Experimental Investigations in Ramsey Avoidance Games: Strategy, Structure, and Scale

Sambhav Dhakal, Joseph Mathew
Mentor: Dr. William Gasarch

August 25, 2025

Abstract

This research paper describes an experimental study of the two-player Ramsey avoidance game, whereby players color edges of a complete graph K_n. One player loses the game upon completing a monochromatic K_m in their own selected color. In this sense the game is designed so that players will force the opposing player to create a forbidden clique. The experiments consist of simulating two algorithms: the composite heuristic agent (Force+Avoid+Random) and the Minimax agent (with depth 2), against itself for several small concentrated configurations of (n, m). The number of empirical proportions is reported with 95% Wilson confidence intervals to ensure statistical fidelity. The empirical reports indicate a substantial performance difference between agents, and the study investigate performance distinctions based on three investigations: (1) the relationship between n and advantage of second player, (2) the affordance of draws for games with large clique size m≥3, and (3) the sensitivity of heuristics with the complex states of the game. The Minimax agent proves to achieve a deterministic win for the second player in all games studied regarding (n, 3). The research ends contemplating the application of reinforcement learning to broaden the research into computationally infeasible graphs.

1 Introduction

An elementary example in Ramsey theory is $R(3, 3) = 6$, which states that no two-coloring of the edges of K_6 will avoid creating a monochromatic triangle [1, 2]. Gasarch's review compares how various authors present the foundational results and differences in approach [2]. Although the theorem is existential, sequential adversarial coloring provides an intriguing opportunity for a game theoretic model. We look at the Ramsey avoidance game where a player tries to force the opponent to create a monochromatic forbidden clique, K_m, with their color. We analyze the game using two different automated strategies: a rules-based composite heuristic (F+A+R), and a search-based Minimax agent with depth=2. To collect data, we run these agents against each other across various small, computationally feasible game scenarios, such as (6, 3) and (7, 3) games, to identify strategies and see how the structure of the game influences win probabilities.

2 Methodology

2.1 Agent Implementations

All experiments were conducted in a custom simulator for the two-player Ramsey avoidance game. Two distinct agent architectures were implemented and tested.

Composite Heuristic (F+A+R): This agent uses a deterministic, rule-based hierarchy to select moves. The agent's goal is to maneuver the opponent into a position where they have no choice but to complete a K m in their own color, thus losing the game. The agent's priorities are as follows:

1. **Force (Offensive):** Due to the fact that the agent is trying to force the opponent into a poor move, the number one priority is to find the move that puts the greatest pressure on the **opponent** by threatening a loss. This pressure is applied by coloring an edge that creates a close-to-complete sub-structure (a near-clique) in the opponent's color. The ForceStrategy function evaluates a constraint score for each possible eligible move based on how many incomplete structures it creates in the opponent's color. The function also rewards moves that put the opponent in immediate danger (i.e., one edge away from a loss) with additional scoring bonuses.
2. **Avoid (Defensive):** In the case that there are no offensive forcing moves, the agent will play a defensive manner. The main aim associated with this move is the agent does not get stuck permanently in a losing trap defined by the opponent. The AvoidStrategy function will compute all of the "safe moves" that will not lead to the ultimate lost movement. The function will subsequently assign risk scores to each safe move as a way of calculating how many almost-complete structures are formed in the own color after the move is played. The **agent** will then select the move with the lowest risk score possible. The score the agent assigns to the move is heavily penalized for moves that place the agent one edge away from a losing structure.

3. **Random:** If neither an offensive forcing move nor a defensive avoidance move is available, the agent selects a legal move uniformly at random from the set of available edges.

Minimax Agent (Depth 2): This agent employs the Minimax algorithm with alpha-beta pruning to search the game tree. The agent was configured with a fixed search depth of two plies (its own move and the opponent's subsequent reply). At the search horizon, a simple utility function assigns +1 for a winning state, -1 for a losing state, and 0 for all non-terminal states. A critical implementation detail is its **deterministic tie-breaking**: when multiple moves are evaluated as equally optimal, the agent consistently selects the one corresponding to the lexicographically first edge.

2.2 Experimental Setup and Reproducibility

To ensure the results can be replicated, we provide the following details on the computational environment and experimental protocol.

Environment:

- **Software:** The simulations were executed using Python 3.12.1. The implementation relies on standard Python libraries.
- **Hardware:** The experiments were run on a desktop computer equipped with an Intel Core i7-11700 CPU and an NVIDIA GeForce RTX 3060 Ti GPU. It should be noted that the simulation is primarily CPU-bound, so the specific GPU model is not a critical factor for performance.

Protocol: The experimental protocol for each tested (n,m) configuration was as follows:

1. **Matchups:** Both symmetric, "strategy vs. itself" pairings (i.e., F+A+R vs. F+A+R and Minimax vs. Minimax) and asymmetric pairings (Minimax vs. F+A+R) were run.
2. **Sample Sizes:** For all matchups, a total of N=1000 independent games were simulated for each configuration.
3. **Data Collection:** For each matchup, the number of wins for Player 1, wins for Player 2, and draws were recorded.
4. **Statistical Reporting:** The raw counts were converted into proportions. We report these proportions along with their 95% Wilson confidence intervals to quantify statistical uncertainty. The source code for the simulator and analysis is available upon reasonable request to the author.

2.3 Experimental Design and Statistical Reporting

The core of our investigation consists of running each agent against an identical copy of itself. This symmetric setup allows us to measure the inherent advantages or disadvantages of playing first (P1) or second (P2) for a given strategy.

- For the **F+A+R vs. F+A+R** matchups, a sample size of N=1000 games was used for each (n,m) configuration.
- For the computationally more intensive **Minimax(d=2) vs. Minimax(d=2)** matchups, a sample size of N=200 games was used.

The primary outcomes recorded were the win rates for Player 1 (P1), Player 2 (P2), and the proportion of draws. To quantify the statistical uncertainty of our results, we compute and report 95% Wilson confidence intervals for all observed proportions. This interval is particularly reliable for proportions near 0% or 100%, which occur frequently in our Minimax trials.

3 Hypotheses

Our analysis is guided by the following three hypotheses.

1. **Increasing *n* Reduces Second-Player Advantage:** In the classic (n,3) game, there is often a theorized second-player advantage. We hypothesize that as the graph size n increases, this advantage weakens, and game outcomes should approach parity, or the likelihood of draws should increase, as the strategic complexity grows.

2. **Larger *m* Leads to Draw-Dominated Outcomes:** For games where the forbidden clique size is larger (m≥4), we conjecture that forcing a loss becomes significantly harder. Consequently, near-optimal play will rarely result in a win for either player, and draws will become the most common outcome, especially for moderate values of n.

3. **Heuristic Brittleness in Complex Game States:** We hypothesize that simple, rule-based heuristics like F+A+R are effective in sparse game states but become "brittle" and perform sub-optimally as the board's edge density increases. Deeper tactical threats, which a search-based algorithm like Minimax can detect, are often missed by such heuristics, leading to a significant performance gap between the two agent types.

4 Results and Discussion

The simulation results, summarized in Table 1, provide evidence to evaluate our hypotheses. The data is sourced from the author-supplied spreadsheet shown in Figure 1.

Table 1: Selected outcomes and 95% Wilson confidence intervals.

(n, m)	**Strategy pairing**	**P1 win (%)**	**95% CI (P1 win %)**
(6, 3)	F+A+R vs F+A+R (N=1000)	49.80	46.71 – 52.89
(6, 3)	Minimax (d=2) vs Minimax (d=2) (N=200)	0.00	0.00 – 1.88
(7, 3)	F+A+R vs F+A+R (N=1000)	37.90	34.94 – 40.95
(7, 3)	Minimax (d=2) vs Minimax (d=2) (N=200)	0.00	0.00 – 1.88
(8, 3)	F+A+R vs F+A+R (N=1000)	39.60	36.61 – 42.67
(8, 3)	Minimax (d=2) vs Minimax (d=2) (N=200)	0.00	0.00 – 1.88
(6, 4)	F+A+R vs F+A+R (N=1000)	7.10	5.61 – 8.89
(7, 4)	F+A+R vs F+A+R (N=1000)	0.80	0.35 – 1.57
(8, 4)	F+A+R vs F+A+R (N=1000)	7.50	5.98 – 9.32
(n, m)	**Asymmetric Matchups**	**P1 win (%)**	**95% CI (P1 win %)**
(6, 3)	Minimax vs F+A+R (N=1000)	55.60	52.51 – 58.66
(6, 3)	F+A+R vs Minimax (N=1000)	41.50	38.48 – 44.58

Analysis of Symmetric Matchups (F+A+R vs F+A+R, Minimax vs Minimax): Contrary to our hypothesis, the provided data does not show a clear reduction in the second-player advantage as n increases for (n,3) games. For the **Minimax agent**, the result is stark: P2 wins 100% of games for $n \in \{6, 7, 8\}$, showing a persistent, deterministic advantage. The 95% CI of [0.00%, 1.88%] for P1's win rate provides strong evidence against parity. For the **F+A+R agent**, the outcome is closer to a balanced game, but P1's win rate drops from 49.8% at n=6 to 37.9% at n=7, suggesting the game becomes more difficult for P1 before stabilizing. This indicates the relationship between n and player advantage is not monotonic for this heuristic.

Analysis of Asymmetric Matchups: The results from the asymmetric matchups provide a clear comparison of the two strategies' relative strengths. When Minimax plays as Player 1 against F+A+R as Player 2, Minimax wins a majority of the time with a 55.6% win rate. Conversely, when F+A+R plays as Player 1 against Minimax as Player 2, the Minimax agent's win rate jumps to 58.5%. These results indicate that the Minimax strategy is superior to the F+A+R heuristic, as it consistently achieves a higher win rate regardless of whether it is Player 1 or Player 2. However, the data also shows that the positional advantage of being Player 2 is significant. When Minimax plays as Player 2 against the same F+A+R agent, its win rate improves by nearly 3 percentage points (from 55.6% to 58.5%), while the F+A+R agent's win rate drops from 55.6% to 41.5% when it is no longer paired against a symmetric opponent.

Analysis of Hypothesis 2 (Draw Domination): The results for m=4 offer strong support for this hypothesis. For the **F+A+R agent**, the outcomes are overwhelmingly draws: 92.9% for (6,4), 77.3% for (7,4), and 62% for (8,4). Wins are rare, aligning perfectly with the conjecture that forcing a loss is difficult. Data for the Minimax agent on these configurations also showed high draw rates ((6,4) and (7,4) were 100% draws in our tests). This suggests that for m $\geq$ 4, the game space is dense with paths that lead to draws when neither player makes an unforced error.

Analysis of Hypothesis 3 (Heuristic Brittleness): The performance disparity between F+A+R and Minimax in the (n,3) games is strong evidence for this hypothesis. While F+A+R plays to a near-even match, Minimax(d=2) identifies a clear, winning path for P2 in every single game when playing against an identical Minimax agent. This suggests that the heuristic, by focusing only on immediate threats (completing or blocking a $K_(m-1)$, misses deeper two-move tactical sequences that the search-based agent can see. As the game progresses and edge density increases, the number of such subtle threats likely grows, making the heuristic's simple rules increasingly inadequate. The Minimax agent's ability to consistently exploit this demonstrates the brittleness of the F+A+R strategy.

Why does Player 1 never win with Minimax vs Minimax? One striking outcome is that when both players use the Minimax (d=2) agent, Player 1 never wins. Player 2 always forces a win in every configuration tested. While our experiments confirm this as an empirical fact, we cannot fully explain why it occurs. It is possible that the deterministic tie-breaking rule, combined with the short search depth, produces a consistent second-player advantage. However, it is also possible that there is an inherent property of the Ramsey avoidance game that grants Player 2 a structural advantage under near-optimal play. To our knowledge, this remains an open question and is worth future investigation.

5 Limitations of the Current Study

The results of this study are indeed useful, but, as always, there are limitations to highlight.

1. **Fixed and Shallow Search Depth:** The Minimax agent was limited to a search depth of two, which restricts its ability to anticipate longer tactical sequences. A deeper search would better approximate optimal play but quickly becomes computationally expensive.

2. **Deterministic Tie-Breaking:** The Minimax agent used lexicographic tie-breaking when multiple moves were equally optimal. This introduces bias by favoring one path through the game tree. A simple alternative would be to replace this with random tie-breaking, which could reduce the bias. We did not implement this here but suggest it as an easy extension.

3. **Heuristic Agent Simplicity:** The F+A+R agent was created for simplicity and does not adapt or learn. Its fixed rules cannot capture deeper strategies.

4. **Computational Infeasibility for Larger Configurations:** Scaling to larger n, m values or deeper Minimax trees becomes computationally infeasible with our current resources. This limits the scope of our experiments to small graphs.

6 Future Work

Building on the current study, future research could proceed in several promising directions.

From the existing study, there are many suggestions for future work. The main downsides to the current study have to do with the simplicity of the heuristic agent and the computational expense of the Minimax tree search algorithm. Testing hypotheses on larger, more theoretical maneuvers (e.g., (17, 3) or (16, 4)) will require a more advanced agent architecture, which can mean progress towards an agent that could learn better. One potential direction would be to develop a Reinforcement Learning (RL) agent similar to the AlphaZero systems, which would be a combination of Monte Carlo Tree Search (MCTS) and a neural network which learned to evaluate the value of the board position and suggest promising moves. In particular, the RL agent would play against itself and discover long, nuanced strategies that were probably not reachable by a human coded heuristics strategy or search depth.

A future opportunity would be to investigate asymmetric matchups which, in general, could help to tease out deeper insights into the relative strengths of various strategies. For example, if the Minimax agent was tested as Player 1 against the F+A+R agent as Player 2, this would determine if Minimax's strategy superiority could mitigate for a potential advantage of the second player (which could be inherent). In general, based on the heuristic agent's brittleness, we would expect Minimax to perform quite well.

To conclude, we would benefit from a more exhaustive study of the heuristic agent itself. Ablation studies looking at "F+R" (Force+Random) or "A+R" (Avoid+Random) agents would allow us to tease apart and measure the value of forcing and avoidance separately, i.e. how much is each contribution used relative to the total. We would also be able to compare these simpler heuristics to the total F+A+R agent and the Minimax agent to better understand the strategic value of each.

This mixed-method strategy with more powerful RL agents, asymmetrical matchups, and ablation studies of heuristics represents an important next step for taking us from empirical observation on small graphs to valid conjectures about the nature of these games at a larger scale.

7 Conclusion

The experimental work showcases a few underlying structural characteristics of the Ramsey avoidance game. A depth-2 Minimax agent achieved a near-sure second-player win in small size games, a result not found with the simpler heuristic agent evaluated. This implies the fragility of rule-based strategies in situations requiring tactics. Further, the data demonstrated that larger sizes of forbidden cliques are very likely to produce draws. Not every hypothesis was supported in its original version, but the results provide clarity on the relationships between graph size, agent complexity, and game outcomes. The results from the asymmetric matchups highlight that while the Minimax agent is the superior strategy, the positional advantage of playing second is still significant and can be decisive in symmetric matchups. Further advancements will require the application of advanced techniques such as reinforcement learning, or exploring matchups between asymmetric agents, to study further on larger scales.

References

[1] Beck, J. (2008). *Combinatorial games: Tic-tac-toe theory.* Cambridge University Press.

[2] Gasarch, W. (2023). Joint review of Ramsey theory books. University of Maryland. Retrieved from `https://www.cs.umd.edu/~gasarch/bookrev/NICK/ramseyall.pdf`