# Modeling the Winkum Game

Istar Abdullahi
*Computer Science (Grinnell College)*
Grinnell, USA
email: abdullah2@grinnell.edu

Erin Jarvis
*Computer Science (Grinnell College)*
Grinnell, USA
email: jarviser2@grinnell.edu

Patrick Sales Garcia
*Computer Science (Grinnell College)*
Grinnell, USA
email: salesgar2@grinnell.edu

Benjamin Cerrato
*Computer Science (Grinnell College)*
Grinnell, USA
email: cerratob@grinnell.edu

Shayak Nandi
*Computer Science (Grinnell College)*
Grinnell, USA
email: nandisha@grinnell.edu

Fernanda M. Eliott
*Computer Science (Grinnell College)*
Grinnell, USA
email: eliottfe@grinnell.edu

***Abstract*—This research began from a reflection: are there any out-of-the-box, unexplored ways to investigate: 1. Group-influenced alliances and strategic decision-making, as well as 2. Reasons for people, institutions, or organizations to decide to commit to certain agreements but ultimately fail to follow through on them? And finally, 3. Can an observer get a hunch on others' intentions as the group dynamics and decision-making unfold and evolve? That reflection guided us to the Winkum Game, which, perhaps due to its complexity and lack of a winner in the traditional sense, has not attracted much interest. We are currently mapping the game's variables (not only its rules but also bridging it to group dynamics) for future computational modeling and experimentation. This could take the form of using either hand-designed agents to simulate specific human-grounded scenarios, or learning agents (applying Reinforcement Learning techniques) to investigate agent-agent dynamics and outcomes. The Winkum Game unfolds interesting power dynamics analogies: since winkers select and launch exchanges (picking sitters), they seem to be better positioned in terms of power, and we expect intriguing scenarios to emerge from flipping roles within the game, helping simulate real-world unpredictability and power dynamics.**

***Keywords-Computational Modeling; Cooperation; Coordination; Decision-making; Group Dynamics; Winkum Game.***

## I. Introduction

What decision-making insights can a game without a winner in the traditional sense reveal? Curious about the Winkum Game's resourcefulness in providing analogies for real-world applications, we decided to investigate ways of modeling it computationally. Given the game's intricate group dynamics, we anticipated that identifying which variables to model would be challenging. Although more difficult than anticipated, the design process has been insightful for decision-making reflection and enlightening for real-world analogies. This is a work in progress, but we are particularly excited to share it with a multi-disciplinary community. More importantly, we invite researchers across areas to partner with us as we computationally model the game, and we hope to attract more interest in the Winkum Game.

In the Winkum Game, there is a circle of chairs, and all but one chair is occupied by a group of players – the "sitters". Another group of players stands behind each chair – the "winkers. Therefore, there should be $n$ players in the game, where $n = (2 \times \text{number of chairs}) - 1$, and we would suggest having at least $n = 7$ for richer analysis. The winker behind the empty chair must wink at a sitter, who should try to escape from their current chair to sit on the empty chair. "Try to escape" because a sitter can only move to the empty chair if: they have been winked at, and are able to avoid, in the standing process, being tapped by the player behind them.

Figure 1 illustrates the game, Figure 2 summarizes the winker's and sitter's decisions within the game, and we provide additional details in Section III. Up until now, we introduced the game using "players" to facilitate understanding. However, from now on, we use "agents" instead. We define an agent as an autonomous figure that undertakes decision-making (which could be either a human or a learning agent).

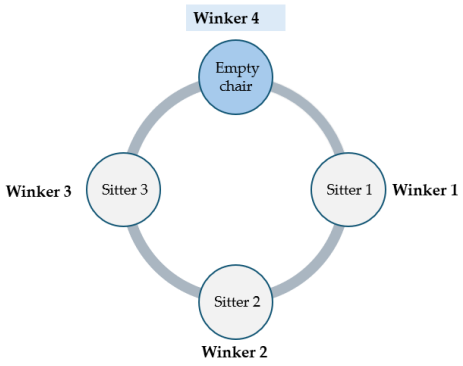


Figure 1. Winkum Game for $n = 7$: four chairs (circles), three sitters, and four winkers. Empty chair color-coded in blue, highlighting who will wink. Players are numbered to facilitate visualization, but sitters dynamically change positions.

Our contributions include spotlighting the game's multi-disciplinary applications and its potential for computational modeling. As this project evolves, we will use it as an analogy to 1. simulate real-world scenarios, e.g., agreements between countries or financial institutions, 2. analyze human behavior, and 3. embed the game and our modeling into a web application to enable people to interact with it. For example, if we refer to the original motivation of the game, which is for people to play and have fun, an instructor could



 ISBN: 978-1-68558-353-8 

use the application to set the game's variables, simulate their classroom, and brainstorm how well the game could work for their students (by visualizing *fun* levels and potential student *isolation* outcomes).

This work is organized as follows: in Section II, we briefly contextualize what makes the game interesting to model; in Section III, we detail the game; and in Section V, we discuss and conclude.

## II. Winkum Game and Networked Emotions

Given the game's strong group dynamics flavor likely impacting humans' internal variables, we saw interesting linkages with *Networked Emotions* – which refer to the view of "emotions as multi-layered processes in which intraindividual processes are tightly coupled and often cannot be separated from interindividual processes" [1]. Thus, in this work, by *layered feeding processes across variables*, we specifically refer to networked emotions as we seek to account for the social nature of emotions and the messy layers of emotion and emotion regulation in strategic decision-making.

Layered feeding processes across variables make the Winkum Game particularly insightful for analyzing group dynamics. For instance, there are three main variables' umbrellas: 1. overall game, 2. winkers, and 3. sitters. What we call *kindness* is among them, a variable that can favor winkers' cooperation (the *Morality-as-Cooperation* Theory [2] inspires our modeling mindset). Specifically, winkers and sitters privately rank each other (e.g., using three categories: not interested, neutral, or interested). If a sitter does not receive enough interest, isolation can emerge (when a sitter is never or rarely winked at, resulting in sitting in the same chair for a long period and being "exposed", since others can easily see the "exclusion or isolation"). Thus, high *kindness* levels make a winker more willing to wink at isolated sitters or even prevent that from happening – therefore enabling a winker to suppress self-interests and cooperate with others [3] [4] (in other words, wink at a "not interesting" sitter).

For now, we named this variable "kindness" instead of "empathy", seeking to ground the modeling on an act's external observation and refrain from assuming an agent's internal drivers; although we do acknowledge that "kindness" has a strong positive connotation, which points us to changing the terminology in the future (for example, a winker may be seen as "kind" but be driven by strategy only).

As we continue to find multidisciplinary pathways to model and interpret the game, our current inspiration sources include: *reputation* [5] as a possible guide for designing the agents' interests, and [6] for kindness, moral reputation, and related considerations. We use as inspiration the "morality-as-cooperation" theory [2], and finally [7] [8] for a hybrid take on Reinforcement Learning (RL) techniques (dry and wet) and the emergence of cooperation in utility-based computational approaches.

## III. Modeling the Winkum Game

The design process has been more difficult than anticipated, and it is currently in its early stages. In this section, we present our reflections on designing and subsequently computationally simulating the Winkum Game. Seeking to put the game "on the radar" of various communities, we focus on a high-level presentation and leave additional technical details for future work. Our next steps are to finish modeling, implementing, and testing the game, as well as to develop a software application (e.g., a Shiny app [9]) to enable people to set the game's variables and simulate different group dynamics scenarios.

We envision this software application providing two main options: enabling people to 1) run different predefined scenarios and 2) set the variables to run customized scenarios. Both options will generate data graphics for users to investigate outcomes and reflect on group-influenced decision-making. Possible visualizations are:

- Agent "winkability" and "tappability";
- How the "fun" variable changes across agents over time (see "fun levels" below);
- How dynamic is the game;
- Heat maps of who is winked at (including isolation);
- Escape rate *vs* tap rate.

Also, as we elaborate on our project in future work, we will explore further linkages between our approach and InvestESG, a multi-agent reinforcement learning benchmark for studying climate investment as a social dilemma [10], as well as possible uses within generative AI systems, such as the GOVernance of the Commons SIMulation (GOVSIM), a generative simulation platform for investigating interactions and cooperative decision-making in large language models [11]. Within that direction, [12] offers interesting reflections on the modeling of opponent shaping.

Below, we summarize the game's components identified so far, and we bridge them to realistic settings whenever possible. Finally, we will start modeling using a Gaussian distribution to set each agent's variables, gaining a sense of the design and game dynamics before implementing RL techniques.

**Game's goal.** Since there is no winner, what is the goal? Looking at it holistically, we can see the game as a way for agents to show and evaluate each other's interests as it evolves, a sort of "intelligence gathering". Through the game, winkers can reveal their interest in specific sitters (stakeholders) or pretend to be interested to deceive others; in other words, we can use it as an analogy for organizations or politicians who show interest in launching commercial, ecological, or political relationships with specific organizations or political agendas, but do not necessarily hold such interests.

Sitters can get a glance of their "popularity" across winkers through how many times they have been winked at in comparison to the others, and show their interest by vigorously trying to move to a winker's chair if winked at, or making less effort to leave a chair if winked at by a less interesting winker (or "commercial partner"). Overall, agents should have "fun" (which could be translated, for example, to "fostering commercial relations or advancing a political agenda"). Still, the game should be dynamic, enabling winkers to wink at different sitters and sitters to explore different chairs; therefore, there should be a balance between "showing interest" and

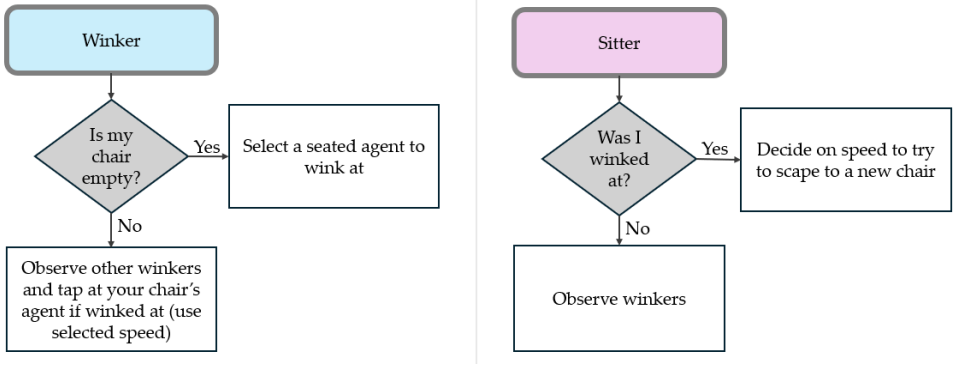


Figure 2. Summarizing decisions *per* agent's role: winkers at left and sitters at right. Note that other strategies are possible; for instance, a winker could, instead of observing other winkers, focus on observing their sitter and tap if noticing any movement, and a sitter could pretend to have been winked at, trying to move to a different chair.

"exploring different options" (offering interesting insights to RL's exploration-exploitation dilemma [13]).

**When does the game end?** Although the game lacks a clear ending mark, it could be set to end once the agents' average "fun" level drops below a certain threshold. At that point, sitters and winkers could switch places and continue playing: a winker could randomly be chosen to remain as winker, and sitters and winkers could be randomly assigned in the first round. Alternatively, new agents could be added, or simply end the game. We assume agents' interests align with making the game last. Therefore, winkers may wink at "isolated" sitters out of genuine kindness, to deceive others, or to prevent the game from ending (by keeping all agents at least minimally engaged); either way, from an external perspective, we simply call it "kindness".

**Kindness.** Before the first game begins, agents have their roles defined (winker or sitter) and privately rank each other (not interested, neutral, or interested). If a sitter is never or rarely winked at, isolation can emerge, and agents' ranks can be dynamically updated as they learn more about each other's interests as the game runs. For instance, sitters may have less fun if they identify peers as being isolated (and, of course, isolated sitters should have their fun levels drop, as well as winkers who are stuck with a sitter, or, on the contrary, have a hard time attracting or keeping sitters on their chair). Therefore, different ranking combinations will likely create distinct group dynamics, influencing fun levels, game duration, and agents' *participation*.

**Actions and participation.** Winkers' actions are to select who to wink at (wink selection) and the tapping speed to keep a sitter. If a winker assumes the corresponding sitter is winked at, they will decide how effective the tapping should be; for example, if a winker has been stuck with a sitter for a long time, they can pretend to try to keep a sitter but facilitate their escape. Game participation comprises observing agents and the winker in charge of winking at each round, and winking. The sitters' actions are to observe the current winker in charge of the winking and escaping speed: if the sitter believes they have been winked at, they will select the escaping speed based on their interest in moving to that specific winker's chair. They can also pretend an understanding to have been winked at, in the interest of moving to a different chair. Whereas game participation involves observing agents, receiving winks, and trying to escape.

**Fun levels.** To allow an agent's internal goals flexibility, it makes sense to set an overall "game's fun" variable in addition to individual agents' "fun" levels. The overall "fun" could either be an average of fun or designed upon something else, such as based on agents *participation*.

**Engagement.** Individual engagement levels increase if an agent participates in the game (and decrease otherwise). Once the engagement is depleted, the agent's fun decreases. The *engagement* variable defines how long an agent can stay engaged without participating before losing interest.

**Escapability and tapability** represent an agent's overall ability (or interest) to escape or successfully tap their opponent, encompassing factors like speed and reaction time.

**Consistency** indicates how reliably an agent performs at their best during a game.

**Environment.** For now, we consider as *environment* the number of chairs, game's rules and agents. Note that we will refine the environment in future work, as we agree with [14] that the environment is an important entity at the application level and that aspects conceptually distinct from the agents themselves should *not* be assigned or hosted within agents.

**Simulation.** A simulation can be defined as letting the game run until it reaches some halt condition (such as having agents switch roles up to a certain number of times and switches being triggered by a drop in average "fun").

**Reinforcement function.** We are still examining the game to define state and state transitions, agent architecture, and suitable RL algorithms. For the reinforcement function, we envision potential within variables, such as: participation, time spent with an agent according to their category, fun levels, winking success (for winkers), and escaping success (sitters).

Finally, additional variables should be considered if we seek to model the game as close as possible to real settings. For instance, one should also consider "visibility", which

is determined by the sitter's seating position relative to the winker, and a sitter seated directly across the winker should be easier to wink at – offering analogies to the ease of making partnerships with "easily reachable partners" even if they are not as "interesting" as others.

## IV. Discussion

The Winkum Game offers a flexible framework for studying strategic decision-making and testing group engagement and dynamics theories. We claim that the Winkum Game can be leveraged for broader research into group interactions, providing insights into how factors, such as active participation, kindness, tolerance for inactivity, and relationships among individual agents shape group performance and morale. For instance, things become particularly interesting once one analyzes questions, such as: How do we measure the authenticity of the fun factor in human settings? Is the fun displayed sincere or performed? Likewise, what about participants who may experience fun even without participating too much in the game? How do we account for the impact of "isolated" participants in the group dynamics?

Once setting up a simulation, to what degree should sitters and winkers be interested in interacting with each other? What is the relationship between overall "fun" and the interest level among participants? In what ways should the group dynamics influence individual ranks? (For example, by identifying a winker's group of interest, a sitter may lose interest in that winker). What does resilience mean in this game? How can an isolated sitter proactively respond to isolation? How to treat isolation and kindness?

Similar to [15], who seeks to model the evolution of cooperative behavior in realistic systems (by applying an evolutionary game in which lack of cooperation likely causes a cascading failure effect), we are particularly interested in the Winkum Game as we envision unexplored potential in investigating real-life group dynamics. As mentioned earlier, the computational modeling requires thorough reflection and multidisciplinary investigation as the agents' internal variables and the duration of the game feed from each other (intraindividual and interindividual processes [1]).

Finally, the Winkum Game offers a rich way to investigate group dynamics and decision-making. In the game, there is no traditional winner, and it focuses on agents having fun as they dynamically change pairs (which could be interpreted as dynamic agreements that unfold over time). At the same time, it is a game that enables agents to model other agents' mental models and strategies. While it may seem "naive" to read too much into this game, it enables us to simulate agents who get strategic insights into other agents' (or opponents') preferences. The game can be used as a metaphor to model different stakeholders' decision-making and simulate various scenarios, power dynamics, and preferences. For example, Commons research [16]–[18] can enlighten our modeling; whereas [19] offers insightful avenues for reflecting on how stakeholders govern themselves and make decisions.

## V. Conclusion and Future Work

Various studies can benefit from computational simulation (e.g., designing simulations of group behavior, enhancing data-collection methods, and providing insights for analysis and visualization to reveal hidden patterns and relationships). Our work aligns with this direction by introducing our initial scaffolding to model the Winkum Game and investigate group dynamics and decision-making. In our approach, we focus on agents' perceptions of fun, individual levels of engagement, and the popularity/likability of agents, all of which reflect the agents' commitment and willingness to keep playing the game. If we assume winkers are better positioned in terms of power, we can simulate real-world unpredictability by flipping power dynamics (e.g., by flipping roles within the game).

We suggest at least three applications of the Winkum Game: 1. Using it as an analogy to simulate real-world scenarios, e.g., agreements between countries, financial institutions, or politicians and political agendas; 2. Analyzing decision-making and group dynamics; 3. Enabling an instructor to set the game's variables, simulate their classroom, and brainstorm how well it would work for their students, which may enlighten considerations of how heterogeneous groups (e.g., age, gender, diversity) can improve group dynamics.

As the use of generative AI systems increases and hybrid spaces emerge, the Winkum Game may offer insights into these new forms of interaction. To conclude, we hope this work will spark multidisciplinary collaborations and attract greater interest in using the Winkum Game as an analogy for investigating group dynamics.

## References


[1] A. Kappas, "Social regulation of emotion: messy layers," *Frontiers in psychology*, vol. 4, p. 51, 2013.

[2] O. S. Curry, D. A. Mullins, and H. Whitehouse, "Is it good to cooperate? testing the theory of morality-as-cooperation in 60 societies," *Current anthropology*, vol. 60, no. 1, pp. 47–69, 2019.

[3] M. Tomasello, "Human culture in evolutionary perspective," *Advances in culture and psychology*, vol. 1, pp. 5–51, 2011.

[4] M. Tomasello and A. Vaish, "Origins of human cooperation and morality," *Annual review of psychology*, vol. 64, pp. 231–255, 2013.

[5] E. Brigatti, "Consequence of reputation in an open-ended naming game," *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, vol. 78, no. 4, p. 046108, 2008.

[6] D. Sperber and N. Baumard, "Moral reputation: An evolutionary and cognitive perspective," *Mind & Language*, vol. 27, no. 5, pp. 495–518, 2012.

[7] F. Eliott and C. H. C. Ribeiro, "Emergence of cooperation through simulation of moral behavior," in *Hybrid Artificial Intelligent Systems*. Springer International Publishing, 2015.

[8] ——, "Moral behavior and empathy modeling through the premise of reciprocity," *HUSO 2015*, p. 42, 2015.

[9] W. Chang et al., *shiny: Web Application Framework for R*, 2026, r package version 1.12.1.9000.

[10] X. Hou, J. Yuan, J. Z. Leibo, and N. Jaques, "Investesg: A multi-agent reinforcement learning benchmark for studying climate investment as a social dilemma," *arXiv preprint arXiv:2411.09856*, 2024.

[11] G. Piatti, Z. Jin, M. Kleiman-Weiner, B. Schölkopf, M. Sachan, and R. Mihalcea, "Cooperate or collapse: Emergence of sustainable cooperation in a society of llm agents," *Advances in Neural Information Processing Systems*, vol. 37, pp. 111 715–111 759, 2024.

[12] J. Duque et al., "Advantage alignment algorithms," *arXiv preprint arXiv:2406.14662*, 2024.

[13] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction (2nd ed.)*. MIT press, 2018 [1998].

[14] D. Weyns, G. Vizzari, and T. Holvoet, "Environments for situated multi-agent systems: Beyond infrastructure," in *Environments for Multi-Agent Systems II: Second International Workshop, E4MAS 2005, Utrecht, The Netherlands, July 25, 2005, Selected Revised and Invited Papers 2*. Springer, 2006, pp. 1–17.

[15] W.-X. Wang, Y.-C. Lai, and D. Armbruster, "Cascading failures and the emergence of cooperation in evolutionary-game based models of social and economical networks," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 21, no. 3, p. 033112, 08 2011.

[16] E. Ostrom, *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.

[17] E. Ostrom, J. Burger, C. B. Field, R. B. Norgaard, and D. Policansky, "Revisiting the commons: local lessons, global challenges," *science*, vol. 284, no. 5412, pp. 278–282, 1999.

[18] A. Agrawal, J. Erbaugh, and N. Pradhan, "The commons," *Annual Review of Environment and Resources*, vol. 48, no. 1, pp. 531–558, 2023.

[19] R. C. Ellickson, *Order without law: How neighbors settle disputes*. Harvard University Press, 1991.