

An Evolutionary Approach to Strategies for the Game of Monopoly[®]

Colin M. Frayn

CERCIA

School of Computer Science

University of Birmingham,

Edgbaston, Birmingham, UK

B15 2TT

<mailto:cmf@cercia.ac.uk>

Abstract- The game of Monopoly[®] is a turn-based game of chance with a substantial element of skill. Though much of the outcome of any single game is determined by the rolling of dice, an effective trading strategy can make all the difference between an early exit or an overflowing property portfolio. Here I apply the techniques of evolutionary computation in order to evolve the most efficient strategy for property valuation and portfolio management.

1 Introduction

Monopoly[®] is primarily a game of skill, though the short-term ups and downs of an individual player's net worth are highly dependent on the roll of a pair of dice. As such, it separates itself from completely skill-controlled games such as chess and Go, where no direct element of unpredictability is involved except that of guessing the opponent's next move. Despite the element of chance, a strong strategy for which properties to purchase, which to develop and which to trade, can vastly increase the expected results of a skilful player over less knowledgeable opponents.

There are many parallels here with real life, where a wise property investor, though largely subject to the whims of the property market, can increase his or her expected gains by a process of shrewd strategic dealing. Much of the skill is involved with appraising the true value of a certain property, which is always a function of the expected net financial gain, the rate of that gain and the certainty of that gain.

The game considered in this work is as faithful to the original rules as possible. Players take it in turns to roll two dice, the combined total determining the number of squares over which they move. Players may acquire new assets either by purchasing available properties on which they randomly land, or else by trading with other players for a mutually agreeable exchange price. Rent is charged when a player lands on a property owned by another player, varying based on the level of development of that particular property. In this study, we use the property names of the standard English edition of the game.

Much anecdotal evidence exists concerning the supposed "best strategies", though very few careful studies have been performed in order to gain any quantitative knowledge of this problem. Because of the inherently stochastic nature of the game, "best strategies" are often described without a sufficient statistical foundation to support them.

In 1972, Ash & Bishop performed a statistical analysis of the game of Monopoly[®] using the method of Markov chain analysis. They evaluated all the squares on the board in order to determine the probability of an individual player landing on each square in any one turn. Furthermore, they gave an expected financial gain for every roll of the dice, given the property ownership situation. The results of this study showed that the most commonly visited group of properties was the orange street consisting of Bond Street, Marlborough Street and Vine Street.

This analysis gave some insights into suggested strategies for the game. For example, encouraging the acquisition of regularly-visited properties. A simple strategy built on a re-evaluation of property value based on the expected gain per turn could well pay dividends. For example, in the standard rules, it takes on average 1400 opponent turns to pay for the purchase cost of Old Kent Road (if unimproved), but only 300 to pay for Mayfair (Ash & Bishop, 1972). One might argue therefore that Mayfair is nearly 5 times under-priced compared to Old Kent Road.

However, there is much more to this than a simple statistical evaluation. For example, when improved to a hotel, both the properties mentioned above (Old Kent Road and Mayfair) take approximately 25 opponent turns to repay their own development costs. So we must find a fair value for these two properties which considers not only their expected (pessimistic) time to repay development costs, but also the potential gain should we be able to purchase all members of the same colour group, and the potential cost of then developing those properties to the required levels. It should also consider other factors such as the mortgage value for each property, strategies for paying to leave jail, when to develop, when to mortgage (and un-mortgage) and how to handle bidding wars.

Clearly we need a more advanced method of obtaining fair prices for all the properties on the board, based on one's own state and that of the opponents.

In this work, I investigate an evolutionary approach to the game of Monopoly[®]. I propose a scheme for representing a candidate strategy (section 2), and present the results a considerable number of games using both a single- (section 3) and multiple-population (section 4) approach. I conclude with the lessons learned from this study (section 5) and the scope for future investigation (section 6).

2 Evolutionary Approach

Evolutionary computation can be applied to the problem of strategy design in the game of Monopoly[®]. It allows the simultaneous optimisation of a very large number of interdependent variables, which is exactly what is required in order to develop coherent fair-price strategies for such a complex environment.

In the case of Monopoly[®], each individual in the population represents a different set of strategies which can be used to make the various decisions required in the game. The representation used in this work consists of four distinct elements;

- Valuations for each property on the board.
- Valuations for each property when held as a member of a street.
- Valuations for the extra value of a property based on the number of houses built on it.
- Extra game-related heuristic parameters.

The first three elements are self-explanatory, though the fourth requires some elaboration. In order to generate a list of required parameters, it was necessary to consider all the decisions made by a human player during the course of a game, and to decide how to encode those decisions as parameter values. The final list was as follows:

- Parameters concerned with whether or not to pay to exit jail. This was modelled as a linear combination of the maximum and average estimated opponent net worth and house counts.
- Parameters concerned with the valuation penalty applied to mortgaged properties, depending on whether they are members of complete streets or not.
- Parameters governing a desired minimum cash position, based on average and maximum opponent net worth and number of houses.

All parameters were stored as floating-point values, and were initialised with random perturbations about the following defaults:

- Properties are worth 1.5 times their face value, but 4 times if members of a street.
- House values are twice their development cost.
- Stay in jail if (maximum opponent worth) + (average opponent worth) + (10*number of houses on the board) is greater than 10000. Otherwise, pay to exit.
- Keep a minimum of 200 pounds in cash, plus 1% of the total and average opponent net worth, plus 5% of the number of houses or hotels.

The exact choice of default values here made no difference to the outcome of the simulation, except that outrageously unsuitable values would cause the evolution process to take longer to settle down to a stable end state.

A detailed interface was also designed, incorporating all the rules of Monopoly[®]. A few slight alterations were made in order to make the game easier to deal with.

- When a player becomes bankrupt, his or her properties are all returned to the bank, instead of auctioning them (which tends to reward the players who happen, by chance, to have a lot of spare cash at that particular time.) In later work we shall use a standard auction at this point instead, to check if this affects the behaviour. It is possible that, by using auctions, we might instead encourage strategies involving more prudent use of resources so that such events might be exploited more effectively.
- Chance and Community Chest cards were picked randomly with replacement, instead of remembering a random initial card order and cycling through these. 'Get out of jail' cards were tracked, and not replaced until used.
- There was no maximum on the number of houses or hotels allowed on the board simultaneously. It is not clear if this affected the strategies, avoiding the need for property strategies dealing with housing shortages.
- Games were limited to 500 turns. If there was no clear winner at this point, then the players were ranked by total net worth.
- Properties not purchased immediately were not auctioned, but remained unsold. See section 6 below for a discussion on this point.

Calculations were run on a 3GHz Pentium-IV machine. For simplicity, all games were started with four players. After careful optimisation, games could be simulated at the rate of approximately 400 per second. This project involved a total of over 377 million games of Monopoly[®], one quarter of which (87 million) are included in the final results.

3 Single Population Results

As an initial test of the evolutionary algorithm approach, I generated a single population of individuals, and ran a

standard evolutionary algorithm, using a population size of 1000 and 24 hours of CPU time. This resulted in a total of 1420 generations completed. In each generation, 100 iterations were played. For each iteration, the 1000 individuals were selected off randomly into groups of four, and each group played one single game of Monopoly®. The game ended when there was only one player left, or after 500 turns. Points were awarded to the individuals based on their position in each game. First place was awarded 4 points, second place was awarded 2 points, third place 1 point and last place 0 points.

The fitness function, therefore, consisted of the sum of the points gathered by an individual over the 100 iterations in each generation.

At the end of each generation, the top three individuals survived by right as elites. 300 survivors were selected using a size 2 tournament selection algorithm with replacement. A further 300 individuals were selected in the same manner to continue to the next generation after undergoing random mutations. During a mutation, ten values from each of the property prices, street prices and house prices were randomly mutated using a Gaussian kernel of standard deviation 10% of the variable's value. All other values were mutated by the same amount with a probability of 50%.

The remainder of the next generation (397 individuals) were generated using a crossover between two tournament-selected parents from the current generation. During crossover, the child acquired each parameter randomly from either parent with equal weighting.

Previous experience with evolutionary algorithms has taught us that the solutions derived in most applications are not very sensitive to these values, and that the above values lie within sensible ranges. Alternative values for the elitism fraction, mutant fraction, crossover fraction etc. were not tested.

At the end of each generation, the best individual was tested in 1000 games against randomly generated opponents, with the same scoring system as detailed above. This was used as a check to ensure that the algorithm was indeed moving towards greater fitness. These testing figures are shown in figure 1. They show a very sharp rise over the first 20-30 generations, from an initial score of 2522. From generation 30-40 up to approximately generation 200, the test scores then slowly declined, before levelling off around 2930. The plots in this paper show the result from a single trial. However, multiple trials were performed during testing, with subtly different algorithmic details, and very similar results were achieved each time.

A speculative interpretation of this behaviour is linked to the manner in which strategies evolve within such a complex evolutionary environment. The population rapidly learned some strong, simple strategies which could be used to good effect against simple, randomly generated opponents. However, after the first few dozen generations, individuals began to develop counter-strategies which partly refuted these 'easy wins'.

Because the opposition from other individuals in the population was now much higher, the simple strategies had to be abandoned, leading to worsening performance against random opponents, but better performance against other members of the population, against whom the fitness function was measured.

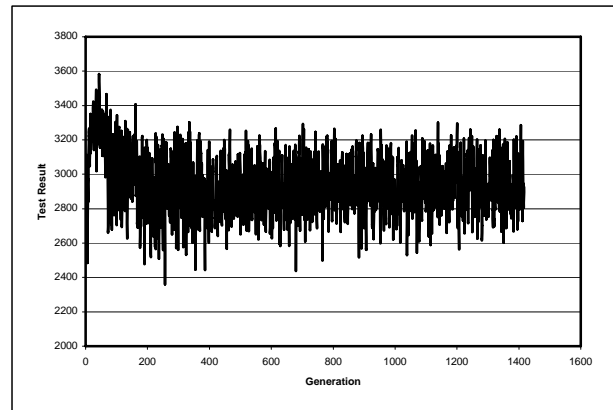


Figure 1 : Test score of best individual as a function of generation (single trial)

This style of learning is common to a wide variety of games. For example, a novice chess player might quickly learn some clever tricks (such as “fool’s mate”) by which he or she can defeat rather inexperienced players. By using such tricks, the novice begins to win a significant number of games because his or her opponents fall for these simple traps.

However, soon the opponents wise up to this strategy, and the novice player can no longer use the same tactics. In fact, these simple ‘trap’ openings often prove rather weak if the opponent knows how to deal effectively with them. If the novice then finds himself facing an unknown opponent then he will not use these same tricks any more, instead using more advanced opening strategies with which he may be far less confident. Against a good player, this will be a better strategy, but against a true novice, using a trick strategy might have given a better chance of winning.

After the end of the run, the best individual from each generation was studied in order to evaluate the degree of learning that had occurred. It was possible to examine how the estimated values of the individual properties and the various other numerical values used in the winning strategies had evolved over time.

Figure 2 shows the change in the estimated property value for “Old Kent Road”, the square immediately after “Go” and the least valuable square on the board, according to face price. In Monopoly®, Old Kent Road is valued at £60, and the individual houses cost £50 to build once a player also owns Whitechapel Road, two squares further along.

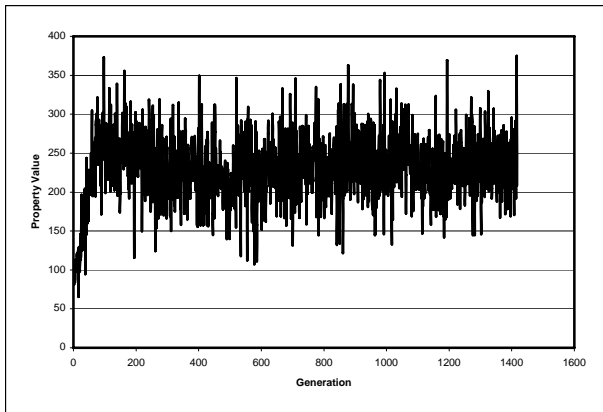


Figure 2 : Perceived value of Old Kent Road as a function of generation

By the end of the run, the average estimated worth of Old Kent Road, averaged over the last 100 generations, was £238, meaning that the evolutionary algorithm valued this property at a mark-up of approximately 297%. Whitechapel Road was valued at £290, or a mark-up of 383%. This property was valued slightly higher because it offers substantially more rental income, especially once developed. Based on individual rent prices, these properties were therefore valued on a forward per-earnings multiple of 119 times and 72.5 times respectively. Clearly, these prices would therefore only be worth paying if the player could expect to own the entire street and develop it with houses, or else prevent an opponent from doing the same.

Figure 3 shows a summary of the average property value for the ownable properties, versus their nominal face value. If a property is valued at less than its face value then the individual player will never purchase that property directly (though it might buy it in a trade from another player for a lesser amount). Note that every property along the lower and left-hand sides of the board appears to be undervalued, and almost all of the remaining properties appear to be overvalued, often enormously so. The only properties that the computer player would buy after “Free Parking” are the two stations (Fenchurch Street and Liverpool Street), the water works and the dark blue properties (Mayfair & Park Lane).

Figures 4 and 5 show the rapid reduction in perceived values of Strand (red property) and Bond Street (green property) over the simulation run. After 1420 generations, the algorithm values these two properties at £90 and £65, at a net discount to their face values of 59% and 80% respectively. During testing with smaller populations, or subtly different selection procedures and fitness functions, I obtained an extremely similar result every time.

For some reason the end result appears to be that the evolved Monopoly® players dislike the red, yellow and green streets. They will never buy a new property on any of these streets. The reason for this is difficult to discern, but it is such a pronounced effect that I suspect

that it is due to the house cost for the upper and right-hand sides of the board. When houses cost £150 or £200 each then it gets very difficult to develop the red, yellow and green streets unless you are already winning by quite a considerable margin. And if you’re already winning then you needn’t bother developing new streets.

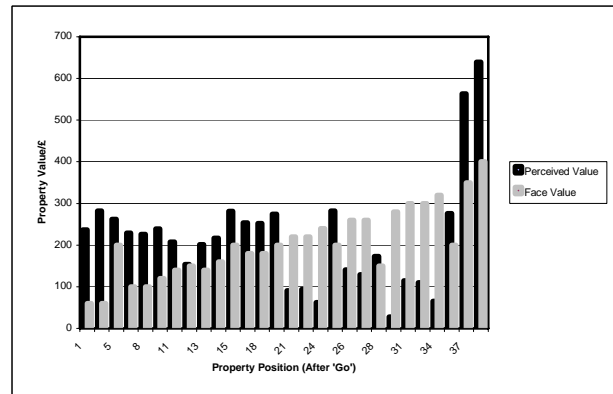


Figure 3 : Perceived value versus face value for all purchasable properties

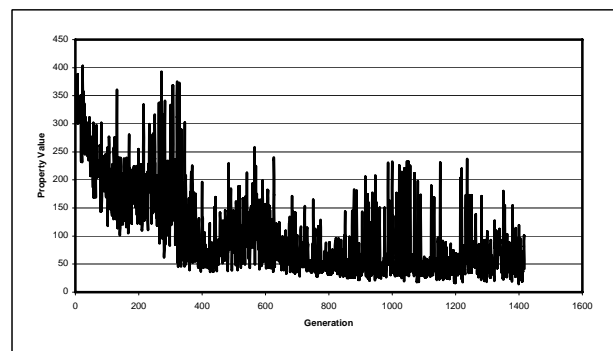


Figure 4 : Perceived value of Strand as a function of generation

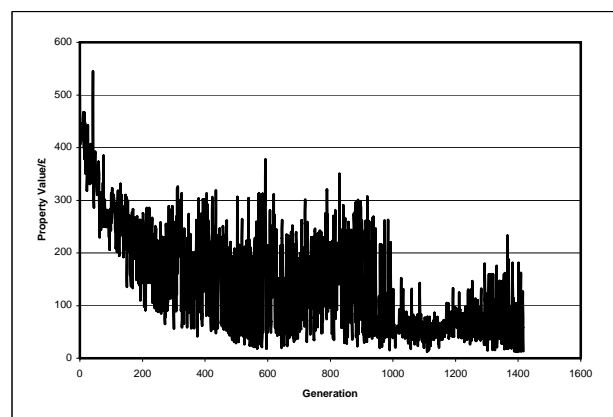


Figure 5 : Perceived value of Bond Street as a function of generation

So the conclusions drawn from this single population experiment seem to show that the best strategy is to gather

the lower value streets as quickly as possible, develop them rapidly and aim to win quickly.

4 Twin Population Results

In order to test these results, I implemented a multiple population approach to test to see whether the results from a single population strategy held up when two or more distinct populations evolved separately with only a very small trickle of individuals exchanging between them.

For this section, I implemented a two population approach with a migration rate of 0.5% at the end of every generation. Each population was set up exactly as the single population above (1,000 individuals, randomly seeded, fitness function and breeding as above). The simulation was run for 800 generations, and the results compared both between the two populations, and also back to the original single population.

The variations between the two parallel populations at the end of the simulation were found to be extremely minimal. When compared to the single population, the variations were slightly larger, but still the results were largely the same. Figure 6 shows the difference between the property valuations in the single population and multiple population runs. The differences between the two populations in the multiple-population simulation were so small that I have just plotted the first population results here.

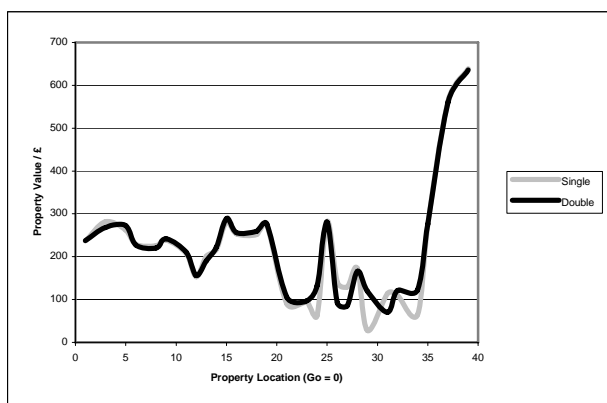


Figure 6 : Comparing final property prices between single population and multiple population simulations

The first thing to note about figure 6 is that the two simulations gave exactly the same results for all of the streets whose perceived value was greater than their face value. Figure 7 shows this feature.

In figure 7, the x-axis represents the ratio of perceived property value divided by the face value. Properties that were perceived to be undervalued on their face value are therefore towards the right in this diagram. Properties which were deemed less valuable than their face price (and therefore would not be bought) are at values less than one. The y-axis represents the logarithm of the disagreement between the perceived property value

derived by the single-pop and 2-pop simulations, as a percentage of the single-pop perceived value.

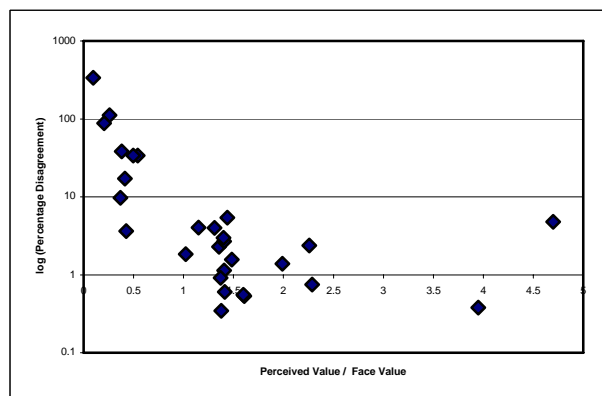


Figure 7 : Disagreement between single-pop and 2-pop simulations as a function of perceived property premium

This plot shows that the largest disagreements by far are caused by those properties whose perceived value was less than their face value. That is to say, those properties which the computer would never buy when it landed on them. In these cases, the perceived property value is useful only for bartering and dealing between players, and the likelihood of any player owning a street of this colour property would be very small. Thus, the variation in the perceived values was occasionally very high. However, for the properties whose perceived value was greater than the face value, the two simulations converged to remarkably similar estimates.

The other result to note is that the evolutionary algorithm, as expected, values the most expensive member of each colour group slightly higher than the other members of that colour group. The perceived value of station properties, all with a face value of 200 pounds, was also slightly variable. Marylebone was, as expected, the most valuable of the four, with an estimated value of £285 taken as the average of the single-pop and 2-pop simulation valuations. This is because there is a chance card moving players to Marylebone without choice. Next came Fenchurch St. At £281.40, King's Cross at £276.50 and finally Liverpool St. At £276.10. However, the variation between the prices here was not large, and there is insufficient evidence to suggest that these prices vary at all from a universal valuation of between £275 - £280.

5 Other strategies

Together with estimated house values, the genome for an individual also contained estimated price premiums for owning a property as a member of an entire street, and developing it with houses. Figure 8 shows the premium, that is the multiple of the basic perceived value, for owning a property as part of a street instead of singly.

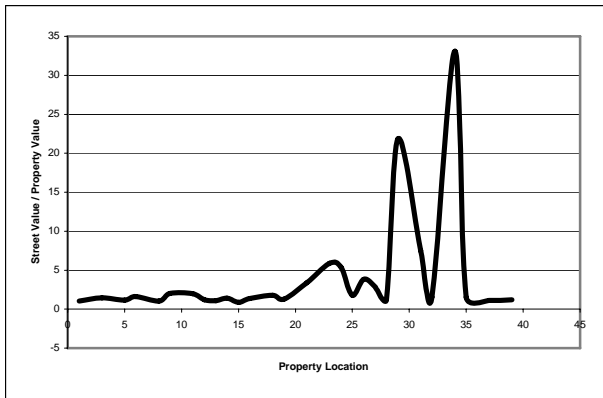


Figure 8 : Street ownership premium for all properties

As is clear from figure 8, all properties up to and including 19 (the lower and left-hand sides of the board) operate at very modest premiums to their individual value. However, the properties of higher face value (that is the red, yellow and green streets, though not the dark blue street, the stations or the water works) have a much higher street value compared to their individual value. This is particularly striking for the most expensive properties in these three zones, namely Trafalgar Square, Piccadilly and Bond Street, which operate at 5, 21 and 33 times their individual values respectively.

This inflated value for these properties as a member of streets reflects their apparently low value as single properties. The evolutionary algorithm grows to dislike these properties, though obviously once two are acquired, it becomes favourable to gain the third. As the chance of the algorithm acquiring two of the properties is very low, this adaptation is probably more of a defensive measure rather than anything else – stopping opponents from gaining the streets rather than aiming to build on those properties itself.

The final set of values used in the genome concern particular financial strategies necessary for accurate play. These values varied enormously between the runs, tending to hint that they were largely irrelevant to the overall performance of any one individual. However, a few rather general conclusions could be drawn.

- (1) It is wise to retain approximately 110-150 pounds in your bank account as a bare minimum. Add to this approximately 5% of the net worth of the strongest opponent.
- (2) The value affecting the minimum amount of money to retain is much more strongly linked to the net worth of the strongest opponent than to the average or total net worth of the players on the board.
- (3) It is almost always a wise idea to pay to get out of jail. If you get to the point when you are staying in jail to avoid paying rent then you've probably lost anyway! However, there are also times when staying in jail can allow you to collect considerable rent from opponents landing

on your properties, without the risk of you yourself being fined.

- (4) Avoid accepting a mortgaged property for a trade unless it makes up a new complete street.

6 Conclusions

This genetic algorithm approach to playing Monopoly® has given a variety of insights into the game. Much of what the simulations discovered has been known for some time, though it is always reassuring to confirm this. However, some strategies are completely new.

In most games, landing on *any* property with a hotel will cause a considerable dent in a player's net worth. Doing this twice will probably spell the end of the game. Therefore, it makes sense to concentrate on the properties that are cheapest to develop, so that you can reach a high level of rent-gathering as rapidly as possible.

For example, for the red properties, reaching the level where you can charge an average rent of nearly £300 would cost £1580 (purchasing all three properties, plus two houses on each – average rent £267). For the brown properties, this only costs £620 (buying two properties and a hotel on each – average rent £350).

With the orange properties – which are the most frequently visited on the board – £1460 can buy you all three properties, plus three houses on each – charging an average rent of £567 pounds. Not only is this £120 cheaper than developing the red properties as above, but it also gives a rent of well over twice the amount. Moreover, these properties are more frequently visited, therefore making the developed orange properties a vastly superior investment. A fine of 567 pounds would considerably dent all but the strongest of opponents.

In addition to the property valuation strategy, three further tips arose from the best evolved strategies.

Firstly, always retain a small reserve of cash to stop you from mortgaging properties. Mortgaging can be useful, but ultimately you are stifling a revenue source, which tends to drop you further back in the game. The penalties derived for mortgaging were very steep – with mortgaged properties sometimes worth as little as 2% - 3% of their perceived un-mortgaged value.

Secondly, don't be afraid to make bids for opponent properties. Human beings often vastly undervalue the cheaper streets – giving an astute player a certain strategic advantage if he or she can initiate a favourable trade. Single, expensive properties can be very useful indeed if they are traded for less expensive properties, even at a substantial concession to their face value.

Thirdly, don't be a coward and stay in jail – fortune favours the bold! Saving 50 pounds in the short term could well cost you the opportunity to pick up on a vital deal later on.

One potential extension to this study is to vary the maximum game length. Setting this well below the

expected survival time for the three losing players (say, 50 turns) would encourage strategies which accumulated wealth very rapidly, but perhaps not in a stable way. This is well worth investigating. Figure 9 shows that most games are either complete by approximately turn 200, or last the full 500 turns. Any game surviving several hundred turns is likely to be in one of two states: either (1) oscillations in power between two or more players, so that the eventual winner is largely random or (2) a stalemate where no player owns any streets nor wants to sell any. It is likely that reducing the maximum game length to, say, 250 turns will not greatly affect the results. However, as the overwhelming majority of games are completed by this stage, the saving in CPU time will be barely noticeable.

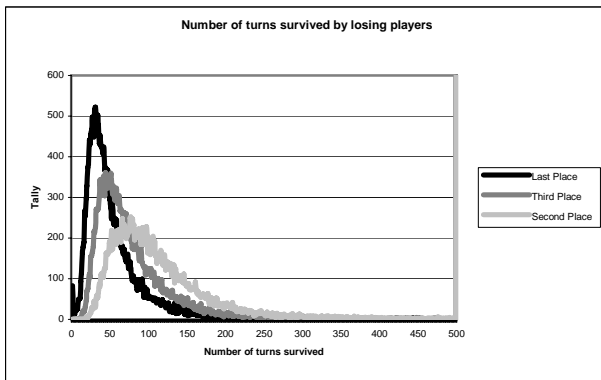


Figure 9 : Average survival period for the losing players in one generation. Note the large spike at 500 turns for games which lasted the full maximum duration.

As an extension, I am investigating the effects of completing the implementation of the full realistic rules. During this study, I used a simplified subset of the rules as I believed that any mild affect on the strategies developed would be more than generously offset by the reduced programming complexity and the increased number of generations that could be run.

Subsequent study has tentatively suggested that the introduction of a more realistic rule set might affect the rules more strongly than I had predicted. Most importantly, the introduction of a full auction system appears partially to prevent the perceived reduction in value for the more expensive properties, and also tends to shift all perceived values upwards. However, such an implementation slows the game speed down considerably, and reduces the learning rate, so a considerable amount more processing time is required in order to investigate this effect more thoroughly. After careful optimisation, and using most of the full rules, the games are running five times slower than reported in the present work. I hope to release a follow-up paper in the future investigating the effects of these changes.

Acknowledgments

I acknowledge Hasbro, Inc. who own the trademarks and publishing rights for the game of Monopoly®. I also

acknowledge Advantage West Midlands, who supplied the funding for my research position and for the formation of the Centre of Excellence for Research in Computational Intelligence and Applications.

Bibliography

Ash, Robert B., and Bishop, Richard L. 1972. "Monopoly® as a Markov process." *Mathematics Magazine* 45:26-29.

Stewart, Ian. 1996. "Monopoly® revisited." *Scientific American* 275:116-119

Stewart, Ian. 1996. "How fair is Monopoly®?" *Scientific American* 274:104-105