Developing The Strategists, an Online, AI-Aware Financial Investment Game for Technology-Enhanced Learning and Research

Shubham Chawla
Arizona State University
schawl17@asu.edu

Bretho Danzy III
Arizona State University
bdanzy@asu.edu

Jai Narula
Arizona State University
jnarula1@asu.edu

Ashish Amresh
Northern Arizona University
ashish.amresh@nau.edu

Chris Bryan
Arizona State University
cbryan16@asu.edu

Abstract

Serious games have been shown to have many positive outcomes across educational and research settings; however, the design of such games is non-trivial. In this paper, we investigate how to design and engineer a serious gaming platform for a financial investment game called The Strategists. Building on the familiar mechanics of Monopoly, The Strategists is a web-based multiplayer game that introduces additional features such as property market dynamics and the flexible integration of AI and data-driven interventions. The platform supports robust data logging, easy deployment, and simplified data management, making it an accessible tool for both educators and researchers. We describe the process of designing, developing, and evaluating The Strategists, and discuss how lessons learned during this process can provide value to diverse communities across both education and research, including business and finance, computer science, and social and behavioral science.

Keywords: serious games, technology-enhanced learning, finance-themed games, human-AI interaction

1. Introduction

Financial and economic investment board games like Monopoly are commonly used as tools for teaching and engaging with financial and entrepreneurial concepts, including risk assessment and strategic decision making (Mangelep et al., 2023; Rosli et al., 2019; Shanklin & Ehlen, 2007), and more broadly can serve as platforms for studying human behavior, game theory, and group dynamics. In parallel to this, the rapid advancement of artificial intelligence (AI) has led to digital games (including digitalized board games) that integrate AI as a mechanism to enhance, augment, or study human gameplay, via real-time interventions such as personalized feedback, adaptive difficulty setting, and strategic analysis and recommendation.

Such "AI-aware gaming experiences" offer new avenues for studying cross-disciplinary topics, including human-computer and human-AI interaction, but they also motivate a need to develop robust and customizable software platforms that will enable such study. Likewise, while it might be appealing to immediately jump to generative AI tools for modeling humans and recommending player actions in serious games, such an approach makes it difficult to create replicable experimental designs (Pérez et al., 2023). In contrast, rules-based AI (which we initially investigate and employ in this paper) provides several benefits like allowing the controlled design and evaluation of game mechanics, supporting replicable learning about how humans interact with AI models, and allowing researchers to pinpoint potential design problems.

A similar problem likewise exists in pedagogy: educators wanting to teach about financial and strategic concepts in a technology-enhanced manner need accessible software platforms for hands-on learning, which have the potential to promote outcomes including inquiry-based learning, active engagement, and deeper and more sustained learning (Pahl & Kenny, 2008).

Specifically for the current paper, we investigate how to design and develop such a platform that can serve for both educational and research purposes, with a novel financial investment game called *The Strategists*, which builds upon the familiar mechanics of Monopoly, whereby players roll a (virtual) dice and take turns navigating a game board circuit. However, *The Strategists* introduces a more sophisticated property investment model that allows players to strategically

invest in percentages of properties, which dynamically influences their market values and rent collections. This can provide a richer simulation of real-world investment scenarios (Brueggeman & Fisher, 2011), making it an ideal platform for studying and learning about decision-making processes in a controlled and reproducible environment.

To help facilitate education and research, *The Strategists* includes robust data logging and flexible storage mechanisms. The platform also adopts a modular architecture design that allows for integrating rule-based AI and other data-driven recommendations (or interventions) to augment gameplay in a personalizable and controllable manner. These features are complemented by an intuitive user interface that presents the game board alongside interactive data visualizations of gameplay information.

In this paper, we discuss the design, development, and evaluation of *The Strategists*, including how we distilled a set of design requirements for such a platform, and how the game's UI, UX, and backend features were refined and validated via empirical evaluations. At a high level, *The Strategists* addresses emerging needs for interdisciplinary solutions in game research by combining elements from game design, AI, visualization, and financial education. This work also contributes to human-computer interaction and serious gaming, as we employ a design study methodology (Sedlmair et al., 2012) to understand and reflect on how to successfully design a flexible multiplayer game platform that integrates AI and data-driven recommendation functionalities in an accessible and customizable manner for the research and education communities. The codebase and documentation for *The Strategists*, including installation and gameplay instructions, is published open source at: https://github.com/shubham1chawla/the-strategists-remastered/

2. Related Work

Digital Games as Platforms for Education and Research. The use of digital games for education and research is a form of "serious games," whereby games are leveraged for purposes beyond simply entertainment (Laamarti et al., 2014). There is rich literature discussing the diverse areas where serious games are deployed for education, which includes language learning, STEM, social sciences, healthcare, vocational training, and more (Lamb, 2024). In particular, serious games can act as a form of experiential learning (i.e., learning through meaningful participation in, and reflection on, real/simulated experiences) and be scaffolded to promote inquiry-based learning (Lameras et al., 2021; Perrotta et al., 2013).

More broadly, digital and serious games can fall under the umbrella of technology-enhanced learning, which has been shown to improve student outcomes such as by promoting deeper and more thoughtful knowledge retention (Booker & Mitchell, 2021) and enhancing creative thinking and engagement (Serrano et al., 2019). Positive outcomes are not restricted only to students; studies have also found benefits for instructors and educators such as increasing efficiency and productivity (Haleem et al., 2022), supporting professional development and training (Scott et al., 2017), and scaffolding teaching in large and/or online classrooms (Serrano et al., 2019). Games are also used for research across a number of disciplines, including computer science, engineering, business, healthcare, and emergency response (Coovert et al., 2017).

Specifically for finance games, Monopoly has remained one of the most popular board games in the world since its introduction almost 100 years ago (Hackett & Coghlan, 2023). Today, Monopoly is commonly used as a tool for teaching financial and entrepreneurial concepts (Mangelep et al., 2023; Rosli et al., 2019; Shanklin & Ehlen, 2007), and several versions of the game are freely available online.

However, these current platforms have critical limitations in the context of serving as research and educational tools, such as not being open source and not supporting flexible and accessible deployment and data collection. Several specific challenges are distilled in our design requirements analysis (Section 3). These are subsequently addressed in implementing (Sections 4–5) and evaluating (Section 6) *The Strategists*.

Integrating AI and Data-Driven Interventions in Digital Games. AI and data-driven recommendations are increasingly being integrated into digital games as a way to personalize the playing experience, such as generating custom interactions with non-playable characters and environments and providing dynamic hints and suggestions to influence player decisions and strategies (Gomez-Alvarez et al., 2024). Specifically for Monopoly, several papers have developed learning-based approaches to recommend optimal decisions during gameplay; while early work formalized this as a Markov decision process (Ash & Bishop, 1972), recent approaches have primarily leveraged reinforcement learning techniques (e.g., Arun et al., 2019; Bonjour et al., 2022).

In parallel, a related area of research is exploring how AI and data-driven interventions impact player decision making, cognition, and social dynamics when playing games. For example, a recent

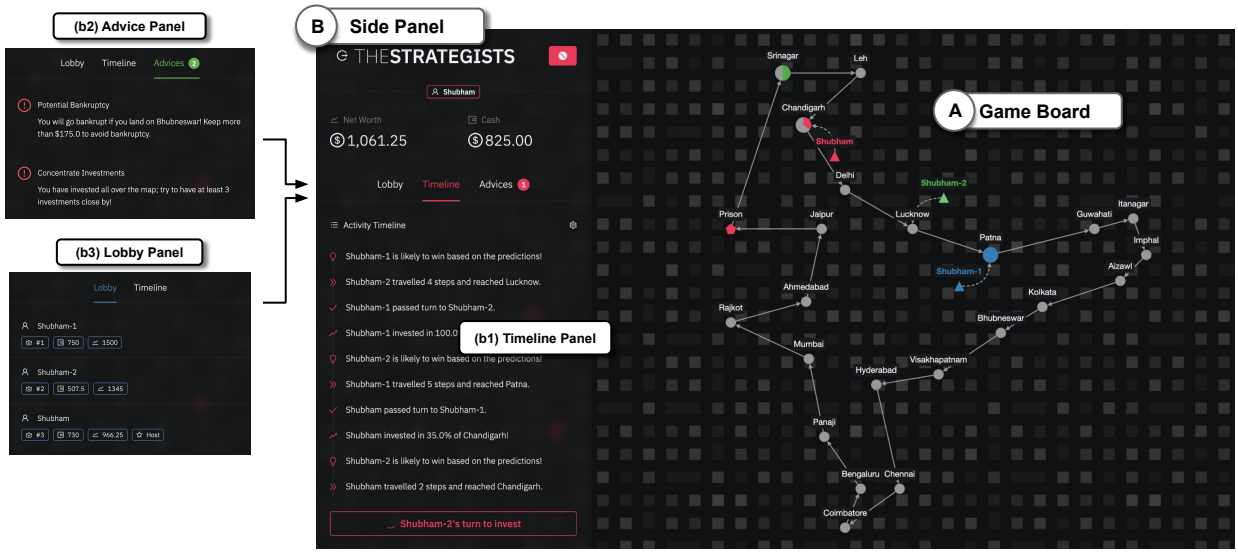

Figure 1. The Strategists' main interface includes two primary panels. (A) The game board shows the game's properties in a circuit, with player icons indicated by colored triangles and property ownership via colored pie slices on properties. Interacting with elements on the game board can trigger various popups (see Figure 3). (B) At left, the side panel shows information for the player and contains a trio of toggleable panels: (b1) the timeline panel shows the history of game actions, (b2) the advice panel shows data-driven player recommendations that have been implemented for the game, and (b3) the lobby panel shows summary information about players.

experiment (Solans et al., 2022) trained an AI model to suggest actions in a single-player exploratory drilling game. The authors found that participants who followed the model's advice scored higher, and these participants placed "moderate" trust in the model without over- or under-relying on it. Interestingly, this result contrasted with participants' stated preferences for accepting the model in a post-study survey, suggesting there are potentially conflicting spaces about how humans trust and operationalize recommendations during gaming. Similarly, another recent study (Sidji et al., 2024) investigated how the use of an LLM assistant impacted factors such as player cognition and social dynamics in Codenames, a team-based cooperative game. Many players felt that the AI assistant challenged "the spirit of the game" and weakened the social connections of their human teammates.

Pertaining to the current paper, we are not aware of any platforms that facilitate studying AI and data-driven interventions within a multiplayer finance/investment game environment. As a software artifact, *The Strategists* can support experiments in AI and human-AI interaction (e.g., see our second evaluation as a demonstration of this, in Section 6), though the current paper, in accordance with the design study methodology (Sedlmair et al., 2012), primarily contributes insights about *how* to effectively design and engineer such a platform.

3. Design Requirements Analysis

The development of *The Strategists* was motivated by several factors, including a growing need for serious gaming platforms to support technology-enhanced education and reproducible research. To help guide the design of our platform, we first distilled a set of high-level design requirements, based on a meta-analysis (e.g., using the research papers discussed in Section 2) of open challenges and gaps regarding research and educational software platforms that support accessible management and the deployment of AI and data-driven interventions. Based on this, we distilled a set of four design requirements DR1–DR4 for developing an online, AI-aware financial investment game for technology-enhanced learning and research.

DR1: Familiar Mechanics with Enhanced Features for Financial-Based Gaming. Given that overly steep learning curves (e.g., for highly complex or novel gaming experiences) can impede educational activities or act as a confound for experiments, we intended to design a recognizable experience with UI/UX mechanics similar to existing finance games like Monopoly. However, to promote a more streamlined and strategic-thinking experience, we modified the

base Monopoly experience by simplifying some aspects of the game (e.g., removing some game mechanics that introduce too much uncertainty or stochasticity into gaming outcomes, such as Chance squares) and developed a set of modified gameplay elements, such as the ability to partially invest in a property.

DR2: Flexible Support for Computational Interventions. As previously discussed, there is an increasing body of research at the intersection of AI and data-driven recommendations and digital games. We wanted educators and researchers to have the ability to design and integrate such computational interventions directly into the gameplay experience, as a way to study how players respond to or interact with these interventions. Our implemented platform is extensible and can easily integrate customizable interventions, including both rules-based methods and trained machine learning models via the use of a modular design that supports extensible Python modules.

DR3: Oversight and Control for Study Proctors and Educators. When conducting experiments, researchers need exact control over study variables to support rigorous and reproducible findings. For a gaming scenario, this would include factors like which players are able to participate, whether computational interventions (and other game features) are enabled/disabled, the order of play, and more. As mentioned in DR2, The Strategist also supports personalizable advice and data-driven recommendations that can be tailored to each player, allowing for highly controllable experimental designs. Likewise, educators might want to customize the gaming experience for their classrooms (such as having students play on a custom or themed map). We implement a number of administrative controls to support these types of constraints.

DR4: Easy Installation and Accessible Management. Finally, the target users who would install and administer this platform are academic educators and researchers, who in many cases might not have software engineering or programming experience. A software tool that is difficult to install and maintain — e.g., requiring the setup of a custom database to store and access collected gaming data, or requiring the installation of niche or bespoke system libraries or settings — would be difficult to use, thus hindering adoption. As such, we adopt a "lightweight" strategy for application deployment, administration, and data management. The platform can be deployed as a Dockerized container, either on a web hosting service (including accessible public platforms such as Netlify or Cloudflare, or to an institution's internal servers via IT support), and games can be set up and initialized and subsequently accessed by players using web browsers. *The Strategists* also leverages Google Drive for ease of storage (saving collected data as CSV files) and does not require setting up a formal database (e.g., SQL), which makes it easier for educators and researchers without a background in computer science or IT to manage and review collected gaming data.

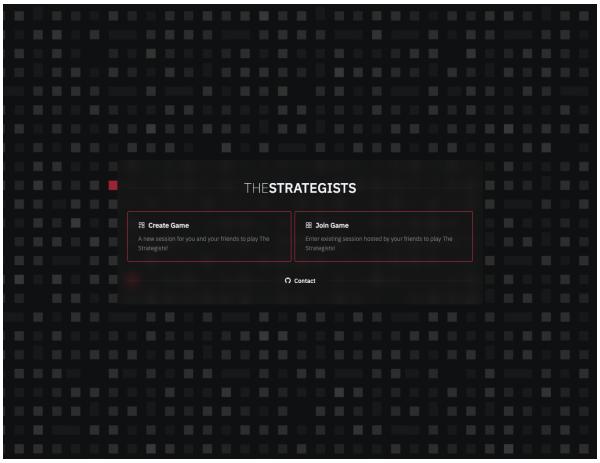

Figure 2. The initialization page allows users to create new games or join an existing one.

4. Interface Description and Gameplay

The Strategists is an open-source, web-based, multiplayer game that builds upon the property investment model of Monopoly. Here, we briefly overview the gameplay mechanics and UI/UX aspects of *The Strategists*. Full documentation, including installation and gameplay instructions, is available with the application's codebase.

Beginning a New Game. To begin a game, a host first spawns a new game from the initialization page (see Figure 2), which generates a four-digit code that can be provided to the game's players. From the initialization page, players can click to join the game (they are prompted to enter the code), which places players in the waiting lobby (see Figure 1(b3)) alongside an empty game board, until the host starts the game.

Gameplay. During gameplay, players take turns using a simulated dice roll to move around properties on a circuit. *The Strategists* supports using custom maps (which are configured and stored as JSON files); for the current paper, we showcase a curated map of 22 Indian cities. Similar to Monopoly, maps in *The Strategists* can contain properties, Jails, and "Pass Go" locations, however the current implementation omits other types of game squares (e.g., Chance, Railroards, Tax Squares, etc.), with the intent of focusing gameplay

on the economic decision-making processes of players.

During their turn, a player has the option to invest in the property on which they are currently located. Like Monopoly, properties are worth different amounts of money, but *The Strategists* extends this concept by allowing players to partially invest in a property (this means that multiple players can own a portion of a property). This variable investment model also inflates the market value of properties, as the value of a property increases based on their investment ownership (the default increase value is set to be linear, though this can be adjusted). The intent is to add both dynamism and better simulate real business and economic realities (Brueggeman & Fisher, 2011). When a player lands on a property with investments by other players, they pay rent proportional of the ownership to each of these investor (this is true even if they also own a percentage of this property). Players also can purchase (or increase) their own investment percentage on the property, up to 100% of the property value. If a player's funds are fully exhausted, (either by spending all of the funds or by paying rent to other players), they declare bankruptcy and their investments are cleared. Similar to Monopoly, players take turns until only one player remains, who is declared the game winner.

The Game Board. Figures 1 shows the game's main interface, including (A) the Game Board and (B) Side Panel. In the game board, properties are shown as labeled gray circles on a circuit, with directed edges connecting to the next location. The x/y coordinates of each property can be specified in the map's configuration file (for example, in Figure 1(A), the properties are arranged to form a rough outline of India). Players are indicated by triangular icons (colored by different hues) with an arrow pointing to their current location; the icon blinks for the currently active player. When properties have investments, their circular node begins to "fill up" with a pie chart glyph; colors and pie slice sizes correspond to the player and ownership amounts. A fully filled circle indicates a property is 100% invested.

Interactions with player or property icons trigger various popup panels (Figure 3); when triggered, these overlay on top of the game board and support corresponding user actions or strategizing. For example, clicking on a property loads the Investments Panel (Figure 3(a)), which provides information about the property (e.g., its market value and investments). When the player is currently located at that property, they are also offered a slider and green 'Invest' action button to invest a desired percentage in the property. Since each property's market value differs, players must pay attention to their proposed investment amount and how much they are willing to commit, especially since new investments increase the market value of that property. When other players land on the property, they pay rent proportional to the invested ownership.

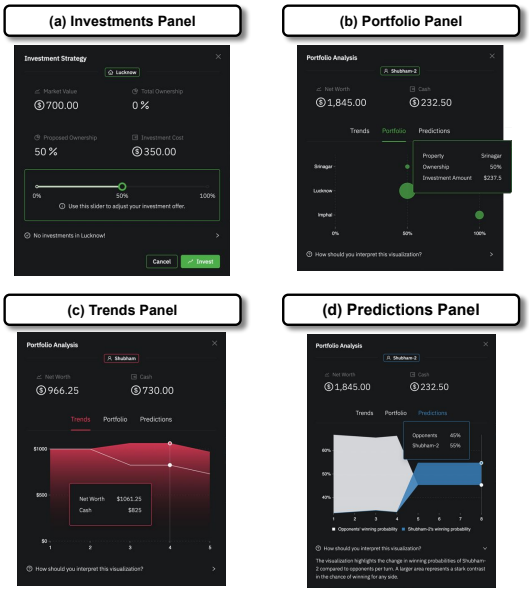

Figure 3. Example popup panels that support different player actions.

Clicking on a player icon loads the Portfolio Panel (Figure 3(b)), which summarizes their net worth and visualizes their property investments with a dot chart that uses circle size to indicate investment amount. This popup supports tabbing over to a Trends Panel (Figure 3(c)), which displays a combination area-plus-line chart showing the player's cash on hand (white line) and the net worth (red area) over time.

If enabled, players can also tab to a Prediction Panel (Figure 3(d)), where predictions about the player's chances of winning are shown. This panel supports integrating rules and/or machine-learning-based models for predictions (in the figure, a Gradient Boosting model was trained based on a curated corpus of player investment patterns and wins, though the panel is model-agnostic and supports updating to newer/different models). The area difference visualization shows the currently selected player's win probabilities against the field (i.e., all other players currently in the game).

The Side Panel. To the left of the Game Board, the Side Panel (Figure 1(B)) shows the player's net worth and cash on hand, and contains three tabs that players can toggle between. The Timeline Panel (Figure 1(b1)) shows the history of game actions. The Advice Panel (Figure 1(b2)), when enabled, displays

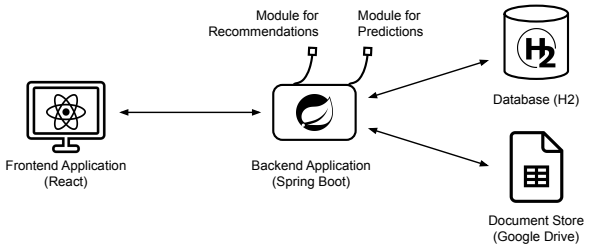

Figure 4. Overview of the platform architecture.

recommendations to the player. Advice items are updated during gameplay and can be triggered based on the current state of the game and player actions, and an icon lists the number of unread advice items to the player based on when the tab was last viewed. The recommendations functionality is highly extensible, and developers can easily add new or different data-driven recommendations that are triggered based on game state and/or player actions (see Section 5 for more details on this process). Finally, the Lobby Panel (Figure 1(b3)) ranks players based on their current net worth (a summation of their current cash and the market value of their investment portfolio) and shows which players have declared bankruptcy.

5. System Architecture and Management

The Strategists is implemented as a web application suitable for deployment on either local or public hosts (for our testing, we used Netlify to host the frontend and Google Cloud Platform to host the backend using a Docker container). Figure 4 shows a high-level overview of the platform's components and how they are connected. The frontend is implemented in React using a combination of Cytoscape, D3.js, and Ant Design, and the backend is built in Java using the Spring Boot Framework. Communication between the frontend and backend is handled via REST API endpoints, with live events sent using Server-Sent Events (SSE). Within the backend, H2 (an in-memory relational database) alleviates the need for hosting and managing a dedicated database (e.g., SQL). Gameplay data is also exported to Google Drive for easy management and review. In-game predictions about each player's winning chances are handled by extensible Python modules that can easily be replaced or updated to support future research and educational activities.

Gameplay Data Collection and Storage. Data collected during gameplay, such as logging information of player actions and any advice information given to players, is stored via syncing to a Google Drive folder. Specifically, a deployed instance of *The Strategists* can be linked to a Google account and Drive folder location, and CSV files will be created and edited in real time during games. Several "default" logging events are collected during gameplay, including information about ownership, investments, debit amounts, debit count, investment numbers, and rent payments, though the system is extensible to logging additional activities. Notably, game logging is designed to preserve player privacy: personally identifiable information (emails, the names players assign themselves, et.c) is not stored. Instead, the system generates numerical IDs for each player and game session.

In addition, if recommendations have been enabled for a game, information about advice items are collected and stored (in a separate CSV) to the Google Drive folder. Table 1 shows an example of several default data points that are collected for advice information, though again the specific logging is extensible to accommodate preferred data points and reporting semantics.

Table 1. Examples of several default advice-related data points that can be logged during gameplay.

Data Field	Description	Value Type
Game Code	A 4-digit alphanumeric code to uniquely identify the game	String
Player ID	A system-generated numeric code to uniquely identify the player assigned to the advice	Numeric
Advice State	Whether the advice is new or followed	String
Advice Type	The type of advice shown to the player	String
Advice Priority	The assigned priority to the advice	Numeric
Additional Parameters	Extra information, if desired, that is assigned to the advice, such as the game turn when the advice is triggered.	String
Advice New Count	A count of how many times the specific type of advice has been created for the given player	Numeric
Advice Followed Count	A count of how many times the specific type of advice has been followed by the given player	Numeric

Implementing Data-Driven and Personalized Recommendations into the Gaming Experience. When enabled, data-driven recommendations are shown in the frontend's Advice Panel (Figure 1(b2)). Advice items are customizable for each player, supporting personalized recommendations and interventions. To support this, *The Strategists* leverages a modular design where recommendation engines and/or trained models can be linked to monitor the game state and events, and generate recommendations during gameplay. At the beginning of a player's turn, the module can be queried (including passing the current gameplay state and information about the players), and the module returns zero, one, or more recommendations that are displayed as advice items.

An example of a simple set of recommendations (using a rules-based logic) is shown in Table 2. When a recommendation is triggered, an advice text string can also be generated that is displayed to the player on the frontend (two examples are shown in Figure 1(b2)). When the module is queried, if no advice items are returned, it is assumed that there are no new recommendations for the player during that turn. When a new piece of advice is received, it is added to the top of the Advice Panel. If this panel is not currently in focus, a numerical counter is incremented on the tab label to indicate the count of unread advice items.

Table 2. A simple example set of rules-based advice items that can be triggered for a player.

Advice	Description	Trigger
Avoid Timeout	The game marks players bankrupt if they timeout more than three times. This advice suggests that players avoid timeouts.	The player's turn timed out in the previous cycle.
Invest Frequently	This advice suggests players invest often to gain a competitive edge.	The player has skipped their turn for the last three turns.
Invest Significantly	The advice suggests that players invest more than 50% to increase the expected rent from the invested property.	The average investment ownership of the player dips below 50%.
Concentrate Investments	The advice suggests that players concentrate their investments on one side of the map to maximize their chance of receiving rent.	The player does not have investments in three properties within a neighborhood of six consecutive spaces.
Potential Bankruptcy	This advice suggests players hold cash as they are in danger of bankruptcy.	A player's cash on hand is lower than the maximum rent they might have to pay in the next board circuit.

Implementing Prediction Capabilities for the Prediction Panel. Similar to how advice logic connects to extensible and updatable modules to generate dynamic recommendations at the beginning of each player's turn, prediction capabilities are similarly controlled by connecting to service modules that return percentage-based predictions of who is likely to win the game. The platform can submit the current gameplay context (e.g., the state of the game board and each player's investments and net worth), and receives back a formatted set of predictions about each player's current chances at winning. These populate the Predictions Panel (Figure 3(d)). During development, we tested a number of prediction models (including Random Forest, Decision Tree, AdaBoost, and Gradient Boosting) trained on a set of 103 games played. During the evaluations, the Gradient Boosting model (which gave the highest accuracy during the validation tests) was used.

6. Evaluation

To evaluate *The Strategists*, we performed two rounds of evaluation, with a goal of demonstrating how the platform can support DR1–DR4 outlined in Section 3. The first evaluation focused on assessing the usability of the platform, including validating the overall design, functionalities, affordances, encodings, and interactions of the game. Based on this, we collected a set of improvements that needed to be made to the platform and implemented a round of design improvements with follow-up target validations to ensure sufficiency. After reaching a stable "version 1.0" system, we conducted an experimental demonstration study to validate that the platform can be deployed and used for research and education activities.

Initial Usability Study. After completing an initial build of *The Strategists*, we first conducted a usability study with six players who played several rounds of the game. The goal with this study was to evaluate the overall playability of the platform regarding the game's design, gameplay, UI/UX, etc. Recruited participants were graduate students and one faculty member at Arizona State University; all were familiar with Monopoly, though they had not previously played *The Strategists*. After a tutorial on the game's rules and mechanics, the study's gaming sessions took approximately one hour to complete. Afterwards, participants provided feedback about their experiences playing the game, with an emphasis on highlighting the positive aspects of the game and seeking to identify any "pain points" or frustrations with the current gameplay experience. Participants were co-located in a quiet office-like room during the study.

Usability Study Results and Refinements. Overall, participants uniformly felt the game represented a familiar experience (e.g., similar to Monopoly) while providing its own unique and engaging twists in terms of gameplay and strategic elements, such as the ability to partially invest in locations. Participant feedback was also highly positive regarding overall system usability of the system, the various implemented features and popups, and the ability to easily access and begin games.

Based on the collected feedback, we also identified several tweaks to improve the gaming experience, such as adding additional encodings to more clearly indicate which player had the current turn. Based on the feedback, we implemented a set of system enhancements (with follow-up targeted evaluations to validate their efficacy), resulting in the "completed" system described in Sections 4 and 5.

Follow-Up Demonstration Study. We next conducted a demonstration study to validate The

Strategist's feasibility as a platform for conducting experimental studies and educational activities. For this, we tested two conditions: games played with predictions and recommendations enabled, and games played with predictions and recommendations disabled. For the "enabled" condition, we used the recommendations list outlined in Table 2 (these were also used in the previous usability testing). We recruited a total of eight player groups for the demonstration study, with four groups assigned to each condition (recommendations enabled/disabled). Each player group contained four players, meaning a total of 32 participants were recruited for this study. Participants were students recruited from Arizona State University and were randomly assigned to a player group. Game rounds were played in a quiet, office-like room, with each player using their own laptop. For each group, a game tutorial was first given, and then the group played three game rounds before completing a post-study survey.

Demonstration Study Results. As the primary goal of this second study was to demonstrate that the platform could collect and store session data, we primarily focus on reporting how the tool supported collecting data in a coherent manner that allowed for investigation of potential research questions or hypotheses (e.g., that could be asked as part of an experiment), as opposed to focusing on analyzing the collected data itself.

During gameplay, game and advice data was automatically collected and stored in CSVs to a Google Drive folder for post-hoc review. This allowed us to test out, for example, diverse questions about how player performance might vary across the enabled/disabled conditions. For example, we could easily investigate if players who receive recommendations exhibit different investment strategies compared to players in the recommendations disabled group, based on analyzing property investment amounts and frequencies during games. We could also investigate if players who followed advice items would be more likely to win compared to players who did not follow advice, based on analyzing game log data.

Alternatively, we could combine both game logging data with measures from the post-study survey, such as examining if a player's perceived trust in AI (a post-study survey question) correlated to greater adherence to recommendations based on logged advice data. Likewise, we could assess how a player's self-perceived financial literacy (another survey question) correlated to looking at and following recommendations. In our specific case, we tested many of these above-mentioned questions via a combination of t-tests and/or Fisher's tests, though the results were generally either insignificant (with $p > 0.05$ values)

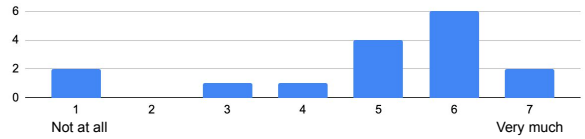

(a) Participants agreement with "I found AI predictions and advice helped me strategize and improve my winning probabilities."

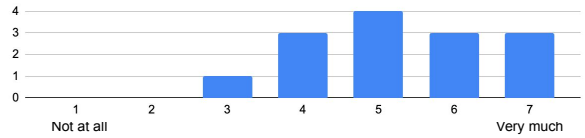

(b) Participant agreement with "After losing a game, I relied more on my instincts than predictions and advice."

Figure 5. Participant response counts from the demonstration study's post-study survey (from the AI "enabled" group) about their perceptions of in-game recommendations based on a 7-point Likert scale.

or showed very small effect sizes, likely due to the small sample sizes tested. Interestingly however, these findings conflicted with self-reported player perceptions of how they used AI. For example, Figure 5(a) shows that participants tended to agree that the in-game AI-driven advice was helpful, but Figure 5(b) indicates such advice began to be ignored when they did not win (keeping in mind, there were four players and only one winner in each game). While such findings are largely formative, and open interesting future questions which we discuss more in Section 7, we note again that statistical analyses and hypothesis testing were not the primary goal of this demonstration study (particularly given that the advice items listed in Table 2 are relatively "simple"). Thus, we still consider this demonstration study a successful validation of *The Strategists* as an experimental platform, as its design and data management approach made it easy to review data and investigate questions in a coherent and rigorous manner.

Interpreting the Studies. The two evaluation studies provide a complementary pair of validations for *The Strategists*, both from a usability perspective and as a feasibility demonstration of the tool for controlled game sessions (e.g., for a research or classroom setting). The takeaways from the usability study help demonstrate that the platform provides an engaging, intuitive, and familiar (yet different) investment-based game experience. The demonstration study shows the tool can collect and organize data across different player groups and conditions, supporting rigorous and reproducible science.

7. Discussion

The Strategists represents a novel, AI-aware, open-source gaming platform for technology-enhanced learning and research. The platform is built around a set of distilled design requirements (DR1–DR4), and surfaces an intuitive and engaging gameplay experience, while providing accessible administration and management control of gaming sessions and data collection. Below, we discuss some of the lessons learned designing, implementing, and validating *The Strategists*, including how the platform can facilitate cross-disciplinary research and pedagogy, and discuss identified opportunities where future work can build on the results developed in this paper.

Supporting Design Requirements DR1–DR4. Based on our process of designing and evaluating *The Strategists*, we demonstrate how the platform successfully fulfills the design requirements DR1–DR4 outlined in Section 3. DR1 is demonstrated by the overall UI/UX of the game, which is familiar to players who have played games like Monopoly (such connections were mentioned by participants during both evaluation studies). DR2 is addressed by the platform's ability to handle AI and data-driven recommendations, via its modular design. DR3 is exemplified in the overall design of the system as a tool for controlled gaming sessions, exemplified by Section 6's demonstration study. Finally, DR4 is likewise shown in the way the system's architecture is designed for easy deployment, access, and data management.

The Strategists as a Boundary Object for Research and Education. "Boundary objects" refer to an entity (such as a concept, practice, or artifact) that is flexible enough to be interpreted differently across disciplines but structured enough to maintain a coherent identity across them (Star & Griesemer, 1989). In the case of *The Strategists*, the design and presence of the AI-aware financial game can serve to bridge across various fields, including business and finance, computer science and AI, social and behavioral sciences, game theory and design, and more, each of which has its distinct perspectives, methods, and knowledge systems.

Through this lens, *The Strategists* can be used to explore diverse research questions and support educational activities. For instance, from a finance and business educational perspective, the tool could be used to teach players about adapting to market dynamics, and researchers could study how such dynamics impact decision making and risk. From a computer science perspective, a researcher could develop and test novel machine or learning models to optimize a player's strategic decision-making.

Towards more complex market dynamics. *The Strategists* currently implements a simplistic version of market dynamics (sometimes called investment-driven price appreciation) (Brueggeman & Fisher, 2011); we are interested in expanding this to more complex setups and scenarios, though there are questions about how best to do this in a way that maintains a streamlined gaming experience and a reasonable learning curve (particularly for education and research contexts, where overly complex game mechanics could confound results).

Additional Future Work and Opportunities. We are actively maintaining and updating the codebase for *The Strategists* (including adding new features and functionalities), and plan to conduct follow-up studies that focus on the decision-making aspects of players (particularly in the context of AI-driven recommendations). Along these lines, we are interested in making the recommendation logic and advice features more customizable and accessible, and testing more advanced models with larger and more robust training datasets (e.g., via reinforcement learning). Likewise, there has been recent burgeoning research into how incorporating Generative AI and chatbots into games impacts player decisions, cognition, and satisfaction (e.g., Sidji et al., 2024). We intend to similarly develop such features to allow for these types of studies.

8. Conclusion

There is a need for serious gaming platforms that support education and research, however the design of such artifacts is non-trivial. In particular, accessible solutions that enable rigorous, controllable gaming sessions while minimizing technical overhead provide significant potential in addressing key challenges faced by educators and researchers. This paper investigates how to design and engineer such a tool, in the context of financial and investment-themed games, establishing both a software platform that can be used by these communities and also serving as a foundation for future interdisciplinary work across serious gaming, computer science, and more.

Acknowledgement

This research was supported in part by the U.S. National Science Foundation through grant DUE-2216452.

References

Arun, E., Rajesh, H., Chakrabarti, D., Cherala, H., & George, K. (2019). Monopoly using

reinforcement learning. *2019 IEEE Region 10 Conference (TENCON)*, 858–862.

Ash, R., & Bishop, R. (1972). Monopoly as a Markov process. *Mathematics Magazine*, *45*(1), 26–29.

Bonjour, T., Haliem, M., Alsalem, A., Thomas, S., Li, H., Aggarwal, V., Kejriwal, M., & Bhargava, B. (2022). Decision making in Monopoly using a hybrid deep reinforcement learning approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, *6*(6), 1335–1344.

Booker, K., & Mitchell, A. (2021). From boring to board game: The effect of a serious game on key learning outcomes. *Journal of Occupational Therapy Education*, *5*(4), 7.

Brueggeman, W., & Fisher, J. (2011). *Real estate finance and investments*. McGraw-Hill Irwin.

Coovert, M., Winner, J., Bennett Jr, W., & Howard, D. (2017). Serious games are a serious tool for team research. *International Journal of Serious Games*, *4*(1).

Gomez-Alvarez, D., Lopez-Franco, M., Carranza, D., Lopez-Franco, C., & Lopez-Franco, L. (2024). The future of gaming: How artificial intelligence is revolutionizing the industry. *ReCIBE, Revista electrónica de Computación, Informática, Biomédica y Electrónica*, *13*(3), E6–10.

Hackett, L., & Coghlan, J. (2023). Why Monopoly monopolises popular culture board games. *M/C Journal*, *26*(2).

Haleem, A., Javaid, M., Qadri, M., & Suman, R. (2022). Understanding the role of digital technologies in education: A review. *Sustainable Operations and Computers*, *3*, 275–285.

Laamarti, F., Eid, M., & El Saddik, A. (2014). An overview of serious games. *International Journal of Computer Games Technology*, *2014*(1), 358152.

Lamb, R. (2024). Serious games. *Oxford Research Encyclopedia of Communication*.

Lameras, P., Arnab, S., De Freitas, S., Petridis, P., & Dunwell, I. (2021). Science teachers' experiences of inquiry-based learning through a serious game: A phenomenographic perspective. *Smart Learning Environments*, *8*(1), 7.

Mangelep, N., Tarusu, D., Ester, K., Ngadiorejo, H., & Bumbungan, S. (2023). Local instructional theory: Social arithmetic learning using the context of the monopoly game. *Journal of Education Research*, *4*(4), 1666–1677.

Pahl, C., & Kenny, C. (2008). The future of technology enhanced active learning: A roadmap. In *Technology enhanced learning: Best practices* (pp. 348–375). IGI Global.

Pérez, J., Castro, M., & López, G. (2023). Serious games and ai: Challenges and opportunities for computational social science. *IEEE Access*, *11*, 62051–62061.

Perrotta, C., Featherstone, G., Aston, H., & Houghton, E. (2013). Game-based learning: Latest evidence and future directions. *Slough: NFER*.

Rosli, K., Khairudin, N., & Saat, R. (2019). Gamification in entrepreneurship and accounting education. *Academy of Entrepreneurship Journal*, *25*(3), 1–6.

Scott, K., Baur, L., & Barrett, J. (2017). Evidence-based principles for using technology-enhanced learning in the continuing professional development of health professionals. *Journal of Continuing Education in the Health Professions*, *37*(1), 61–66.

Sedlmair, M., Meyer, M., & Munzner, T. (2012). Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, *18*(12), 2431–2440.

Serrano, D., Dea-Ayuela, M., Gonzalez-Burgos, E., Serrano-Gil, A., & Lalatsa, A. (2019). Technology-enhanced learning in higher education: How to enhance student engagement through blended learning. *European Journal of Education*, *54*(2), 273–286.

Shanklin, S., & Ehlen, C. (2007). Using the Monopoly board game as an in-class economic simulation in the introductory financial accounting course. *Journal of College Teaching & Learning*, *4*(11).

Sidji, M., Smith, W., & Rogerson, M. (2024). Human-AI collaboration in cooperative games: A study of playing Codenames with an LLM assistant. *Proceedings of the ACM on Human-Computer Interaction*, *8*(CHI PLAY).

Solans, D., Beretta, A., Portela, M., Castillo, C., & Monreale, A. (2022). Human response to an ai-based decision support system: A user study on the effects of accuracy and bias. *arXiv preprint arXiv:2203.15514*.

Star, S., & Griesemer, J. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's museum of vertebrate zoology, 1907-39. *Social Studies of Science*, *19*(3), 387–420.