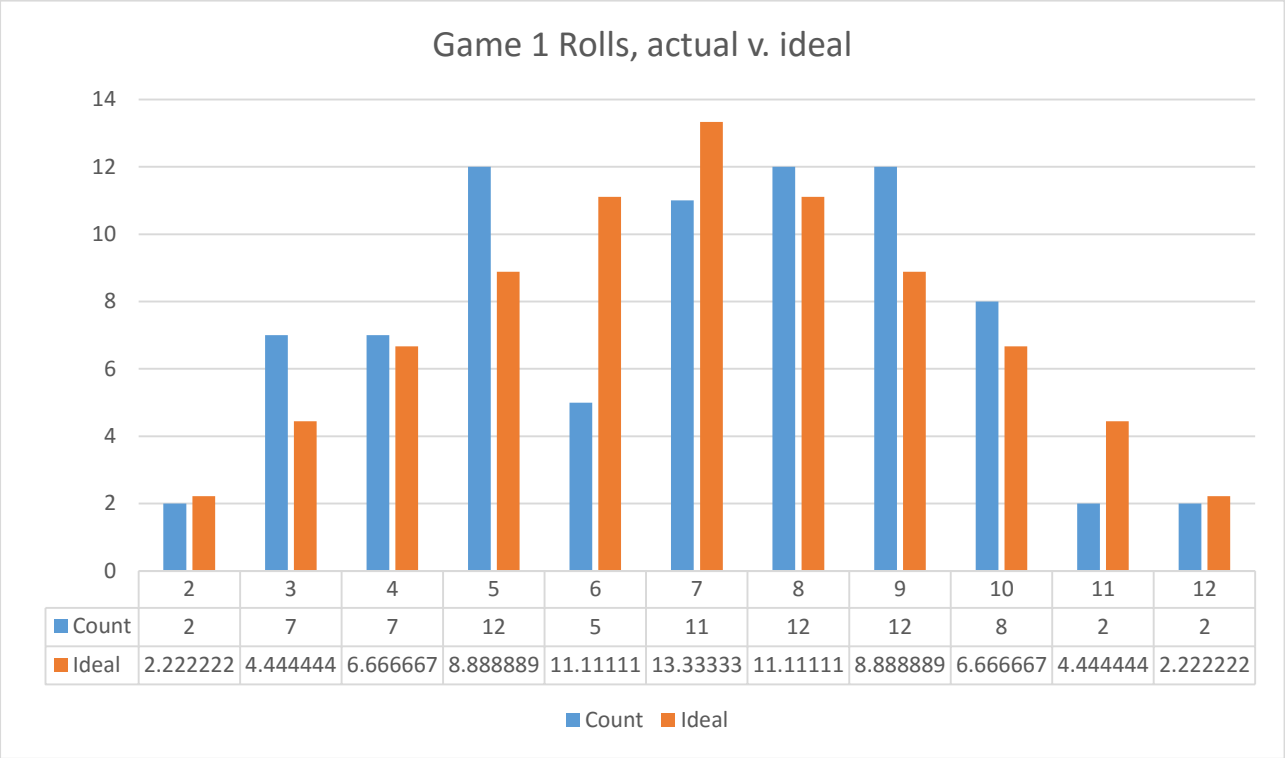

	2	3	4	5	6	7	8	9	10	11	12
Count	2	7	7	12	5	11	12	12	8	2	2
Ideal	2.222222	4.444444	6.666667	8.888889	11.11111	13.33333	11.11111	8.888889	6.666667	4.444444	2.222222

Figure 19 - Game 1 Rolls, actual v. ideal.

As can be seen in Figure 19, the rolls went in the favour of the orange and white players, as every time a 5 was rolled, that was a road segment equivalent for orange (as orange was getting one wood and one brick every time a 5 was rolled in the game), 9 and 10, which were overly productive helped white, and 6's which should have ideally happened 11 times in this game only happened 5 times, which hampered the red player's efforts. Red initially earned the longest road trophy, but it was taken by orange in late game. As I discuss elsewhere, the event of the longest road trophy being stripped from the player who first earned it happened in every primary game observed for this study but one. It should also be noted that it was the white player who won this game, with an ore-wheat based strategy.

A discussion must also take place as to how game data from primary games was recorded. As has been mentioned elsewhere, a consistent recording system for detailing what took place in the Settlers games conducted for this study had to be developed, as no such system previously existed (unlike Chess or Go, which have standard notations for the game).

The left-hand side of the sheet has 5 columns, the leftmost to store each roll as it happens in the game, the next 3-4 columns record what resources each player got each time the dice were rolled, and the right most column (which took up the right-hand-side of the sheet) stores all

actions taken in that turn.

So, for instance, assume the game board in Figure 20 is in play:

Figure 20 - Game 18 Board, end of round 9

So, during the beginning of round 10, Player 1 (who is blue), rolls a 4, and then player 2 (white) rolled a 6. That round would get recorded as follows:

Roll	Blue	White	Orange	Red	Actions
4	0	2S	0	2W	1x2(B-W); 1x(2S-O); 1City, pts 6
6	O	2O	2W	2O	2x1(W,S-O); 2x3(W,S-O); 2x(4O-2W); 2(Rd Bld); 2(Rd*2); 2(Lg Rd); pts 10

Table 2 - Demonstration of Settlers of Catan game recording system.

Deciphered, this turn, blue rolled a 4, the blue and orange players received no resources that turn, white received 2 sheep, and red received two wheat. The blue player then commenced to trade with white one brick for one wheat, blue then proceeded to trade two sheep with the bank for an ore, upgraded a settlement to a city, and then had 6 VP on the board. White then rolled a 6, blue collected an ore, white 2 ore, orange 2 wheat, and red 2 ore. White then committed the following actions: He traded one wheat and one sheep to blue for an ore. White made the same

trade with orange. White then traded with the bank 4 ore for 2 wheat (using his ore port). White then played a road building card, which allowed him to build 2 road segments, which allowed white to claim the longest road trophy, which in turn pushed white to 10 victory points, winning this game.[74]

In this system, the when recording resources, the following notation was used: O = ore, W = wheat, Wo = wood, B = brick, and S = sheep. Given the similarity in orthography between a standard nought and the letter O, a slashed zero was used to record when no resources were received by a player any turn. A slashed zero with a subscript R was used to indicate when the robber prevented a player from receiving any resources that turn, whereas an (R) was used to indicate that resources were lost in a turn because of the robber, but other resources were still attained.

As I mentioned above, the building of cities was recorded as (player number)City, the same was done for settlements, using the abbreviation "Sett", so if for instance blue built a settlement, it was encoded as: 1Sett. Roads that were built were encoded as (player number)Rd, with (#)Rd*n to record multiple roads built in a single turn.

When development cards were purchased, they were recorded as (player)Dev. When action cards were played, I would abbreviate the action card, as (player number)(action card). Monopoly cards played were encoded as (player)(Mono-__) where the blank was filled in with what resource was being monopolised by the card. If white played a monopoly card and called for brick, that would be encoded as: 2(Mono-B). Year of Plenty cards are abbreviated as YOP, and if they were played, they'd be encoded as (player)YOP(the resources obtained). The road builder card was encoded as Rd Bld, so if white played the road builder card, it would be recorded as 2(Rd Bld). Trophies, when earned, were encoded as (player)(trophy), where the largest army trophy was abbreviated (Lg Amy) and longest road was abbreviated (Lg Rd). When the robber was moved to another hex from the one it was on, it would be recorded as: Rob(tile letter) <- (player number) to note that the robber was moved to the particular tile marked by that letter, and the player number indicates the player that was robbed. When a player used a knight to move the robber, it was recorded as (player playing knight)Kt(Rob(tile letter) <- (player number)), so for instance if the blue player in this game used a knight to rob from the

[74] In game 18, white was played by Sameer Kanal, utilising the strategy that I set for him (he was running late, so whilst Sameer played white for that game, I acted for him in laying his initial settlements).

orange player from the hex with tile M (the grain hex with the 9), that would be encoded: 1Kt(Rob(M)<-(3)). If no actions were taken in a turn, the action column was left blank.[75,76]

See appendix 2 for the example of a full game that was encoded using the system I developed. The example is Game 18, since it was the second shortest game of the primary games recorded, and had enough actions taken by players that it exhibits the recording system as applied to a game.

One of the things I noticed is that 3-player games seemed to take a lot longer than 4-player games. The box plot of the two groups in Figure 21 I think is telling:[77]

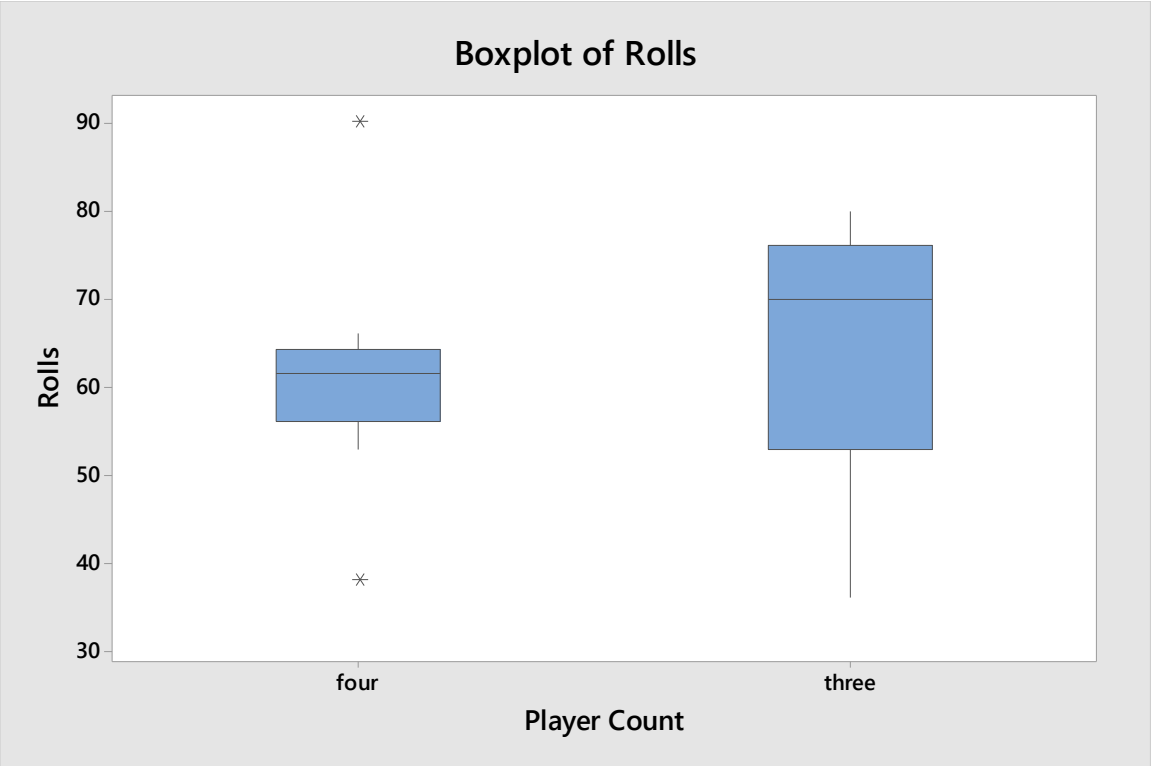

Figure 21 - Boxplot of Rolls of 3-player and 4-player games

When discussing something I noticed with one of the experienced players I consulted as a reference as to why 3-player games seemed to take longer than 4-player games and on the basis of the evidence I have, that was the case (as can be seen in the boxplot above), although it

[75] Player 4's turns were bolded to record the fact that after every round, a picture was taken as reference to the game state after that round. When recording those games, I used a pad of graph paper and a pen, I used a hashmark after the line of Player 4's roll to indicate the rounds.

[76] Each of the number discs are also uniquely lettered from A – R inclusive, so when encoding robber placement, it's an identifier.

[77] 3-player: Q1: 53 Q2: 70, Q3: 76, Whiskers Range [36,80]
4-player: Q1: 56 Q2: 61.5 Q3: 64.25, Whiskers Range [53,66]
Performing an ANOVA using Minitab, the results it gave were as follows, with a pooled SD of 13.2761:

Count	N	Mean	StDev	95% CI
four	12	61.08	11.75	(53.00, 69.17)
three	7	64.29	15.69	(53.70, 74.87)

should be noted the data isn't statistically robust given it's only 19 games worth of data.

He remarked that the explanation of why that was the case was probably because of the fact that the economy of three-player games was much reduced in comparison to 4-player games, because fewer resources went out with each turn.[78] I can think of no better explanation and that one made total sense to me. Admittedly, it was something that I hadn't thought of, nor did any of the literature I reviewed mention the phenomenon. So, credit for that thought is his alone.

Initial Heuristic-based trading model

The initial trading algorithm I developed heuristically is as follows:

Are both players under 7 pts

if yes, trade probability 2/3

assume initial offer at ½ probability (assuming non-sheep being offered), 1/3 if sheep being offered

if rejected, assume a second card (of same or differing type) will be offered, 2/5, any of non-sheep type, 3/5 sheep will be offered as the second card

Is one player at 7 pts

if yes, is second at 6 pts

if yes, the probability of trade 2/5ths, this holds only if a player at 6 pts can make a VP gain of 1+ immediately that turn.

if not, probability of trade = 0

Is one player 7pts < x

if yes, the probability of trade = 0

Given my previous experiences with Settlers of Catan, this set of trading rules seemed reasonable, at least initially. Sheep are generally speaking the least valuable commodity. This viewpoint is so pervasive that there are both memes on the internet, as well as a t-shirt that is made that has written on the obverse: "No one wants your f*ing sheep".[79]

To elucidate this point, let me lay out an example. Assume that the 10 VP a player acquires to win is composed of the following developments: 3 cities, 2 settlements, 7 road segments (5 purchased), and 4 development cards, of which all 4, or at least 3, are knights, and is awarded the Largest Army trophy.

[78] Brian Strom

[79] Buzzfeed references that in a post entitled "19 Pictures That Are Too Real For People Who Play Settlers Of Catan" which was published in 2015. That same list also depicts the shirt mentioned. A version of that shirt can be found at: https://www.etsy.com/listing/250424639/nobody-wants-your-fing-sheep-t-shirt

This would require, over the course of the game to acquire the following resources: 8 brick, 8 wood, 13 ore, 13 wheat, and 7 sheep. This requires 49 total resource cards, of which the 7 sheep are of course 1/7th of the total, and are the least required of the 5 resource types. And of course, this pathway is more heavily dependent on sheep than is the norm.[80]

A more common pathway to victory would ignore development cards totally. An example of such a pathway would be: 3 cities, 2 settlements, 8 road segments (6 purchased), no development cards, and holds the Longest Road trophy. Such a pathway would require 9 each brick, wood, ore, and wheat, and only 3 sheep. So 39 total resource cards needed for all development, and only 1/13th of the required are sheep.[81] A pathway that is similar to that one, but has one less road segment (and no longest road trophy) and another city would have the following resource requirements: 12 ore, 12 wheat, 9 brick, 9 wood, and 3 sheep. This case would require 45 total resource cards, and still only 3 sheep. That is, of course, 1/15th of the total.

The reason I went through those examples is that whilst the requirements for sheep are often minimal in Settlers games, the production of sheep often is not. As I demonstrated previously, 2/9ths of the resource generating hexes are sheep as are 2/9ths of the probability mass dictated by the enumerated discs. However, given that sheep demands are usually much less than their production, there is usually a glut of sheep on the market relative to their demand, which of course drives down their value.

More data are needed to refine this trading algorithm, but at least with initial data acquired for this study, it seems to hold well against the data and can be conditionally recommended for use until it is refined. From the data collected in primary games for this study, it is more conservative in trading than was observed, but is consistent with those results.

7) **What Does This Data Indicate**

As a part of this study, I performed various forms of analysis. One was the utilisation of a classification learner to assess whether or not the demographic and personality criteria I picked for the strategy survey was, in fact, useful for explaining the strategic competence of people.

The criteria used for that initial classification learner are as follows:

[80] This pathway is of course the median path as suggested by the 19 primary games conducted for this study.
[81] As it happens, this pathway is also the one that requires the fewest total resources of the various configurations available.

- Gender
- Mach-IV test score.[82]
- Years of formal education
- Risk tolerance.
- Merging years played and games played into an engineered experience score:

  Exp = Ceiling((years played)*ln(games played)).

  This was needed because the solicitation for the number of games played were (often) ranged estimates, so, treating that value with the natural logarithm function minimises the exaggeration, and without utilising a ceiling function, there is an implication of greater precision than in fact really exists.[83]

The summary statistics for the MACH-IV scores amongst survey participants is in Table 3

	Number	Mean	Mode	Q0	Q1	Q2	Q3	Q4	Low Mach	High Mach
Values	167	59.68712	64	2	53	60	67	94	0.544919	0.4550

Table 3 - MACH-IV scores of the survey takers

Using the quartile values of the set of people surveyed, I was able to create categorical bins that map to the Christie and Geis breakdown for Low Mach v High Mach (since quartiles 1 and 2 correspond to Low Mach, quartiles 3 and 4 to High Mach). As you can see from the table above, 54.5% of survey participants are Low Mach individuals, and 45.5% are High Mach individuals. I am not sure if there is a degree of selection bias here, as data doesn't exist to answer whether or not Settlers players are less or more Machiavellian than the general population, however, this is still instructive.

On the other hand, gender does have a relationship with MACH-IV score, and that does have enough entries to make a statistically valid claim. Summary statistics are in Table 4:

	Q1	Q2	Q3	IQ Range	Lower Whisker	Upper Whisker
Male	51	61	69	16	34	92
Female	50	56	65	15	41	82

Table 4 - 25%/75% breakdown

As can be seen, men are more Machiavellian than women are, at least in this sample. If this implies anything, it's that women, caeteris paribus, are nicer people than men are, which is why

[82] As I explain elsewhere, Mach-IV is a test of the Machiavellianism trait (a score is assigned from 1-100, x>60 places you in "High Mach", x:[1,60] places you in "low Mach").

[83] Lewis Fry Richardson used a similar system in his conflict modelling system, given that the estimates for wounded and killed were sometimes very rough at best. That is discussed in section 5.

they are on the net less ruthless.[84] Looking at the distribution in boxplot (Figure 22) is no less instructive. What the boxplot demonstrates is that not just are men, caeteris paribus, more Machiavellian than women, but that men have more variance in this regard than women.[85]

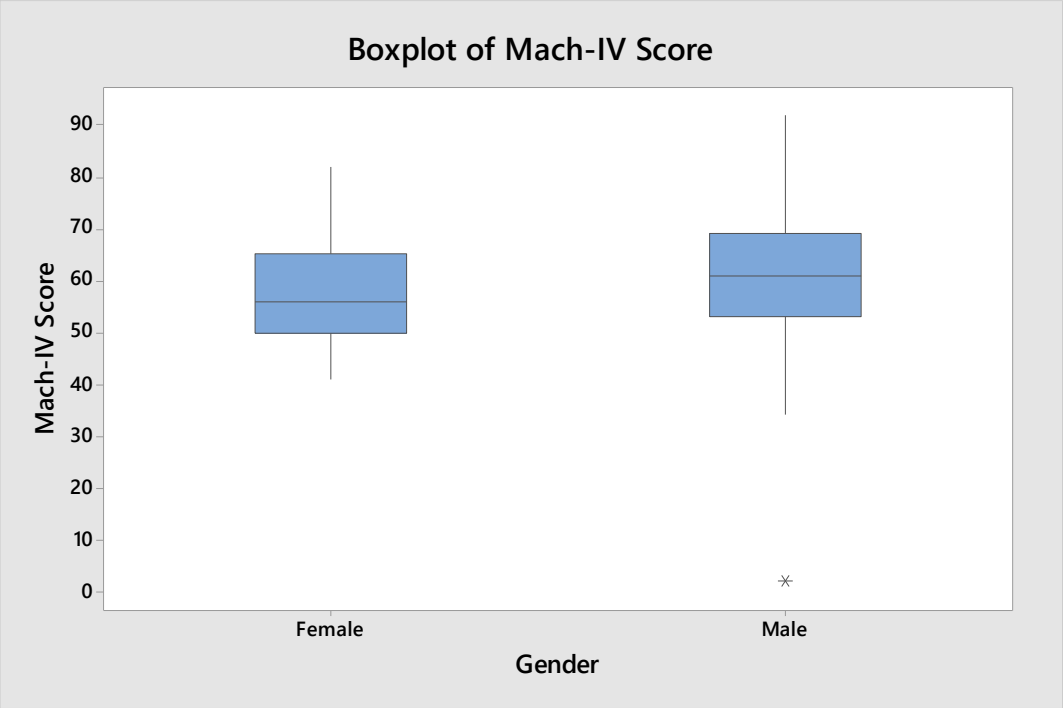

Figure 22 - Boxplot of Mach-IV scores, broken down by gender

[84] All things being equal.

[85] Utilising Minitab to conduct a one-way ANOVA, it generated the following (with a pooled SD of 11.5664):

Gender	N	Mean	StDev	95% CI
Female	50	57.58	9.45	(54.35, 60.81)
Male	113	60.62	12.38	(58.47, 62.77)

This is informative, but not conclusive. A larger sample pool is needed to confirm or disprove.

Risk tolerance comes from the questions I asked in the survey based on Tversky and Kahneman, of which I classified subjects into three groups, A = Risk-adverse to those willing to choose the zero variance option in both questions (guaranteeing that 200 lives are saved each time, but of course also guaranteeing that 400 lives are lost each time), B = middle for those who chose the high variance option for one of the two questions, and C = Risk-seeking being those who were willing to choose the high variance option both times (that is to say that they were willing to accept the potential for saving no one with a 1/3rd chance of saving everyone).

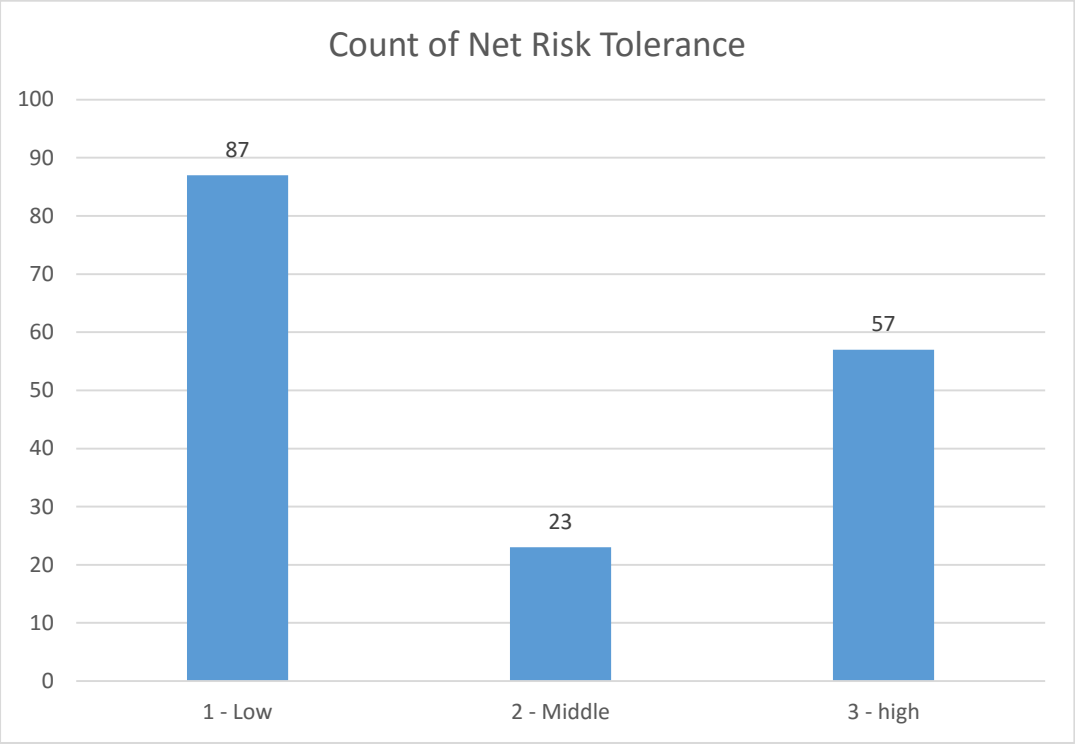

Figure 23 - Net Risk Tolerance Count

The middle case wasn't intentional, however, I had to do something with the participants who conditionally risk seeking, so it made sense to treat them as moderately risk-seeking.

As for the experience score, it ranged from 0 to 125, with a vast majority of the experience values skewed to the left (as depicted in Figures 12 and 13). The summary statistics for experience in Settlers of Catan are as follows:
Min: 0; Quartile 1: 7; Quartile 2: 14; Quartile 3: 32.5; Quartile 4: 125; IQ Range: 26

Below is the summary report of the experience score data:

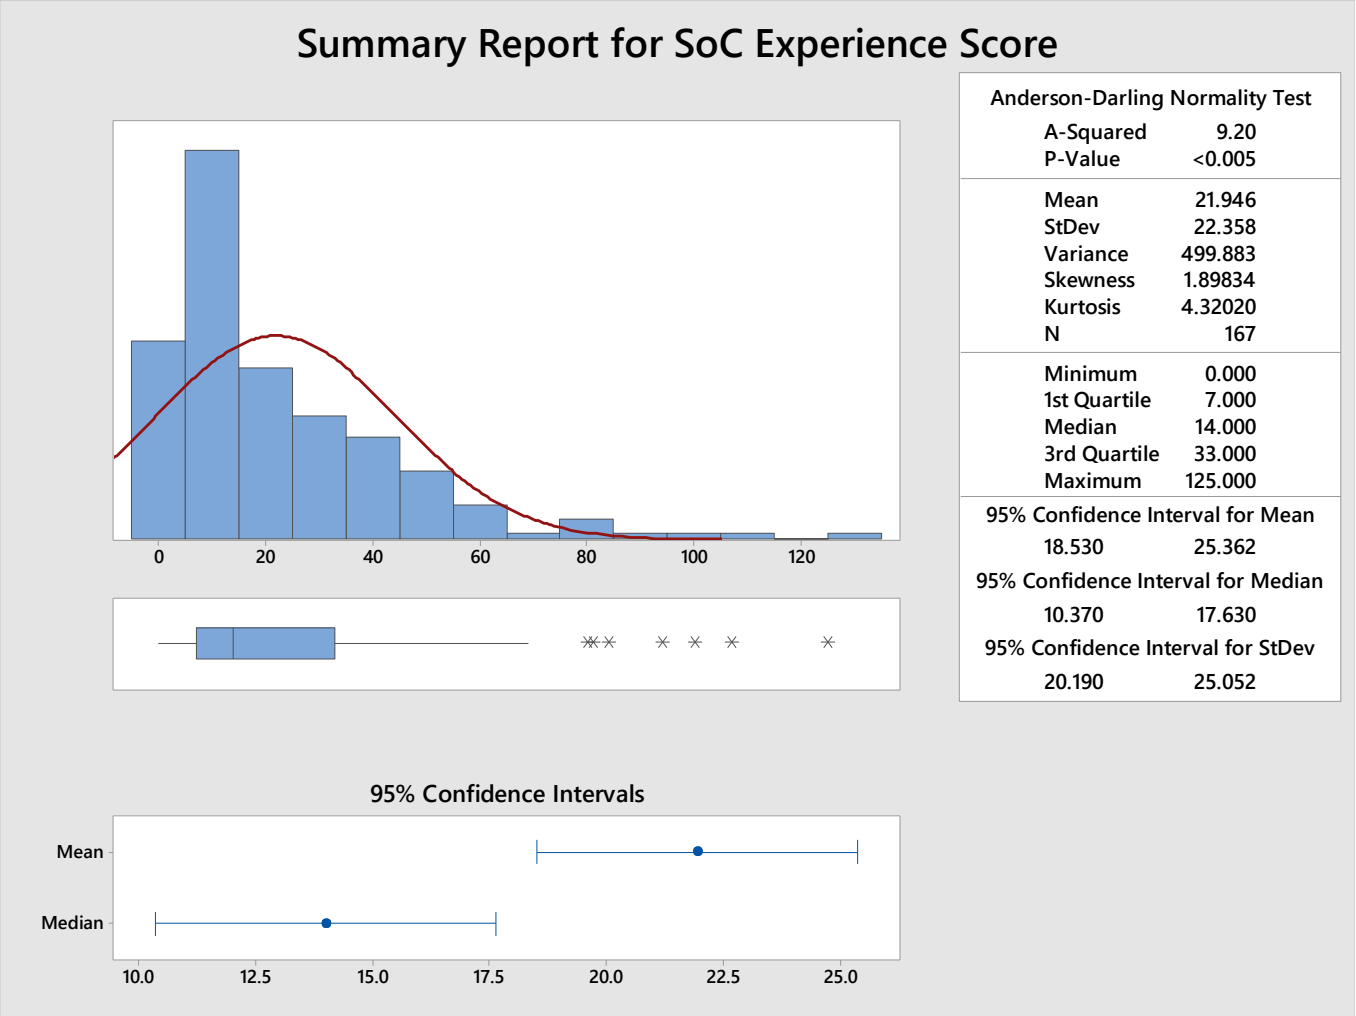

Figure 24 - Summary Report for Settlers of Catan Experience Scores

Sufficed to say, most players of Settlers (at least those that were surveyed for this study) are less experienced than more, and if you exclude the most inexperienced players, this distribution looks quite exponential in nature. It should be noted that although the primary author of this study was in the upper decile of experience, I did not possess the highest experience score (I was second at 108). As can be seen from figure 13, a second-order polynomial fitted to the experience score data has an R^2 of 0.9005. A fourth-order and sixth-order polynomial are also provided for contrast, but I believe those to be over-fitting the data and shouldn't be considered. This plot I think plainly indicates the fact that the rank-ordered experience data is well-ordered and demonstrative of the usefulness of the experience score formula I created.

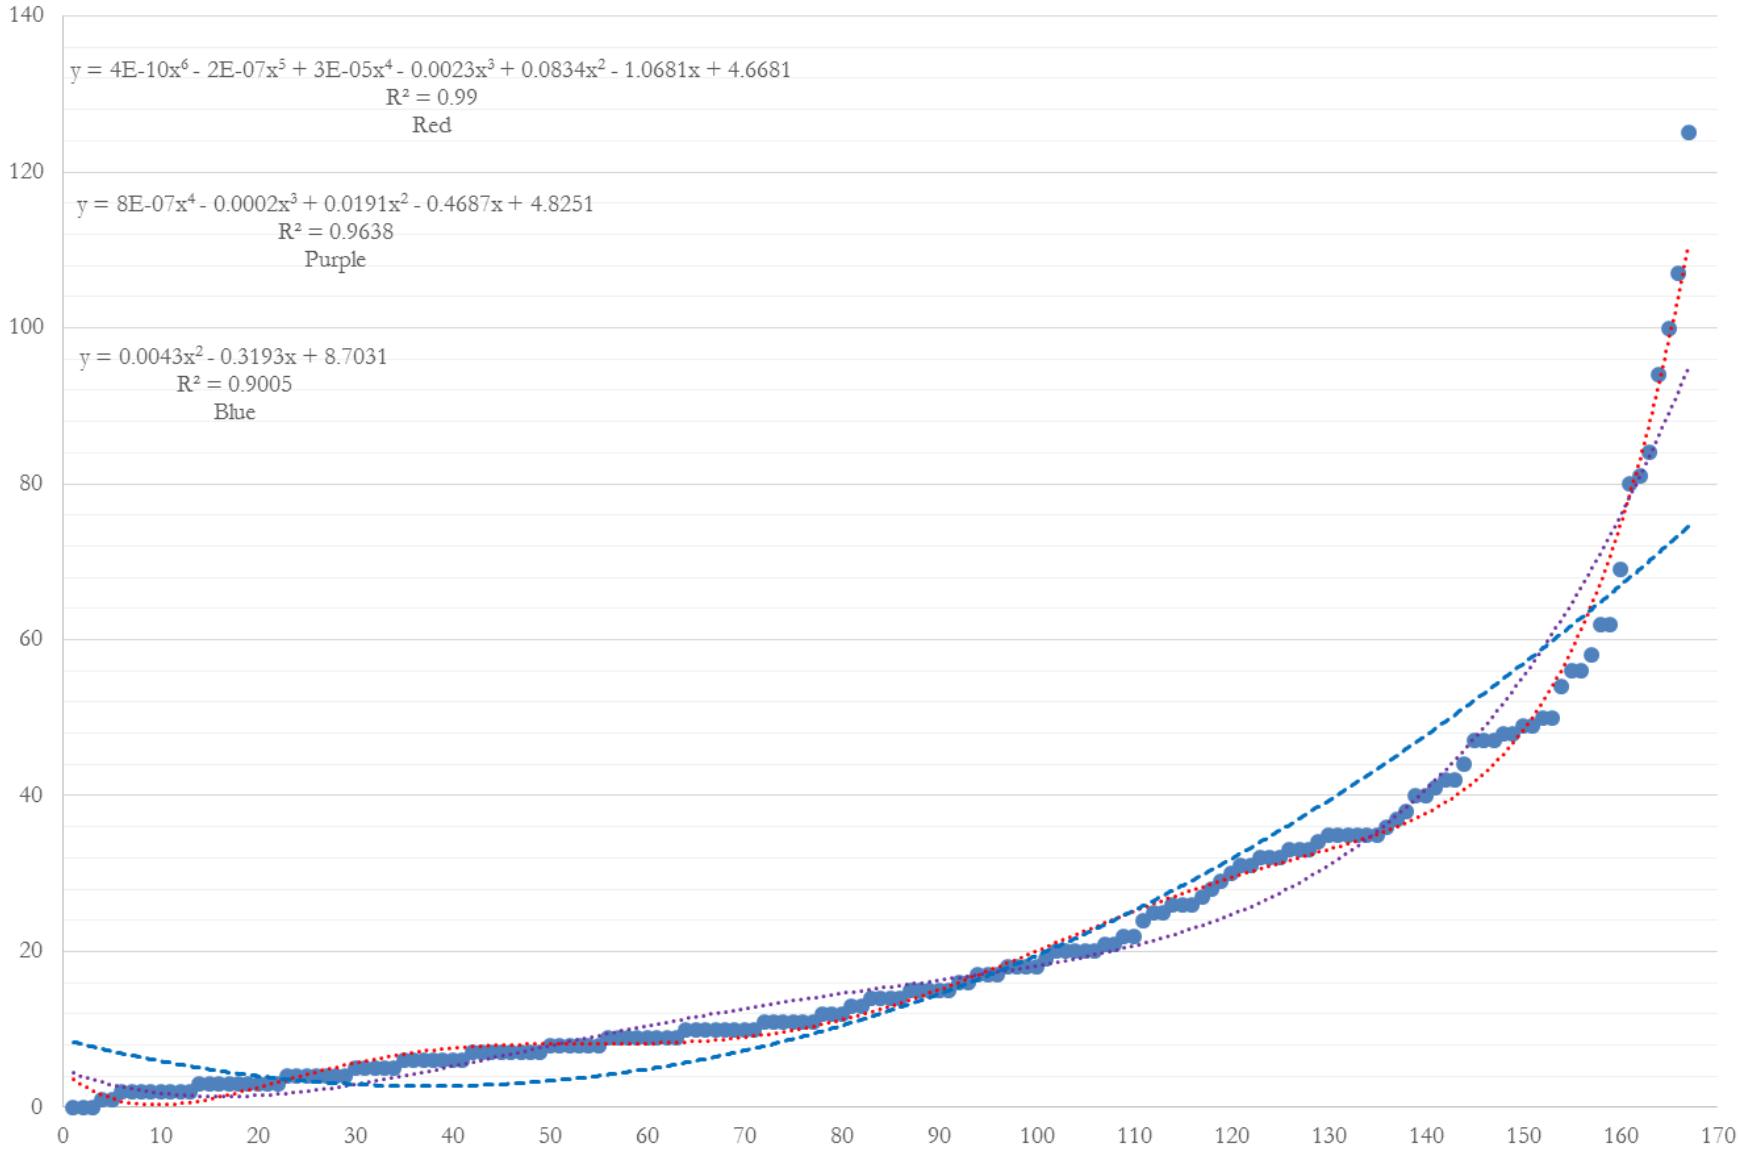

Figure 25 - Settlers of Catan Experience Scores as ranked from lowest to highest

In order to analyse Strategic Thinking Quotient in Settlers, used some of the questions I asked in my survey that that covered:

- Knowledge of strategy types (8 pts total, 4?s)
- Development Card Issues (3 pts, 1?)
- Position Order preferences (5 pts total, 2? (4 and1 pts))
- Basic Trading Preferences (4 pts total, 2?s)

This breakdown leads to a 20-point scale on which subjects have been graded and given points awarded per question, range of potential scores is x:[0,20].

The vast majority scored perfectly on the knowledge of general strategy questions.

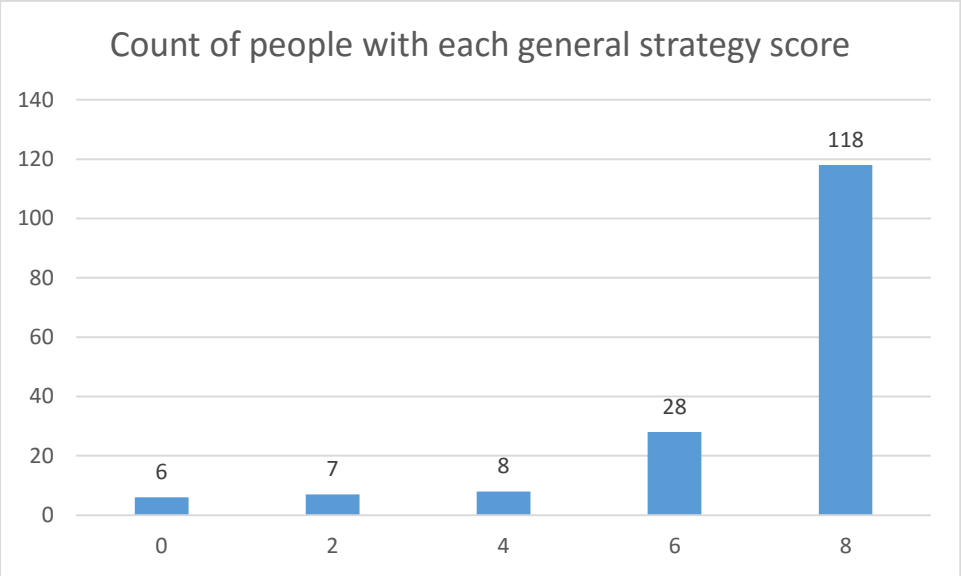

Figure 26 - Count of people with each general strategy score

As can be seen, 118 (70.7%), scored perfectly on this section, with another 28 (16.8%) scoring as proficient with general strategies.

In order to validate the basic strategic competence questions, I plotted how people scored on these and plotted their experience level. Experience is not the same as strategic competence, and that fact held in this study as well, however, experience did correlate to strategic

competence, which is to be expected.

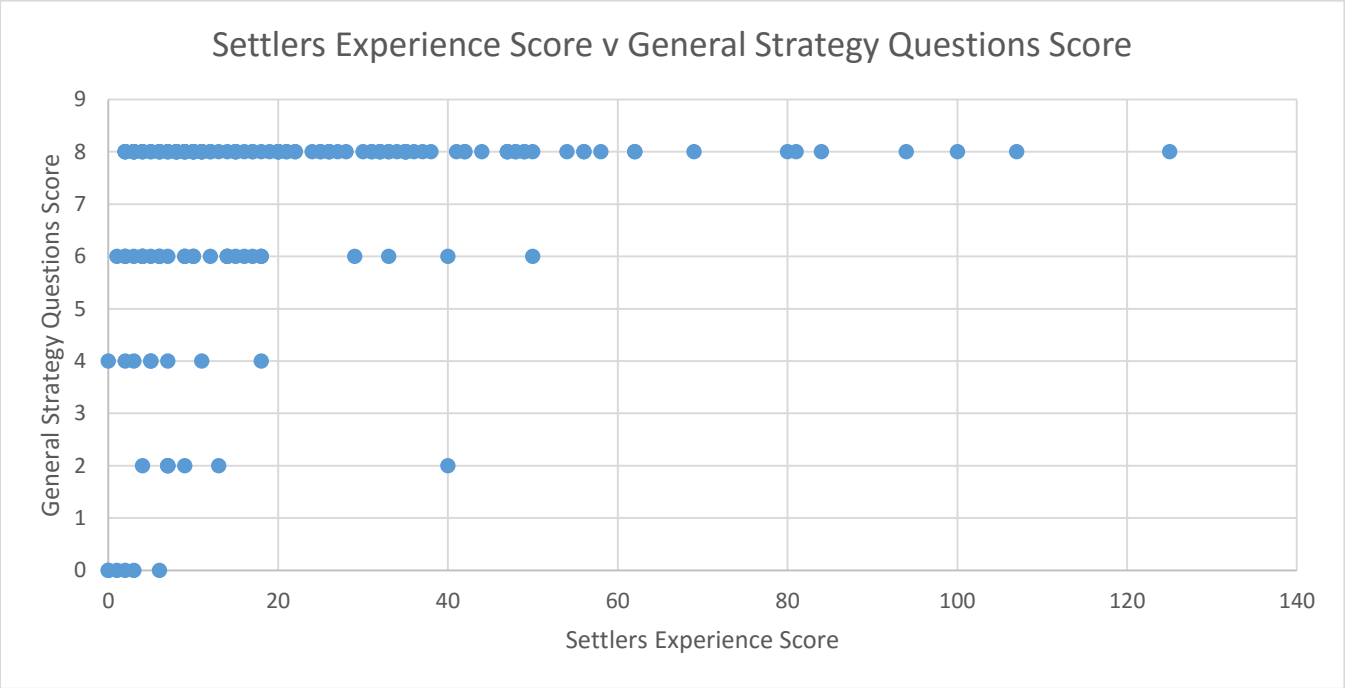

Figure 27 - Players' Settlers Experience Scores plotted against their General Strategy Questions Score

As can be seen in figure 27 above, as should be expected, those that scored poorly on the basic strategy questions were those with little experience in Settlers of Catan. However, it doesn't require that much experience to become familiar with the basic strategies of Settlers. However, of course, the more experience one has, the higher their average score. For instance, every survey taker with experience of 50+ scored 6 points or better (n = 16), and every survey taker who had experience of 54 or better (n = 14) scored a perfect 8 on the four basic strategy questions.

Scoring on development card issues was based on a single question: "Do you view any development card held for 3 or more turns as a victory point generating card?", which is based on the fact that players with VP generating cards, in my experience, leave them face down, so whilst there is the presumption of it being a VP card, or something that will bring VP (since a well-timed action card can allow for resource acquisition that will allow for construction, and, a player needs three knights played to acquire the Largest Army trophy). What I found was revelatory, if not expected. Only 81 (48.5%) of survey takers answered in the affirmative, 86 (51.5%) did not.

The scoring on position order preferences was arguably the most subjective of the questions asked, but the intent was to drive at what Szita *et al* noted in their paper, that whilst playing

against random players, that all things being equal, players in positions 1 and 4 were in the best position in a 4-player game, and that players in positions 1 and 3 were in the best positions in a 3-player game. The two questions were worth a total of 5 points, and the 4-player question was more heavily weighted.

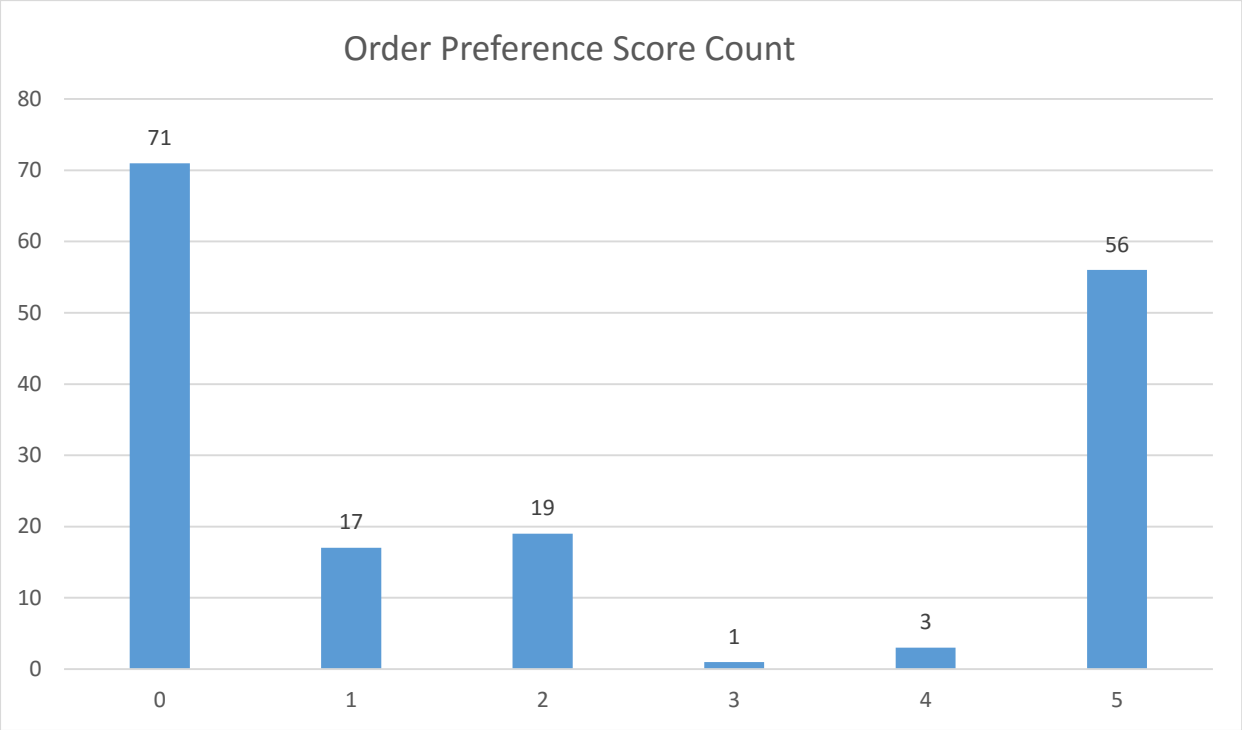

Figure 28 - Order Preference Score Count

Modally, as can be seen in Figure 28, the most common score was a 0, with 42.5% of respondents getting this score. The next most common was a perfect score, with just over 1/3 of respondents getting such a score. Anything in between is based on partial credit for the 4-player question (as in a player who demonstrated some thought in their answer if not totally correct though) as well as full credit for the 3-player question.

Opponent trading scoring was based on two questions and were worth 4 points total (2 points per question):

- Will you trade with your opponent if your opponent is at 8 pts?, and,
- Will you trade with your opponent if your opponent is at 9 pts?

Perfect scores of 4 were given for those who answered no to both questions, and of course, if one was answered right and the other wrong, they were awarded two points. Over 91% of respondents scored zero (n = 152), only eight scored 2 points and only 7 scored 4. My initial thought was that this was related to people's Machiavellianness or lack of the same, however, at least on the basis of this data (as can be seen in Figure 29), it would appear not to be the case

(however there were too few who scored either 2's or 4's to be statistically confident in this assertion).[86]

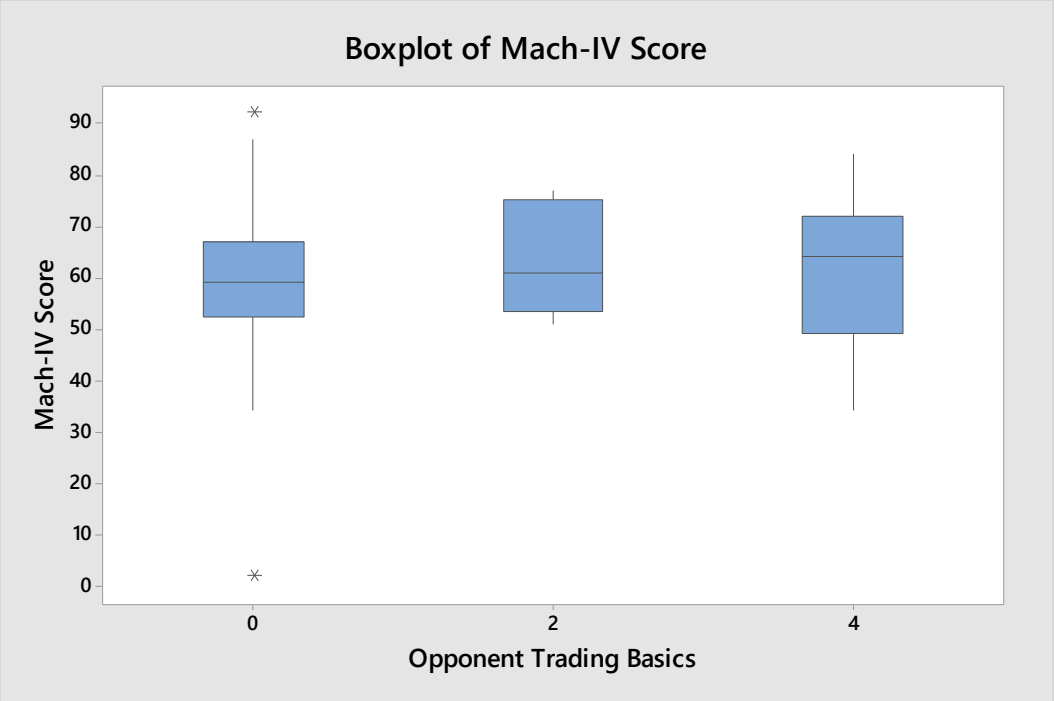

Figure 29 - Boxplot of MACH-IV Scores, organised by trading scores

[86] Trading N Mean StDev 95% CI
0 148 59.399 11.488 (57.507, 61.290)
2 8 63.25 10.44 (55.12, 71.38)
4 7 61.71 16.09 (53.02, 70.41)

Pooled StDev = 11.6502

The distribution of total scores is shown in figure 30:

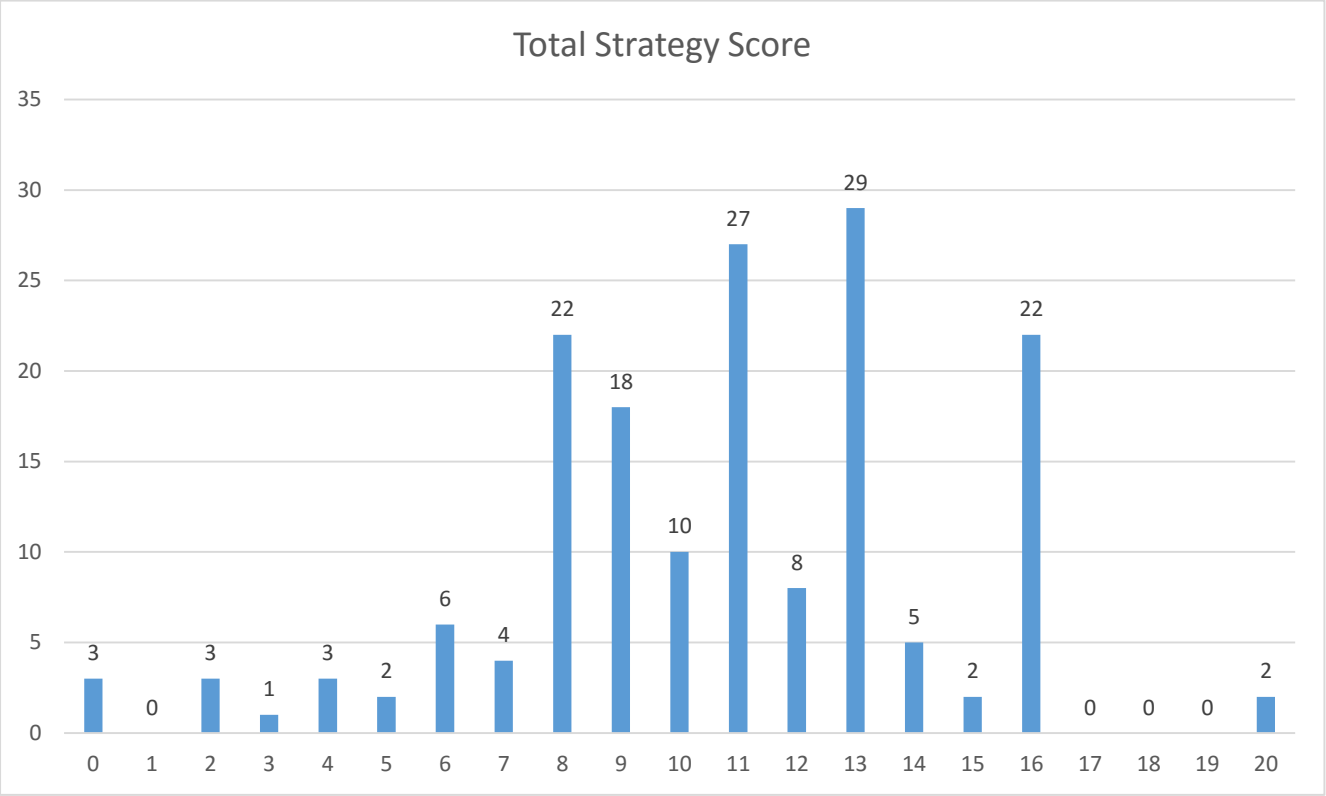

Figure 30 - Total Strategy Score Distribution

As can be seen, scores ranged from 0 to 20, with the vast majority of takers being assessed a score between 8 and 16 inclusive. Utilising those scores, I then binned all survey takers into one of six categories. Those categories were:

- x<8: incompetent (Level 0)
- x = 8: basic competence (Level 1)
- 9-10: high basic (Level 2)
- 11-12: intermediate (Level 3)
- 13-15: advanced (Level 4)
- 16-20: expert (Level 5)

As can be seen from Figure 31, this allowed relatively even binning of survey takers, with a majority clustering in the middle categories (something that seems logical).

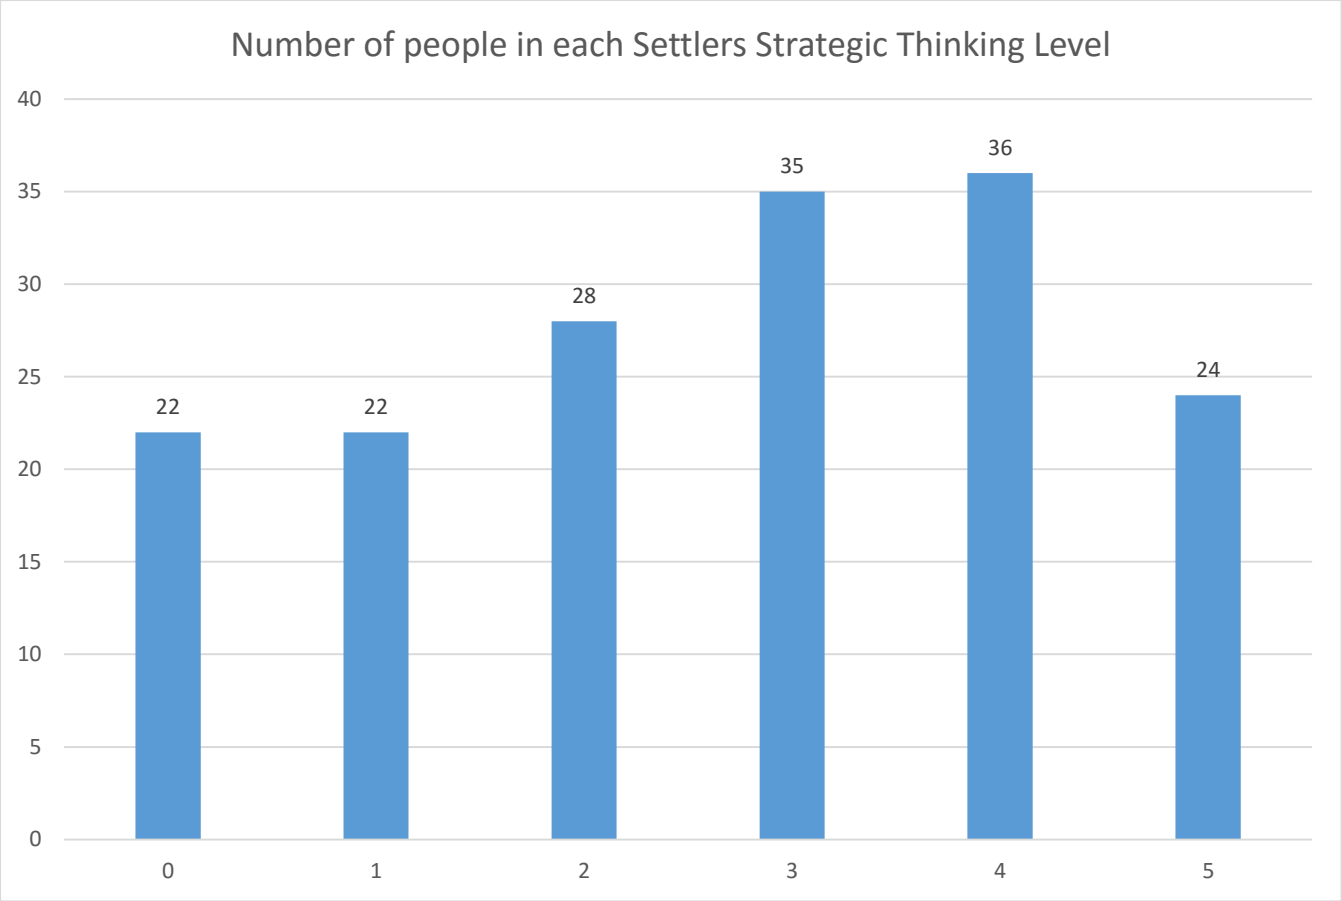

Figure 31 - Histogram of the distribution of Settlers of Catan Strategic Thinking Levels

It should be noted that there was no relationship between gender and strategic competence. My data indicates that players of both genders were equally strategically aware. The summary statistics are as follows:

	Q1	Q2	Q3	IQ Range	Lower Whisker	Upper Whisker
Male	8	11	13	5	2	20
Female	8	11	13	5	2	16

Table 5 - Settlers Strategic Survey Scores, separated by gender

The only reason for the discrepancy between the extent of the lower and upper whiskers seen in Table 5 is that the only two survey participants with perfect scores were, in fact, both men. In boxplot form, the distribution of scores is demonstrated in Figure 32:

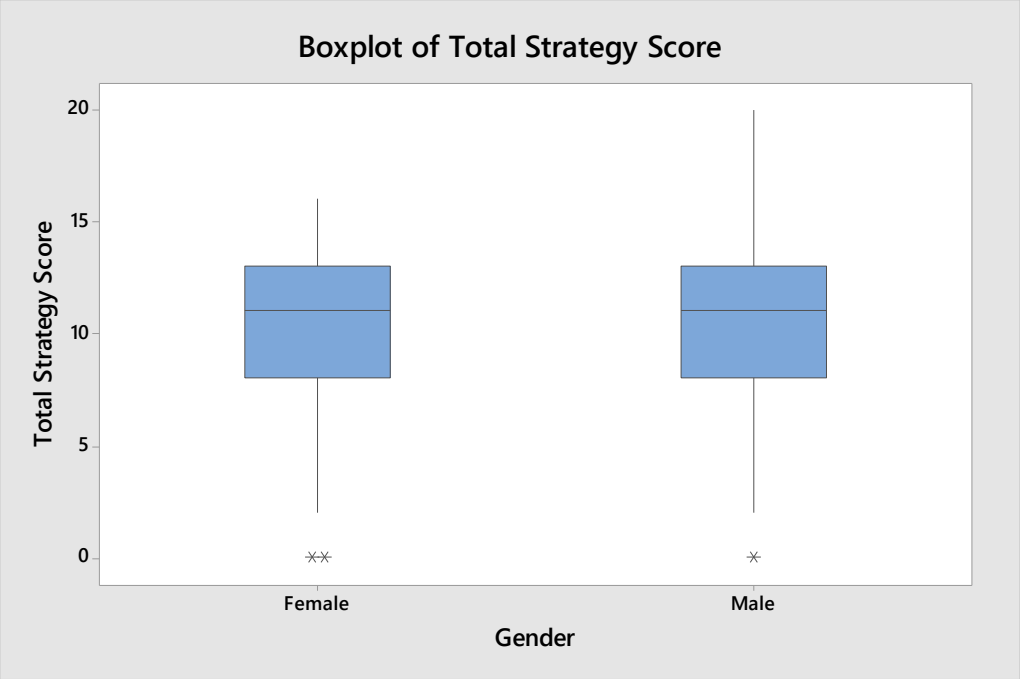

Figure 32 - Boxplot of Total Strategy Scores, separated by gender

In summary, what I found was that these few criteria in and of themselves were reasonably good at providing enough data for an accurate profile according to the classification learner in MATLAB that was utilised for this study, correctly selecting the right experience level in 80.8% of cases.

The parallel coordinates plot of this learner is shown in Figure 33 and I think is generally illuminating

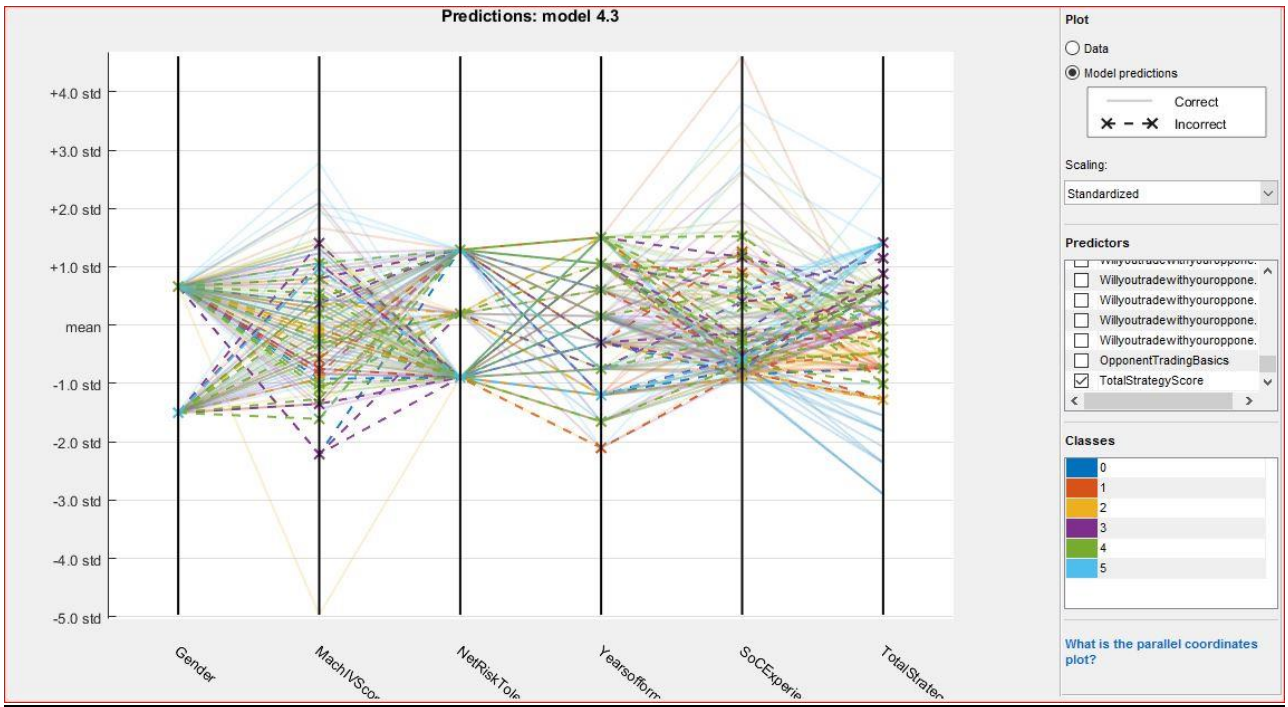

Figure 33 - Classification Learner Parallel Coordinates Plot (produced in MATLAB)

A parallel coordinates plot is a visualisation method used to plot individual data elements across multiple dimensions, where each of the dimensions corresponds to a vertical axis and each data element is displayed as a series of connected points along the axes. This system was invented in the 1970s by Alfred Inselberg as a means to visualise multiple dimension data.[87]

The confusion matrices are I think even more illuminating. First, the confusion matrix with no adornment is shown in Figure 34.

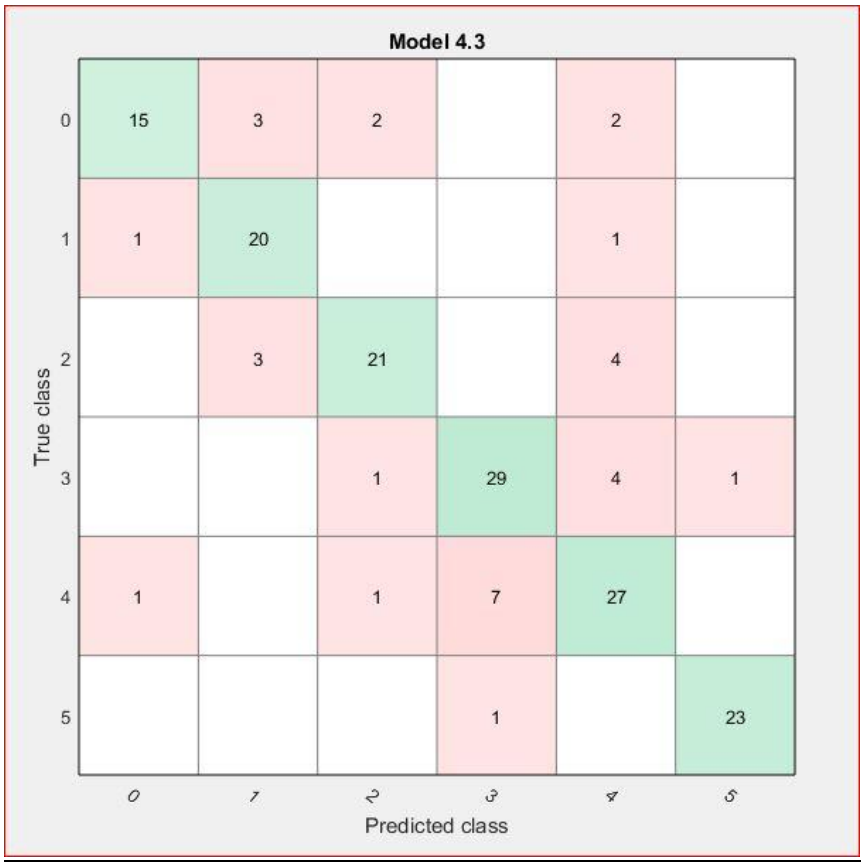

Figure 34 - Basic Confusion Matrix (produced in MATLAB)

As you can see, the vast majority of predicted strategic levels of individuals was, in fact, their true level. And now, for the confusion matrix shown in Figure 35 which demonstrates the rate of true positives to false negatives.

[87] There are several good basic references on a parallel coordinates plot. I referenced the one on Juice Analytics' blog (http://www.juiceanalytics.com/writing/writing/parallel-coordinates), which in turn references the Perceptual Edge's blog (http://www.perceptualedge.com/articles/b-eye/parallel_coordinates.pdf), which I also referenced, which was written by Stephen Few, a noted data visualisation consultant who has taught at the Haas School of Business (UC Berkeley) and elsewhere.

Figure 35 - True Positive v False Negative Confusion Matrix

Lastly, for the confusion matrix with the demonstration of positive predictive value to false discovery rate is shown in Figure 36:

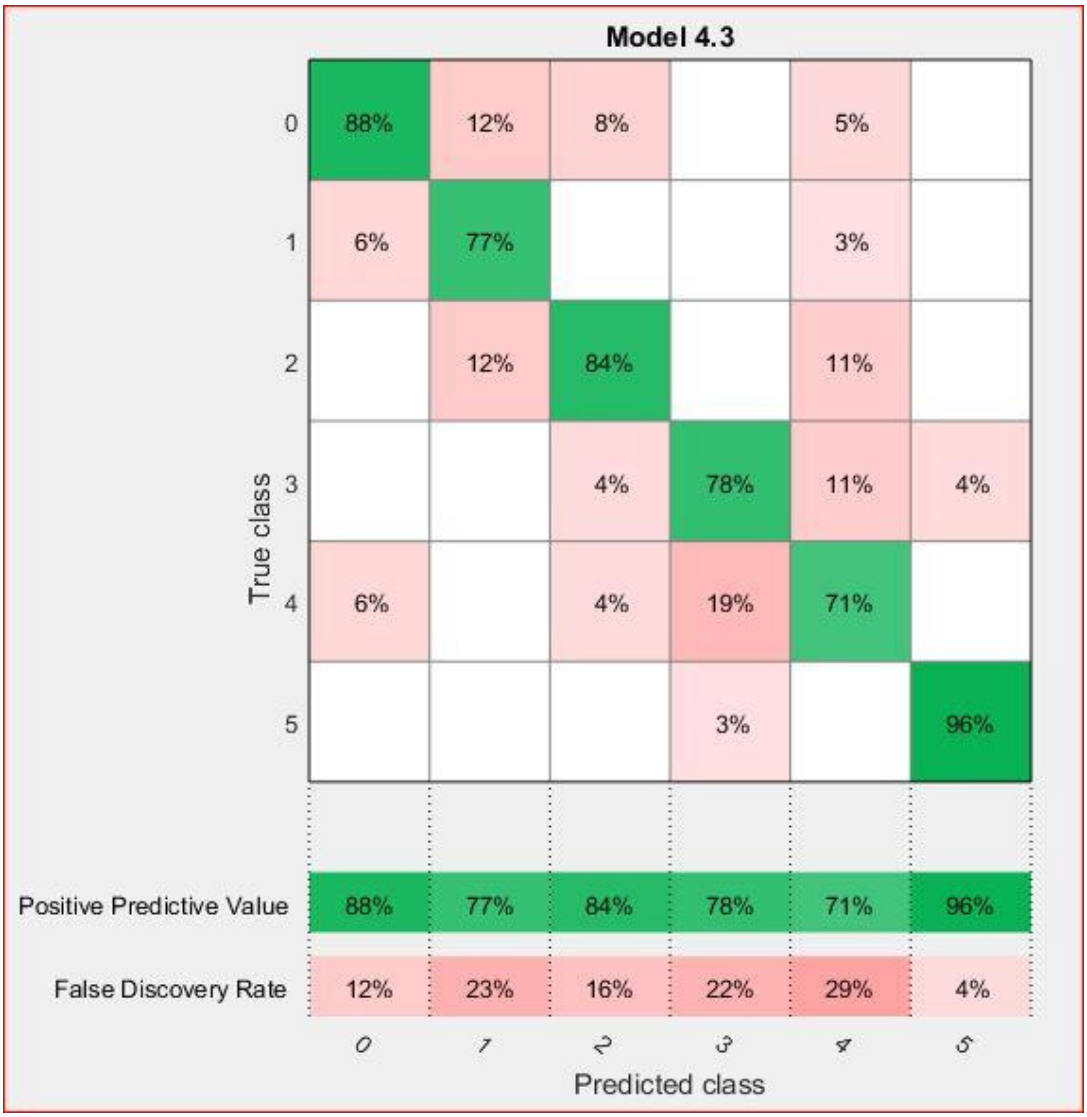

Figure 36 - Confusion Matrix, Positive Predictive Value to False Discovery Rate (produced in MATLAB)

Anyway you look at it, the few criteria I selected were good at predicting the strategic competence of participants. However, care must be exercised, given what Nate Silver said with respect to political prediction models: "But there are some big problems with fundamentals-based models. Namely, while they backtest well — they can "explain" past election results almost perfectly — they've done poorly at predicting elections when the results aren't known ahead of time."[88] That same issue exists in any model based on fundamentals, regardless of what it is describing. Point being is that in the future, any efforts in this regard, must be forward focused, in order to ensure that it the model IS robust at predicting competence, and not just being something that fits well after the fact.

[88] Silver, "How I Acted Like A Pundit And Screwed Up On Donald Trump", *Fivethirtyeight.com*, 18 May 2016

8) **Conclusions and Lessons Learned**

This research explores the potential for the incorporating behavioural information in decision problems, by designing and piloting an experimental protocol for eliciting and recording such information in the context of Settlers. Many lessons were learned and these should provide useful guidance for continuing research in this vein.

First (and surprisingly), because this research uses human subjects, we became aware that we required approval by UVA's Institutional Review Board (IRB). Whilst the approval process is relatively straightforward, given steps involved, and the time that was taken, it would have been ideal to have filed for IRB approval as soon as I had the idea that I would be interviewing subjects and recording gameplay for this study.[89]

Second (and in hindsight, not surprisingly), issues with my subjects abounded[90] Scheduling issues, inconsistent communication, etc. As noted previously, I created the secondary game type of observed games in the attempt to increase the number of data points given that the number of primary games I was able to conduct was under a third of what I had intended. Ample time (at least four to six months) should be budgeted if a larger dataset is to be acquired.

Third, given the means by which I solicited people to participate in taking the strategy survey, the sample skews heavily with respect to education. Compared with national averages, the sample pool was well over-represented by those with a baccalaureate degree or better in STEM subjects, comprising almost 59% of survey takers (n = 102). It may not be entirely unrepresentative of those who play Settlers of Catan, but this undeniably is a limitation in the dataset I assembled and in its validity in attempting to express how larger populations would act and may in fact bias the results in other ways (e.g. willingness to trade or Machiavellianness). This issue also is relevant with respect to soliciting player-participants for this study, since they were acquired in the same way that survey takers were. A similar limitation is often noted in experimental research in cognitive psychology, wherein a high majority of subjects for those studies are students enrolled in psychology courses. I acknowledge this is a sample of convenience, without having a specific remedy.

Fourth, there were limitations on the amount and type of data that one person is able to collect,

[89] I must express my appreciation in this regard to my roommate last year, Scott P. Laughery, Ph.D., given it was he who suggested that I check with IRB in the first place.

[90] This limitation was, to be fair, probably self-imposed

especially when, in order to ensure that there were enough players to play, not just did I have to record data, but that I had to act as one of the player-participants as well (in 15 of the 19 primary games that I conducted this was the case). It was hard to maintain focus as a player, or as a data recorder, when I had to split my attention between two roles. The course of certainly research would benefit from recruiting and training a knowledgeable assistant, perhaps an undergraduate or another grad student willing to volunteer their time.

Fifth, the lack of primary games conducted for this study greatly reduced the not just the quantity of data, but the quality of the data as well. It took all 20 primary games to perfect the Settlers recording protocol, so it would have been nice to have used that system to record more games, ensuring the highest quality of data was collected.[91] Using this system in the future should allow collection of a much larger sample with the same amount of time and effort and is recommended.

Sixth, the data affirms is that the Ellsberg paradox holds, which in and of itself isn't exactly revolutionary (given that the Ellsberg paradox has been around a bit more than a half century).[92] However, what I think is the potentially big news about all of this is that because of the Ellsberg paradox, and the fact that people's trading preferences are not logically consistent and violate the postulates of subjective expected utility, but that because of this fact, there is in fact no way to axiomatically describe (thus program) an entirely consistent and uniform decision making process without "fudging" (that is, forcing a certain decision every time one of these logically inconsistent decision pairs comes into play, and use of a probabilistic-based decider in this instance, is a "force"). Aspirations for future research should be framed accordingly.

This is important because in the land of artificial intelligence, even if humans aren't Vulcans (that is to say totally rational computational machines that act with perfect logic), the general assumption is that that you can still make a complete system of decision rules that are logically consistent even if they act sub-optimally. I assert that making such a system of decision rules is, in fact, impossible given the forcing function described above (which is at best a hollow approximation).

Seventh, another thing discovered is that an arguably major flaw exists in the distribution of number

[91] As I said previously, I may only have the data of 19 primary games recorded, but I did record 20 games, I lost the sheets for game 2, so games 2 – 19 as recorded were actually 3 – 20.

[92] As I describe earlier in this study, but will reiterate here, the Ellsberg paradox is a paradox in decision theory in which people's choices violate the postulates of subjective expected utility, where in effect, people prefer taking on a quantified risk they know more than a risk in which the odds are ambiguous, even if those odds are more in their favour than the risk with known odds.

discs. As I pointed out in my discussion of section 2, the number discs compose a set *N*, which looks like:

N={2,3,3,4,4,5,5,6,6,8,8,9,9,10,10,11,11,12}.

This distribution is of course based on the dice distribution. When evaluated against the idealised distribution, as is shown in Table 6, I found there to be a sum of squared error of 4.

Roll	2	3	4	5	6	8	9	10	11	12	
Discs	1	2	2	2	2	2	2	2	2	1	18
SE	0.16	0.64	0.04	0.16	1	1	0.16	0.04	0.64	0.16	4

Table 6 - Standard Error of the Actual Disc Distribution

As I can demonstrate in Table 7 and Figure 37, a small change in the distribution would reduce the sum of squared error from 4, to 0.8, which is an 80% reduction in the SSE. That change is to take one of the two discs assigned to the numbers 3 and 11, and reassign them to 6 and 8 to better reflect the central tendency of the dice distribution.

Roll	2	3	4	5	6	8	9	10	11	12	
Discs	1	1	2	2	3	3	2	2	1	1	18
SE	0.16	0.04	0.04	0.16	0	0	0.16	0.04	0.04	0.16	0.8

Table 7 - Standard Error of the Ideal Disc Distribution

I assert that because of this discrepancy between how the discs should have been enumerated and how they actually were, there is a decrease in realism and quality of gameplay.

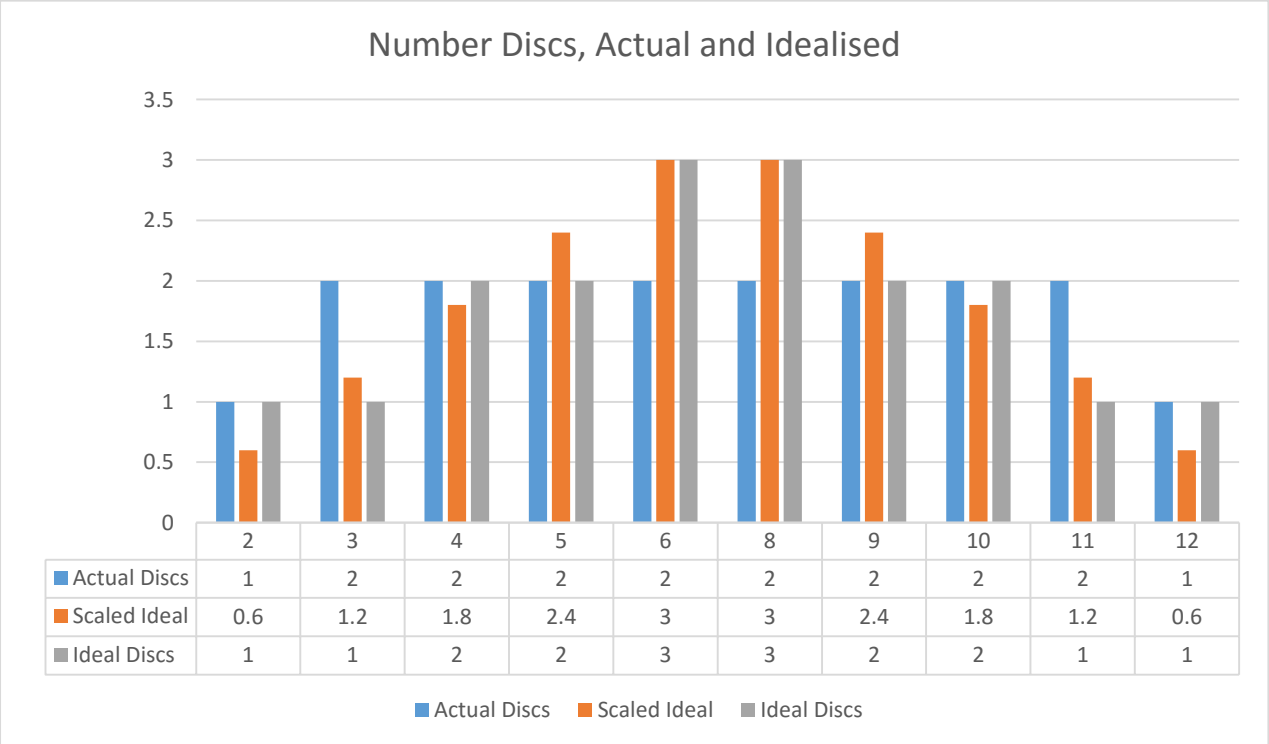

Figure 37 - Comparison of Actual and Idealised Number Discs to Scaled Dice Distribution

Figure 37 shows the actual disc distribution, the scaled ideal distribution, and the ideal disc distribution (which of course results from the fact that you can't have partially enumerated hexes nor partial discs). So, given the limitation of dealing in the natural numbers and only having 18 resource hexes (thus 18 enumerated discs), the ideal disc distribution I proffer is as optimal as it can get with respect to minimising the sum of squared errors. Making such a change to more realistically express the probabilities associated with the game would improve gameplay by improving realism.

9) **Where Can We Go From Here**

I think the simplest and most immediate thing that can be done is to collect more and better data. Nineteen primary games, whilst instructive, as there is a rough shape of things, is far from enough data to make wide-ranging conclusions, especially since only 12 of those 19 games were 4-player, which meant that there wasn't enough of the same type of data collected (either in 3-player games or 4-player games) to perform rugged statistical analyses of the game types, although there did appear to be differences in 3-player games and 4-player games, as I noted previously in section 5. The general strategic surveys themselves, whilst instructive, were certainly a rough instrument and could use refinement also, both on personality traits surveyed as well as an assessment of strategy and trading practices.

Additional psychometric tools, such as the five factor model of personality (also called the Big Five personality traits) or the HEXACO model of personality structure developed by Ashton and Lee should be brought into any future modelling attempts in this sphere.[93] Not only is there a relationship between the "honesty/humility" factor in the HEXACO model (it is the "H"), and Machiavellianness (as has been borne out by in a study conducted by Jones and Paulhus).[94] There are potentially interesting relationships to find, such as the relationship between certain personality traits and particular strategic choices (e.g., Are the extroverted more likely to trade? or Are the duplicitous more likely to use a monopoly card after trading away all of a particular resource? being two interesting questions).

Furthermore, targeted samplings of underrepresented populations would need to be conducted in order to ensure that the data collected with respect to Settlers play is widely reflective of how actual

[93] They were the primary contributors to "A Six-Factor Structure of Personality-Descriptive Adjectives: Solutions From Psycholexical Studies in Seven Languages.", sole authors of "Empirical, theoretical, and practical advantages of the HEXACO model of personality structure" as well as a bunch of other papers that weren't referenced for this study

[94] Chapter 7 of, *Handbook of Individual Differences in Social Behaviour*;

gameplay is done. The surveys used also need refinement, in particular the general strategy survey (which I had to modify as I was soliciting responses in order to remove confusion, as a few of the incomplete surveys were due to the fact that the player order preference questions weren't a sharply worded as they could have been initially).[95] And utilising the game recording method I developed and honed for this study, all future games of Settlers of Catan can be recorded in a consistent manner, which would benefit future research involving Settlers of Catan.

One of the results of this study is that a methodology to map personality characteristics to a strategic thinking quotient was developed, as illustrated through use of Settlers of Catan. Projecting strategic competence is difficult because of the issues of randomness, and the fact that someone can do everything that is correct theoretically, but still that does not always yield victory given that aleatory element. I have demonstrated how to use the game to map characteristics, so is generalizable, what has been demonstrated in one context, can be used in many others.

Where this is all headed is to a machine-learning model for AI agents. Collecting more game data along with all the demographic and personality data can be done in order to train the models to match certain personality types or styles of play (such as a risk-averse wood-brick, or high-Mach-highly-educated trader, for instance).[96]

Breaking down Settlers of Catan into its many subgames, and solving the perfect equilibria for those games, is needed in order to provide accurate weightings to the decision paths of the primary game. MCTS (or any probabilistically based methodology) is only as good as the data framing it, which is why the value exists for taking game theoretic tools and reapplying them in this context.

One of the papers discovered during the literature review discussed the use of a reinforcement learner in modelling Settlers of Catan. It did not however discuss utilising a game theoretic based analysis to build an MCTS-based model. Given that there do exist subgame perfect equilibria in Settlers of Catan (as I established in section 5 of this study), that means that such a model can be created. However, in order to do that, accurate recording of game states and large numbers of games recorded is required for such an analysis to take place, since without proper measurement of

[95] The end of game surveys were modified to collect data that I was asking of players, but had forgot to spell out on the end-of-game survey sheet.

[96] This application was borne out from a discussion with a friend of mine, Anthony Fazio, who is a senior software engineer at Microsoft, before which he took a degree in computer science from Princeton. He was also one of the Settlers strategy survey takers.

payouts in various scenarios, there is not a way to properly weight one's options.

So whilst this study is a step forward with respect to strategic modelling, what could be accomplished up until now has been limited by the paucity of accurate data and a means to record it. It is hoped that this study has created a foundation to take it to the next level.

10) References

- Abramson, B.; *The Expected Outcome Model of Two-Player Games*, PhD Dissertation, Columbia University, USA, 1987
- Ariely, D.; *Predictably Irrational: The Hidden Forces That Shape Our Decisions*, HarperCollins, 2008
- Ashton and Lee; "Empirical, theoretical, and practical advantages of the HEXACO model of personality structure". *Pers Soc Psychol Rev* **11** (2): 150–66, 2007
- Ashton, M., *et al*; "A Six-Factor Structure of Personality-Descriptive Adjectives: Solutions From Psycholexical Studies in Seven Languages.". *Journal of Personality and Social Psychology* **86** (2): 356–366, 2004
- Aumann and Hart (editors); *Handbook of Game Theory with Economic Applications*, Elsivier, 1992
- Balla and Fern, "UCT for Tactical Assault Planning in Real-Time Strategy Games", *Proc. 21st Int. Joint Conf. Artif. Intell., 40–45*, 2009.
- Bouzy, B.; "Old-fashioned Computer Go vs Monte-Carlo Go"; *IEEE Symposium on Computational Intelligence and Games,* Honolulu, 2007
- Branca and Johansson, "Using Multi-agent System Technologies in Settlers of Catan Bots", *Agent-based Systems for Human Learning and Entertainment,* Honolulu, 2006
- Browne *et al*., "A Survey of Monte Carlo Tree Search Methods", *IEEE Trans. Comp. Intell. AI Games, 4(1):1–43*, 2012.
- Chaslot *et al*, "Adding Expert Knowledge and Exploration in Monte-Carlo Tree Search", *Proc. Adv. Comput. Games*, LNCS 6048, 1–13, 2010.
- Chaslot *et al*, "Progressive Strategies for Monte Carlo Tree Search", *New Mathematics and Natural Computation 4 (3)*: 343–359, 2008
- Chaslot, *Monte-Carlo Tree Search*., PhD dissertation, Maastricht Univ., Netherlands, 2010
- Christie, R. and Geis, F.; "Studies in Machiavellianism", Academic Press, 1970
- Clark, K.; "Green Bay's Board-Game Obsession", *Wall Street Journal*, 15 January 2015 (http://www.wsj.com/articles/the-packers-of-catan-green-bays-board-game-obsession-1421346102)
- Coulom, "Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search", *Proceedings of the 5th International Conference on Computers and Games*, 2006.
- Dixit and Nalebuff; *Thinking Strategically*; W. W. Norton & Co., 1991.
- Ellsberg; *Risk, Ambiguity, and Decision*, PhD Dissertation, Harvard University, USA, 1962, published in 2001 by Garland Publishing.

- Guhe and Lascarides, "Game Strategies for The Settlers of Catan", *2014 IEEE Conference on Computational Intelligence and Games*, 2014
- Gunnthorsdottir, A., McCabe, K. and Smith, V., "Using the Machiavellianism instrument to predict trustworthiness in a bargaining game", *Journal of Economic Psychology 23*, 49-66, 2002
- Harford, T.; "Why we still love board games"*, FT Magazine*, 17 July 2010, accessed 29 June 2016, (http://www.ft.com/cms/s/2/1aab09a4-8fb2-11df-8df0-00144feab49a.html)
- Hayes, B.; "Computing Science: Statistics of Deadly Quarrels"; *American Scientist*, Vol. 90, No. 1 (JANUARY-FEBRUARY 2002). It can be found here: http://www.jstor.org/stable/27857587
- Holmes, J.; "Business Strategy Isn't Military Strategy", *The National Interest*, 5 April 2016, accessed 28 June 2016, (http://nationalinterest.org/feature/business-strategy-isnt-military-strategy-15681)
- Holmes, J.; "Everything You Know about Clausewitz is Wrong", *The Diplomat*, November 2014 (http://thediplomat.com/2014/11/everything-you-know-about-clausewitz-is-wrong/)
- Hunt; "Lewis Fry Richardson and His Contributions to Mathematics, Meteorology, and Models of Conflict", *Annual Review of Fluid Mechanics*, Vol. 30: xiii-xxxvi, January 1998
- Jones, D., and, Paulhus, D.; "Chapter 7. Machiavellianism", in Leary, M. and Hoyle, R; *Handbook of Individual Differences in Social Behaviour*; The Guilford Press. pp. 257–273, 2009
- Kocsis and Szepevari, "Bandit based Monte-Carlo Planning". *Machine Learning: ECML 2006, 17th European Conference on Machine Learning, Berlin, Germany, September 18–22, 2006, Proceedings. Lecture Notes in Computer Science 4212*, 2006
- Lai, *Giraffe: Using Deep Reinforcement Learning to Play Chess*, M.Sc. Thesis, Imperial College London, UK, 2015
- Luttwak, E.; *Strategy: The Logic of War and Peace*; Belknap/Harvard University Press
- Machiavelli, N. (translator: Lynch, C), *Art of War*, University of Chicago Press, 2003
- Méhat and Cazenave, "Combining UCT and Nested Monte Carlo Search for Single-Player General Game Playing."*, IEEE Trans. Comp. Intell. AI Games, 2(4):271–277,* 2010.
- Metz, C.; "What the AI Behind AlphaGo can teach us about being human"; *Wired Magazine*; 19 May 2016 (http://www.wired.com/2016/05/google-alpha-go-ai/)
- Monte Carlo Tree Search (MCTS) research hub, http://www.mcts-hub.net/index.html
- Möring, *Games and Metaphor – A critical analysis of the metaphor discourse in game studies*, PhD dissertation, IT University of Copenhagen, Denmark, 2013
- Moyer, C.; "How Google's AlphaGo Beat a Go World Champion", *The Atlantic*, 28 March 2016 (http://www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611/)

- Ness *et al.*, ed; *The Science and Simulation of Human Performance* (Advances in Human Performance and Cognitive Engineering Research, Volume 5); Emerald Group Publishing, 2004
- Nisan *et al.*, ed.; *Algorithmic Game Theory*, Cambridge University Press, 2007
- Osborne, M. J.; *An Introduction to Game Theory*; Oxford University Press, 2004
- Pfeiffer, M.; "Reinforcement Learning of Strategies for Settlers of Catan"; *Proceedings of the International Conference on Computer Games: Artificial Intelligence, Design and Education*, Reading, 2004.
- Raghuram *et al*, "On the Behavior of UCT in Synthetic Search Spaces", *Proc. 21st Int. Conf. Automat. Plan. Sched.*, Freiburg, Germany. 2011.
- Rapoport, A.; "Lewis F. Richardson's mathematical theory of war", *Journal of Conflict Resolution* vol. 1 no. 3 249-299, September 1957
- Richardson, L. (ed. Wright and Lienau); *Statistics of Deadly Quarrels*, Boxwood Press, 1960
- Richardson, L.; (ed. Rashevsky and Trucco); *Arms & Insecurity: A Mathematical Study of the Causes and Origins of War*, Boxwood Press, 1960.
- Richardson, L.; *Mathematical Psychology of War*, W. Hunt, 1919
- Saleem and Raees Natiq; *A Multi-agent player for Settlers of Catan - A MAS Player in game playing*; MSc Thesis, Blekinge Institute of Technology, Sweden, 2008
- Shannon, C.; "Programming a Computer for Playing Chess"; *Philosophical Magazine* **41** (314), 1950
- Shoham, Y.; "Computer Science and Game Theory", *Communications of the ACM* **51** (8): 75–79, 2008
- Silver (David) *et al*; Mastering the game of Go with deep neural networks and tree search; *Nature* **529**, 484–489, 28 January 2016 (http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html)
- Silver, N.; "How I Acted Like A Pundit And Screwed Up On Donald Trump", *Fivethirtyeight.com*, 18 May 2016 (https://fivethirtyeight.com/features/how-i-acted-like-a-pundit-and-screwed-up-on-donald-trump/?ex_cid=538twitter)
- Stopera; *19 Pictures That Are Too Real For People Who Play Settlers Of Catan*, *Buzzfeed*, 9 March 2015. (https://www.buzzfeed.com/daves4/i-dont-need-your-dang-sheep?utm_term=.sn5exGQV0#.ts7zY2dwj)
- Szita *et al*, "Monte-Carlo Tree Search in Settlers of Catan", *Proc. Adv. Comput. Games, 21–32*, 2010.
- Tam, P; "An Old-School Board Game Goes Viral Among Silicon Valley's Techie Crowd", 17 December 2009, *Wall Street Journal*

- Thielscher, "A General Game Description Language for Incomplete Information Games", *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, 994–999, 2010
- Thomas, R.; *Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic*; PhD dissertation, Northwestern University, USA, 2004. (This study used SoC as the basis of the work)
- Thompson, J. Mark; "Defining the Abstract", *The Games Journal*, 2000. (http://www.thegamesjournal.com/articles/DefiningtheAbstract.shtml)
- Trunda and Bartak, "Using Monte Carlo Tree Search to Solve Planning Problems in Transportation Domains", *Advances in Soft Computing and Its Applications: 12th Mexican International Conference on Artificial Intelligence* (MICAI 2013), Castro *et al* (eds.)
- Tversky, A., and Kahneman, D. (1981). "The framing of decisions and the psychology of choice", ***Science*** **211** (4481): 453–458.
- United States Marine Corps, *Marine Corps Doctrinal Publication (MCDP) 1: Warfighting*
- United States Marine Corps, *Marine Corps Doctrinal Publication (MCDP) 1-1: Strategy*

11) Acknowledgements

I hope I do not forget to acknowledge anyone whom I owe (however if I do, that fault is mine alone, and I hope I'll be forgiven). As with any collaborative process, and this was undeniably collaborative process, many hands have played a part in the formation of this study and they will be mentioned here (all veterans mentioned within are noted by their branch of service). I must also bear responsibility for any and all errors carried within.

The students in SYS 3021 in Fall 2014, I definitely appreciated serving as your mentor and guide, as limited as I may have been in that role. You all were and are interesting and fun people, and you helped make my first semester at UVA worth a lot, and were undeniably more worthy of UVA than UVA was of you. Shout outs especially go out to those who I've been able to develop friendships with after, including: Joseph Fitzsimmons, David Guinn (USN), Merritt Hale (USN), Conor Mettenberg (USN), Eric MacBlane, Anne Moenning, Dimitri Ofiesh, Stephen Walshak, Callum Weinberg, and Gregory Waldrip. Emily Bisallon (USN) was not my student, but given she was dating David who was, and how often I saw her, she might as well have been.

The professors, instructors, and mentors I have had, undeniably I have gained so much from you. At UW, this list includes: John B Sullivan, PhD; Peter W Soverel, PhD (USN); Ana Mari Cauce, PhD; Anthony J Gill, PhD; Christopher Jones, PhD; Adam D Moore, PhD; Alexander F Hanning, Rene Singleton; Phillip J Hunt (USAF); and L Lincoln Johnson. At UVA, this list includes: William T Scherer, PhD, who taught me to ask the right questions and go outside my lane; Alfredo Garcia, PhD; Brendan J Boler (USA); my thesis committee members (Mark R Floryan, PhD; James W Lark 3d, PhD; committee chair Michael C Smith, PhD) and my thesis adviser, K Preston White, Jr, PhD, all to whom I am grateful beyond measure, especially Prof White who went along with this crazy idea. Acknowledgement must also go to Barry G Silverman, PhD, at the University of Pennsylvania, who has elected to take me on as a doctoral student under his tutelage, so that I can continue the research that I have begun under Prof White.

The player-participants of the games I observed and those who took my survey, without you, I'd have had no data upon which to do analysis, especially to my friends whom I annoyed to no end about taking the survey.

The friends and relatives who have been my sounding board and helped kept me sane, serving as a lifeline to the lonely and crazy life that is being a graduate student, and many of whom have served

as readers to this work or otherwise supported me in the adventure that is my life (and especially my life at UVA) and with this study: Aaron D Schwitters (USA); Aaron P Tatryek, JD; Alexander B Berezow, PhD; Alexandra C L Truzzi; Alissa L R Maxwell, Andre B LaRoche, JD, Anthony J Fazio, Ariel Hargrave (USMC), Ariel J Greene (IDF); Ashley N Parcells; Augustine F Eck (USA); Bonnie R Brown, PhD; Brandon M Ray (USN), Brian Strom (USN/USMC), Cameron Vest (CLAS 2016), Christina Leonhard-Melief, Ph.D., Claire L McQuin, Craig J A Truzzi, Daniel T Anderson, Daniel Mow, JD (USA); Danielle Flatt, JD; David A Nash-Mendez, Demetrios G Lagos, Rabbi Elie Estrin (USAF), Elizabeth Anderson; Evan M Roberts, JD; Gavriel G Jacobs, JD (USMC), Jacob S Shepherd (USA), Jason D Hansman (USA), John B B Maxwell (USA), John R Rutledge (USA), Josef G Stute, Josiah S. Byers, Justin W Morris (CLAS 2016), Katerina G Lagos Tsakopoulos, DPhil; Kayleen M Hooley; Marc L Lamothe (USA); Mary Grace White, JD; Matthew J Schneiderjan, MD (USA); Michael P Eisenberg, Michael L Malinowski (SEAS 2012, SEAS 2016), Nicholas A Fusso, Nicholas R Kumleben, Fr Raphael Mary (Wesley) Salzillo, OP; Robert C Olson, Sameer Kanal, Scott P Laughery, PhD (CLAS 2016), Shir Lerman, PhD, Trond Nilsen, PhD; William T Mari, PhD (USN). And Zachary J Charat.

Julia Abelev, JD (Law, 2013), who brought me into the fold of Mr Jefferson's University, the debt I owe you TJ is massive and in many ways indescribable, Bryan R Lewis (CLAS 2013, SEAS 2015), it may have been random chance that lead you to be assigned as my departmental mentor, but you helped me find a home in the department, and lastly to my parents, Robert and Sophia Everett, without whom, I couldn't have done a thing.

This study is dedicated to them, and more importantly, to those who are currently serving and those that will serve in the Armed Forces of the United States and her allies, and for law enforcement officers throughout the world, for whom it is hoped that this study will help lessen the burdens they bear by enhancing the training models they use, instructing better ways how to plan and act, and lastly to those who have served in those military and police forces, for whom have lived the reasons why enhanced training models are needed.

12) **Appendices**

Appendix 1: Table of All Dice Rolls in Primary Games

Appendix 2: End-of-game player summaries

Appendix 3: Example Recorded Game

Appendix 4: IRB Oral Consent Script

Appendix 5: End of game surveys (modified version)

Roll	Game 1	Game 2	Game 3	Game 4	Game 5	Game 6	Game 7	Game 8	Game 9	Game 10	Game 11	Game 12	Game 13	Game 14	Game 15	Game 16	Game 17	Game 18	Game 19	Sums
2	2	2	2	0	0	3	1	1	7	1	1	2	5	2	3	1	0	1	1	35
3	7	5	4	1	1	4	3	3	1	4	2	4	6	9	6	8	6	2	3	79
4	7	4	4	4	2	2	5	3	2	6	4	3	2	6	3	7	6	3	11	84
5	12	5	8	8	5	6	5	7	9	5	6	6	6	3	5	13	9	5	7	130
6	5	11	6	7	2	15	6	10	2	9	9	5	8	6	4	13	8	7	7	140
7	11	14	14	15	6	12	7	14	7	12	14	13	10	12	12	19	8	2	14	216
8	12	8	8	6	6	8	5	14	8	5	3	5	10	8	6	9	9	8	5	143
9	12	4	4	17	7	10	13	7	9	10	11	8	3	6	6	5	6	4	7	149
10	8	4	5	10	4	3	4	9	9	6	3	4	4	6	5	6	6	3	5	104
11	2	2	3	3	3	5	3	1	6	3	1	2	5	3	6	6	5	2	4	65
12	2	1	2	4	0	2	1	7	1	1	1	1	3	1	3	3	2	1	2	38
Roll Total	80	60	60	75	36	70	53	76	61	62	55	53	62	62	59	90	65	38	66	1183
Player Count	3	4	3	3	3	3	4	3	4	4	4	3	4	4	4	4	4	4	4	

Appendix 2: End-of-game player summaries

Game #	Player	Colour	Order	Strategy-1	Strategy-2	Change	City	Sett	Bld Total	Trophy	Long Rd	Total Rds	Kts	VP Cds	Act Cds	Dev Cds	VP	
1	*1*	*Red*	*1*	*5*	*2*	*yes*	*0*	*3*	*3*	*1*	*9*	*9*	*5*	*0*	*0*	*5*	*5*	
1	*2*	*White*	*2*	*2*	*3*	*no*	*4*	*2*	*10*	*0*	*9*	*10*	*3*	*0*	*0*	*3*	*10*	
1	*3*	*Orange*	*3*	*1*	*6*	*no*	*1*	*5*	*7*	*1*	*14*	*14*	*2*	*0*	*2*	*4*	*9*	
2	**4**	**Blue**	**1**	**3**	**5**	**yes**	**4**	**0**	**8**	**0**	**2**	**4**	**3**	**0**	**1**	**4**	**8**	
2	**1**	**Red**	**2**	**6**	**5**	**no**	**1**	**3**	**5**	**0**	**7**	**7**	**0**	**0**	**0**	**0**	**5**	
2	**5**	**White**	**3**	**4**	**1**	**no**	**2**	**3**	**7**	**1**	**10**	**10**	**0**	**0**	**0**	**0**	**9**	
2	**3**	**Orange**	**4**	**2**	**0**	**no**	**4**	**1**	**9**	**0**	**4**	**9**	**1**	**1**	**0**	**2**	**10**	
3	*7*	*Blue*	*1*	*5*	*0*	*no*	*0*	*5*	*5*	*0*	*5*	*7*	*3*	*2*	*0*	*5*	*7*	
3	*6*	*Orange*	*2*	*4*	*6*	*no*	*2*	*2*	*6*	*2*	*6*	*8*	*3*	*0*	*3*	*6*	*10*	
3	*1*	*Red*	*3*	*2*	*5*	*no*	*2*	*2*	*6*	*0*	*4*	*7*	*0*	*0*	*1*	*2*	*6*	
4	*1*	*Red*	*1*	*1*	*6*	*no*	*0*	*4*	*4*	*1*	*10*	*10*	*3*	*2*	*1*	*6*	*8*	*Dev Cd St*
4	*7*	*Blue*	*2*	*4*	*5*	*yes*	*0*	*5*	*5*	*1*	*6*	*10*	*4*	*2*	*4*	*10*	*8*	*Dev Cd St*
4	*6*	*Orange*	*3*	*3*	*6*	*no*	*4*	*0*	*8*	*0*	*2*	*5*	*2*	*2*	*0*	*4*	*10*	
5	*9*	*Blue*	*1*	*1*	*5*	*yes*	*0*	*2*	*2*	*0*	*3*	*4*	*0*	*0*	*0*	*0*	*2*	
5	*1*	*Red*	*2*	*2*	*3*	*no*	*0*	*3*	*3*	*0*	*4*	*5*	*2*	*0*	*1*	*3*	*3*	
5	*8*	*White*	*3*	*1*	*6*	*no*	*3*	*1*	*7*	*1*	*3*	*6*	*3*	*1*	*0*	*4*	*10*	
6	*9*	*Blue*	*1*	*2*	*3*	*yes*	*1*	*4*	*6*	*1*	*4*	*7*	*4*	*2*	*0*	*6*	*10*	
6	*1*	*Red*	*2*	*2*	*6*	*yes*	*0*	*3*	*3*	*1*	*10*	*13*	*2*	*1*	*1*	*4*	*6*	*Dev Cd St*
6	*8*	*White*	*3*	*6*	*0*	*yes*	*2*	*5*	*9*	*0*	*7*	*10*	*2*	*0*	*2*	*4*	*9*	
7	**4**	**Blue**	**1**	**1**	**4**	**yes**	**3**	**1**	**7**	**1**	**8**	**9**	**1**	**1**	**1**	**3**	**10**	
7	**5**	**White**	**2**	**1**	**5**	**no**	**2**	**2**	**6**	**0**	**4**	**5**	**0**	**0**	**0**	**0**	**6**	
7	**3**	**Orange**	**3**	**3**	**2**	**no**	**2**	**2**	**6**	**0**	**2**	**4**	**0**	**0**	**0**	**0**	**6**	
7	**1**	**Red**	**4**	**1**	**6**	**no**	**2**	**2**	**6**	**0**	**7**	**7**	**1**	**0**	**0**	**1**	**6**	
8	*5*	*White*	*1*	*4*	*6*	*no*	*1*	*4*	*6*	*1*	*7*	*7*	*2*	*2*	*3*	*7*	*10*	*Dev Cd St*
8	*3*	*Orange*	*2*	*4*	*2*	*no*	*2*	*1*	*5*	*0*	*6*	*7*	*1*	*0*	*0*	*1*	*5*	
8	*1*	*Red*	*3*	*2*	*5*	*no*	*0*	*4*	*4*	*1*	*4*	*5*	*4*	*0*	*2*	*6*	*6*	*Dev Cd St*
9	**6**	**Orange**	**1**	**4**	**6**	**no**	**1**	**5**	**7**	**1**	**10**	**11**	**0**	**1**	**0**	**1**	**10**	
9	**7**	**White**	**2**	**1**	**5**	**yes**	**0**	**5**	**5**	**0**	**9**	**12**	**1**	**1**	**1**	**3**	**6**	

9	1	Red	3	2	6	yes	2	2	6	0	5	9	1	0	0	1	6	
9	4	Blue	4	3	6	no	0	5	5	0	6	7	1	1	0	2	6	
10	1	Red	1	1	3	yes	0	4	4	1	8	8	1	0	0	1	6	Dev Cd St
10	10	Blue	2	2	5	yes	3	0	6	0	3	4	5	2	2	9	10	Dev Cd St
10	11	Orange	3	3	0	no	1	3	5	0	4	6	2	2	0	4	7	
10	5	White	4	5	2	yes	4	0	8	0	3	5	1	1	1	3	9	
11	1	Red	1	1	6	no	0	5	5	0	9	9	1	0	0	1	5	Dev Cd St
11	10	Blue	2	5	2	yes	0	4	4	0	4	5	1	0	1	2	4	
11	11	Orange	3	2	0	no	1	4	6	2	10	10	3	2	1	6	11	Dev Cd St
11	5	White	4	1	0	no	0	5	5	0	7	8	0	0	0	0	5	
12	*3*	*Orange*	*1*	*4*	*2*	*no*	*2*	*2*	*6*	*0*	*4*	*6*	*1*	*0*	*0*	*1*	*6*	
12	*5*	*White*	*2*	*5*	*0*	*no*	*1*	*5*	*7*	*1*	*7*	*9*	*0*	*1*	*0*	*1*	*10*	
12	*1*	*Red*	*3*	*2*	*0*	*yes*	*0*	*3*	*3*	*1*	*3*	*4*	*4*	*2*	*0*	*6*	*7*	*Dev Cd St*
13	6	Blue	1	4	6	no	1	3	5	1	4	6	3	0	0	3	7	
13	3	Orange	2	4	5	yes	4	0	8	1	6	7	0	0	0	0	10	
13	5	White	3	5	2	yes	3	1	7	0	6	7	2	1	1	4	8	
13	1	Red	4	2	6	no	2	1	5	0	3	5	1	1	1	3	6	Dev Cd St
14	1	Red	1	2	6	no	2	1	5	0	6	8	0	1	0	1	6	
14	3	Orange	2	4	5	no	2	4	8	1	11	12	3	0	2	5	10	
14	5	White	3	5	0	no	0	5	5	0	4	7	0	3	1	4	8	Dev Cd St
14	11	Blue	4	2	0	no	3	0	6	0	8	8	2	1	1	4	8	Dev Cd St
15	1	Red	1	5	1	yes	0	5	5	0	8	10	0	0	0	0	5	
15	3	Orange	2	4	5	no	2	3	7	0	6	8	0	0	0	0	7	
15	5	White	3	5	0	no	2	4	8	1	9	10	0	0	0	0	10	
15	11	Blue	4	4	0	no	0	3	3	0	4	5	2	2	2	6	5	Dev Cd St
16	14	Orange	1	3	6	yes	2	2	6	1	8	8	3	0	3	6	8	Dev Cd St
16	12	Blue	2	3	0	no	4	0	8	1	10	10	0	0	0	0	10	
16	13	White	3	1	2	yes	2	3	7	0	6	9	2	0	0	2	7	
16	15	Red	4	6	0	yes	1	3	5	0	5	8	3	2	0	5	7	Dev Cd St
17	12	Blue	1	2	0	no	3	2	8	1	4	8	4	0	0	4	10	Dev Cd St
17	13	White	2	3	5	yes	0	5	5	0	8	9	0	1	0	1	6	

17	**14**	**Orange**	**3**	**1**	**3**	**no**	**1**	**4**	**6**	**1**	**10**	**10**	**1**	**1**	**0**	**2**	**9**	
17	**15**	**Red**	**4**	**2**	**5**	**yes**	**1**	**3**	**5**	**0**	**5**	**6**	**2**	**0**	**0**	**2**	**5**	Dev Cd St
18	**17**	**Blue**	**1**	**6**	**5**	**No**	**1**	**3**	**5**	**0**	**2**	**4**	**2**	**1**	**1**	**4**	**6**	
18	**19**	**White**	**2**	**2**	**0**	**no**	**2**	**1**	**5**	**2**	**7**	**7**	**3**	**1**	**1**	**5**	**10**	Dev Cd St
18	**16**	**Orange**	**3**	**2**	**0**	**no**	**2**	**2**	**6**	**0**	**4**	**5**	**0**	**0**	**0**	**0**	**6**	
18	**18**	**Red**	**4**	**1**	**5**	**yes**	**1**	**2**	**4**	**0**	**4**	**6**	**2**	**0**	**1**	**3**	**4**	
19	**18**	**Red**	**1**	**1**	**3**	**yes**	**0**	**4**	**4**	**0**	**10**	**10**	**2**	**2**	**0**	**4**	**6**	
19	**17**	**Blue**	**2**	**6**	**0**	**yes**	**3**	**2**	**8**	**1**	**8**	**10**	**3**	**0**	**1**	**4**	**10**	
19	**19**	**White**	**3**	**1**	**5**	**yes**	**0**	**3**	**3**	**1**	**11**	**12**	**1**	**1**	**2**	**4**	**6**	
19	**16**	**Orange**	**4**	**2**	**0**	**no**	**4**	**1**	**9**	**0**	**2**	**6**	**0**	**0**	**0**	**0**	**9**	

Appendix 3: Example Recorded Game

Roll	Blue	White	Orange	Red	Actions
Open	2O, Wo	O, S, W	W	O	
6	O	O	W	O	
8	S	O	0	0	2DEV
7	-	-	-	-	Rob(P)<-1
9	**0**	**0**	**B, W**	**W**	**4DEV**
11	0	S	0	O	1x3(O-W); 1Dev
5	B	W	W	0	2x3(O-W); 2Dev
10	Wo	Wo	Wo	Wo	3Rd, 3City
8	**S**	**O**	**0**	**0**	**4Kt(Rob(O)<-2)**
8	S	O	0	0	1Rd
2	0	0	0	B	2Kt(Rob(G)<-2)
6	O	O	2W	O	3x1(W-S)
9	**0**	**0**	**2W, (R)**	**W**	**4Rd**
8	S	O	0	0	1Dev
5	B	W	W	B	2x3(O-W); 2City
8	S	O	0	0	3x2(W-Wo)(R); 3x2(2W-Wo)
6	**O**	**2O**	**2W**	**O**	**4City**
3	0	0	2Wo	S	1YOP[2W]; 1Sett; 1Dev
3	0	0	2Wo	S	2City
6	O	2O	2W	O	3x(4W-B); 3Sett; 3x2(W-O); 3x1(Wo-O)
8	**2S**	**2O**	**0**	**0**	**4Dev**
4	0	2S	0	W	1x4(S-W)(R); 1Kt(Rob(M)<-3);
10	Wo	2Wo	Wo	2Wo	2x3(O-W); 2Kt(Rob(A)<-1);
6	O	2O	2W	O	3City;
6	**O**	**2O**	**2W**	**O**	**4YOP[2B]; 4Sett; 4Rd**
8	2S	2O	0	0	1x(4O-B); 1Rd;
12	Wo	0	0	0	2x(4Wo-B); 2Rd; 2x3(O-Wo); 2x(4O-B)
9	0	0	2B, 2W	2W	3x1(W-S); 3(Rd*2); 3x2(W-Wo)(R); 3x(2W-B)
7	**-**	**-**	**-**	**-**	**Rob(F)<-2; 4Dev**
4	0	2S	0	W	1x2(2S-B)(R);

5	B	2W	2W	B	2Sett; 2Dev
10	Wo	0_R	0_R	2Wo	3x1(W-Wo); 3Sett
8	**2S**	**2O**	**O**	**0**	**4Kt(Rob(E)<-3); 4Rd;**
5	B	2W	2W	B	1Sett; 1x2(B-2O); 1Dev;
11	O	2S	0	0	2Dev; 2Kt(Rob(G)<-O); 2(Lg Amy)
9	0	0	2W, (R)	2W	
5	**B**	**2W**	**2W**	**B**	**4x1(W-S); 4Sett**
4	0	2S	0	2W	1x2(B-W); 1x(2S-O); 1City, pts 6
6	O	2O	2W	2O	2x1(W,S-O); 2x3(W,S-O); 2x(4O-2W); 2(Rd Bld); 2(Rd*2); 2(Lg Rd); pts 10

Appendix 4: IRB-SBS Oral Consent Text

As you know, I am an engineer from the University of Virginia. I am conducting a study on efficient resource allocations, using Settlers of Catan as a framework, and I would like to ask you some questions about that. I will not collect any personally identifying information, and all data I collect will be written down, but I will relay back what I write down to assure accuracy of data. I will not reveal the content of our conversation beyond myself and people with whom you are playing whom I trust to maintain your confidentiality. I will do everything I can to protect your privacy, but there is always a slight chance that someone could find out about our conversation. Now I would like to ask you if you agree to participate in this study, and to talk to me about Settlers of Catan and your strategic thoughts and actions. Do you agree to participate?

Appendix 5: End of game questions

Game Number:

Date (year/month/day):

Net ID:

Player Number:

What colour were you?

What was your position in the game?

What was your starting strategy this game (1 wood-brick, 2 ore-wheat, 3 blended, 4 rare resource, 5 abundant resource, 6 number diversity)?

Did you intend to overlap strategies, playing more than one at a time, if so, which did you use? If not, write n/a.

Did you change strategies mid-game? What prompted this change?

How many victory points did you have end of this game?

What was the composition of your VP holdings (how many settlements, cities, VP cards and which trophies did you hold if any)?

- Settlements
- Cities
- Trophies
- Knights
- VP Cards
- Action Cards
- Total Road Segments
- Longest Road Built.