Play-Purchase Dynamics in Virtual Environments

Seung Yoon Lee
The Chinese University of Hong Kong
seungyoonlee@cuhk.edu.hk

K. Sudhir, Kosuke Uetake*
Yale School of Management
k.sudhir@yale.edu, kosuke.uetake@yale.edu

October 2026

Monetizing virtual environments (e.g., games, metaverse platforms) requires understanding how players make interlinked decisions to continue playing and to make purchases. We develop a dynamic structural model of player behavior that captures these play–purchase interactions: players choose whether to continue playing and to purchase in-game tools that increase their probability of winning in both current and future games. These purchases dynamically alter the utility of gameplay in the present and future. Using data from a single-player golf game, we uncover two latent segments: *win-seekers*, who receive higher utility from winning the current game and have lower price sensitivity, and *progress-seekers*, who receive higher utility from level-up rewards and have higher price sensitivity. Counterfactual simulations of game design interventions—such as whom to give free tools and when, or dynamic difficulty adjustments—reveal opposing effects across players and within the same player over time. A free tool can cannibalize purchases, yet help players continue playing and reach more difficult levels where they purchase additional tools. Our findings explain when and why current tools complement or substitute for future purchases and how firms can induce dynamic complementarity by designing the game environment.

Key words: virtual environments, gaming, retention, monetization, personalization, dynamic structural model

---

* We thank the game company (which prefers to remain anonymous) for the data and hosting the first author as an intern. The first author is deeply indebted to Jiwoong Shin and Ryungha Oh for their unstinting support and insights; she thanks Joonhwi Joo, Fei Teng, Ankit Sisodia, and Minkyung Kim for sharing resources and providing helpful feedback. The authors thank Soheil Ghili, Vineet Kumar, Nitin Mehta, Eva Ascarza, and Shubhranshu Singh for their feedback. They also thank participants in the marketing seminars at CUHK, Harvard, HKUST, NUS, Nova SBE, Yale Quantitative Marketing Brown Bag Seminar, and the 2023 AMA-Sheth Foundation Doctoral Consortium for their helpful comments.

1. Introduction

Virtual environments are digital spaces where individuals engage with simulated real or imagined scenarios through various digital interfaces, including computers, smartphones, and virtual or augmented reality (VR/AR) devices. These environments span a wide range of applications, including entertainment, education, and health. Among these, the gaming industry is currently the largest and continues to grow in market size. Worldwide revenues for the video games market are projected to reach approximately $561 billion in 2026 and grow at an average annual rate of 6.3%.[1] Revenue per user for mobile games is estimated at $70 in 2026.[2] Advancements in VR/AR technologies and the growth of metaverse platforms are expected to further accelerate consumer spending in virtual environments.

Virtual environments possess ideal properties for studying consumers' decisions to continue playing and to purchase tools in response to dynamic features of the environment—such as game difficulty and rewards linked to play and progress. First, they operate as self-contained "worlds," where player actions are fully and continuously observable. Second, the environment is controlled, and both players and designers have more transparent knowledge of the reward structure and how it evolves over time. This observability of both behavior and contextual rewards makes virtual environments uniquely well-suited for empirically studying the dynamic interlinkages between play and purchase decisions. Further, purchases are not merely goods that provide immediate utility. When players acquire durable tools, which are items that remain available for repeated use across levels and enhance gameplay performance, they increase their probability of winning in both current and future games. In this way, consumer decisions to continue playing and to purchase tools are dynamically interlinked. That is, tool purchases alter the utility derived from play, which in turn affects future play and purchase decisions.

The transparency of both player choices and the "rules of the game" in virtual environments allows us to study when and why current tools complement or substitute for future purchases. For

[1] https://www.statista.com/forecasts/1344668/revenue-video-game-worldwide/

[2] https://www.statista.com/outlook/amo/media/games/mobile-games/worldwide/

example, suppose the firm discounts the tool today. On the one hand, players may substitute the discounted tool for a purchase they would have otherwise made later on their own. Furthermore, because tools increase players' win probability, they can help players reach higher levels without needing another purchase. On the other hand, tools can help players continue playing and reach more difficult levels, where they may purchase additional tools to increase their win probability. This creates the potential for dynamic complementarity, where an initial tool purchase leads to future purchases. Which of these effects – *dynamic substitution or complementarity* – dominates in aggregate? Could there be opposing effects – across players and within players at different points in time? While the questions above focus on tool giving, the same framework can be applied to other aspects of game design (e.g., dynamic difficulty adjustment, level progression speed) to understand how firms can induce dynamic complementarity by designing the game environment.

To answer these questions, we develop a dynamic structural model that captures how consumers generate utility through play and how this utility evolves with the game environment, as players dynamically and heterogeneously optimize their decisions to continue playing and to purchase tools. The model captures key aspects that are common in such environments: (i) the level progression design that increases in difficulty, resulting in player attrition, (ii) the opportunity to purchase tools that enhance players' productivity, so that players can endogenously balance their level of challenge at a cost, and (iii) short-term rewards for immediate success that sustain interest in the game even as they seek long-term rewards of reaching the next level. Our model extends existing frameworks that have typically focused separately on purchases, as in dynamic durable goods models, and effort or participation decisions, as in incentive compensation models, by modeling play and purchase decisions together. This makes the model applicable not only to gaming but also to the design of other virtual environments, such as digital learning platforms and digital health applications, where purchases improve the productivity of or payoff from future actions.

We estimate the model using data from individual play and purchase choices in a free-to-play single-player mobile golf game. Our comprehensive dataset includes detailed match-level information on players' actions, environments, rewards, and progression throughout their entire gaming

experience. We also obtain detailed records of players' in-game tool purchase transactions, which allow us to examine the relationship between the timing of tool purchases and their impact on player performance. Because player performance depends not only on tool ownership and game difficulty but also on underlying player skill, accounting for heterogeneity in ability is important for our analysis. In many empirical settings, individual ability is difficult to observe directly. Our detailed gameplay records allow us to estimate persistent player ability from early-game outcomes and incorporate this heterogeneity into the model.

Our estimation strategy extends and adapts the two-step estimation framework in Hotz and Miller (1993). First, we estimate player ability from gameplay at early levels and incorporate player ability heterogeneity in the player win probability function and the first stage estimation of the conditional choice probability (CCP). We accommodate latent class heterogeneity and use the expectation-maximization (EM) algorithm within the two-step framework, following the approach in Arcidiacono and Miller (2011). We estimate the structural parameters in the second stage by estimating the value function for each ability and latent segment type. Estimates reveal two latent segments of players: (i) *win-seekers*, who receive higher utility from winning the current game and have lower price sensitivity, and (ii) *progress-seekers*, the largest share of players, who receive higher utility from level-up rewards and have higher price sensitivity.

We use the estimated model to run counterfactual simulations where we examine the dynamic impact of game design interventions to understand when and why current tools complement or substitute for future purchases. Specifically, we simulate the effect of providing free in-game tools—an extreme form of targeted discounting—to isolate how purchases influence future play and purchase decisions. A free tool can cannibalize purchases, yet help players continue playing and reach more difficult levels where they purchase additional tools. We show that higher retention does not necessarily lead to higher monetization: purchases can decrease even when players continue playing more.

Whether a free tool complements or substitutes for future purchases depends on player ability, preferences, and the player's game progression state. We show that even within the same player,

the timing of the intervention (e.g., early vs. late levels) can produce opposing purchase effects, underscoring the value of a structural model to explain these dynamic interlinkages. By jointly adjusting free tool giving and game difficulty, the firm can induce dynamic complementarity even for players whose purchases are cannibalized by the free tool under the current game design. While we illustrate the value of our approach through the free tool-giving and dynamic difficulty adjustment counterfactuals, the framework naturally extends to other dynamic game design elements that exploit the dynamic interlinkages between play and purchase decisions.

The rest of the paper is organized as follows. §2 positions the paper with respect to the related literature. §3 describes the data and the empirical setting. §4 develops the model, and §5 describes the estimation strategy. §6 reports the model estimates, and §7 discusses the findings from counterfactuals. §8 concludes.

2. Related Literature

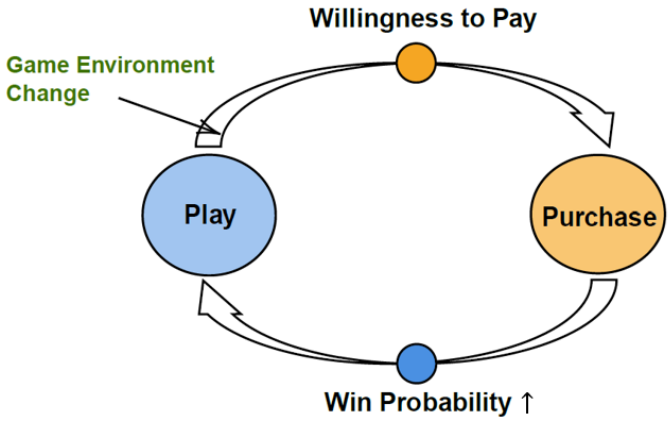
Our paper contributes to several streams of literature. First, we contribute to the growing literature on the video games market. The shift in consumer focus to the online and mobile gaming landscape has enabled researchers to collect and utilize detailed usage records. This has given rise to a stream of empirical papers that study user engagement and gameplay behavior (e.g., Huang et al. 2019, Nevskaya and Albuquerque 2019, Zhao et al. 2022, Castelo-Branco and Manchanda 2023, Chen 2023), monetization strategy and welfare (Appel et al. 2020, Runge et al. 2022, Ascarza et al. 2025, Haenlein et al. 2023, Wang et al. 2023, Amano and Simonov 2024).[3] Despite the significant and growing market size of the gaming industry, there is a lack of research documenting consumer spending behavior on in-game items, which are the most common and major driving source of revenue in the gaming market. Our work is related to the small but growing literature that empirically investigates in-game retention and monetization design involving in-app purchases. Among these, Amano and Simonov (2024) develop a structural model of player play and purchase

[3] There is also a stream of analytical papers studying the design of in-game product monetization (e.g., Guo et al. 2019, Chen et al. 2021, Meng et al. 2021, Harutyunyan and Koca 2024, Miao and Jain 2024, Sheng et al. 2024)

dynamics, focusing on the impacts of loot boxes and their welfare implications—a feature of games that has attracted regulatory scrutiny due to concerns over its resemblance to gambling.

Our paper contributes to this literature in several important ways. To our knowledge, we are the first to articulate and formalize an *active* and *dynamic* consumer utility generation framework that explicitly incorporates the interlinkage of play and purchase decisions in gaming environments. By modeling how durable tools affect continued play and progression, we explain when and why current tools complement or substitute for future purchases. The framework can also inform the design of other virtual environments, such as digital learning platforms and digital health applications, where purchases improve the productivity of or payoff from future actions. We illustrate our conceptual framework in Figure 1. Purchases improve players' probability of winning, changing the utility derived from play. This creates dynamic interlinkages between play and purchase decisions, where purchases change the utility of play, which in turn impacts players' future play and purchase behavior.

Figure 1 The Play and Purchase Loop in the Gaming Environment



Thus, our framework extends previous literature that has examined effort or participation decisions and purchase behavior separately. This includes various contexts such as salesforce management, digital health, online education, and gamified loyalty programs.[4] What differentiates our

[4] For example, an agent's achievement state in relation to sales compensation rewards can motivate performance (effort), which in turn increases monetization for the firm (e.g., Steenburgh 2008, Chung et al. 2014), and sales training can be used to manage salesforce retention and performance (Chung et al. 2021). On the purchase side, previous research has focused on how loyalty programs and status achievement affect consumers' purchase behavior (e.g., Orhun et al. 2022).

empirical setting is its ability to enhance observability of both play and purchases in real-time. By modeling consumers' play and purchase decisions together, our paper presents a comprehensive framework for modeling novel virtual environments.

Unlike traditional digital marketing contexts, digital and virtual environments are redefining how firms interact with consumers by enabling real-time engagement, personalization, and dynamic product (game) design. These self-contained digital spaces allow firms to not only observe consumer behavior but also design and modify incentives, rules, and experiences. Virtual environments, such as gaming, offer a unique advantage by enabling firms to collect granular data on user engagement and monetization in real-time. Within this category, firms can continuously adjust and personalize offerings to optimize engagement and retention. Such findings have consistently been documented in gaming and gamified settings such as e-training and online learning environments (e.g., Santhanam et al. 2016, Huang et al. 2021, Huang et al. 2023, Leung et al. 2023).

In gaming, dynamic difficulty adjustment (DDA) is a widely used design technique where game challenges are adaptively modified in real time based on a player's skill level and behavior (e.g., Hunicke 2005, Xue et al. 2017, Zohaib 2018, Huang et al. 2019, Zhao et al. 2022). In particular, using field experiments, Ascarza et al. (2025) demonstrate that personalized DDA interventions can significantly enhance both player retention and monetization. Our structural approach offers a valuable complement to field experiment approaches for effective DDA. While A/B testing is well suited for estimating treatment effects, structural models provide a theory-driven framework to simulate long-term outcomes across a wide range of counterfactual strategies. They are particularly useful in capturing rich interactions across multiple dimensions of player heterogeneity—such as ability, price sensitivity, and preferences for progression—and in accounting for the dynamic and cumulative effects of interventions over time. This capability allows designers to anticipate the behavioral impact of different strategies and prioritize those most likely to succeed. This can make subsequent downstream experimentation more targeted and cost-effective.

Our work is also related to reinforcement learning (RL) approaches to dynamic policy optimization. For example, Ma et al. (2026) combine state representation learning with offline RL

to optimize personalized difficulty adjustments in mobile gaming. While standard RL methods focus on learning policies that maximize a specified reward objective, our goal is to estimate the behavioral preferences underlying players' play and purchase decisions and use these estimates to understand their responses to alternative interventions. In this respect, our approach is closely related to recent work connecting inverse reinforcement learning (IRL) with structural dynamic discrete choice estimation. Kang et al. (2025) develop a gradient-based IRL approach to estimating dynamic discrete choice models, while Lee et al. (2026) use adversarial IRL to estimate a dynamic choice model with unobserved heterogeneity. These methods offer complementary approaches to estimating behavioral primitives.

3. Data and Empirical Setting

This section describes the game, provides details of the data, and presents model-free evidence that informs our model development.

3.1. Description of the Game

Our empirical setting is a popular free-to-play single-player mobile golf game with over two million registered users. In this game, players engage in one-hole game matches, where the objective is to complete the hole with fewer shots than the opponent. Each game lasts around three to five minutes, and players are assigned their matches by the game platform once they enter the game. Players accumulate points from winning the match (and lose points from losing), and the collection of these points is required to unlock higher levels in the game. The game design ensures a sequential progression where higher levels demand the accumulation of more points to be unlocked. The game has a total of 11 levels.

The level-progression system of the game is designed to increase difficulty through several design features. We report the points system of the game in Table 1. The expected points for each game given a player's win probability are calculated based on the game's win-loss points schedule, which varies by level. We illustrate the level difficulty design embedded within this points system using levels 6 to 11 as examples in Figure 2. In the figure, each line represents a different level, demonstrating that as players advance to higher levels, the expected points for a given win probability

generally decrease under the same win rate. Figure 3a helps further illustrate this point by highlighting that as the level increases, the win probability required to at least break even (i.e., zero expected points) also increases, requiring players to have higher win rates at subsequent levels. Finally, higher levels require a higher points accumulation criterion to level up, making progression increasingly demanding, as shown in Figure 3b. Overall, the points system of the game is designed to discriminate on player ability, ensuring that only the most able players progress to the top levels.

Table 1 Points and Level Progression Design of the Game

Level	Win	Lose	Total Available Points	Cumulative Points Collection
1	+4	-1	25	-
2	+6	-2	75	25
3	+8	-3	125	100
4	+10	-4	175	225
5	+12	-7	225	400
6	+14	-10	300	625
7	+16	-13	375	925
8	+18	-16	450	1300
9	+20	-20	550	1750
10	+22	-24	700	2300
11	+24	-36	900	3000
Final Lvl Clear				3900

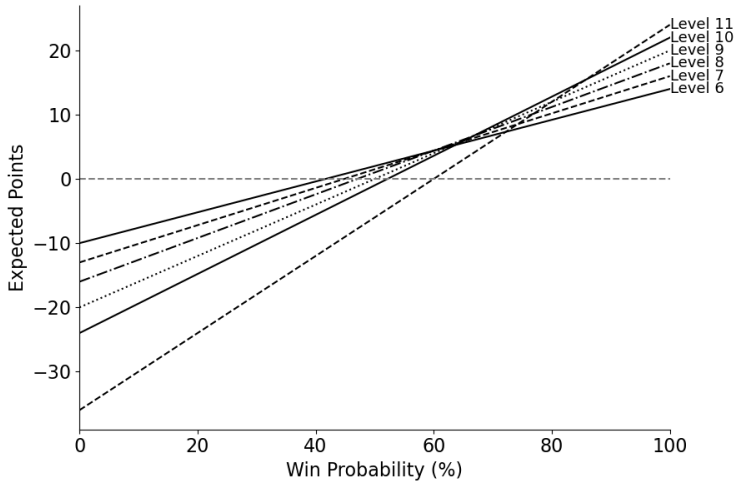
Note: The *Win* and *Lose* columns indicate the number of progression points awarded or deducted for winning or losing a game at each level. *Total Available Points* refers to the number of points a player needs to collect at that level to level up. *Cumulative Points Collection* indicates the total number of points a player has accumulated upon entering that level. Once a level is unlocked, the player's progression points do not fall below the corresponding cumulative threshold.

Players can enhance their win rates by purchasing tools. The in-app purchases consist of durable ability enhancers (i.e., golf clubs) that allow players to improve their win rates. For example, a paid golf club enables greater range and ball guide precision when taking a shot. By incurring a monetary cost, players can increase their chances of progressing through the levels that they might not have had otherwise. By allowing players to self-select the balance between the need to win and the cost of purchasing ability enhancers, the firm can effectively monetize across different player segments of different abilities and price sensitivity. The most popular in-game tool offerings cost $9.99 (generating around 60% of revenue), and around 90% of total durable tool transactions are

generated from product offerings between \$9.99 and \$19.99.[5] Higher-priced offerings provide higher quality golf clubs that can help improve player win rates.[6] While the game also features in-app advertising, the firm primarily generates revenue through in-app purchases.

The game offers a suitable setting for a single-agent model. From the player's perspective, it is effectively a single-player experience, in that in the game of golf, players' performance—the number of shots players take to complete a hole—is independent of the opponent's actions, making the opponent an exogenous factor influencing the game outcomes. The players do not strategically select which opponent to play against, but rather get assigned by the platform.[7] We observe opponent Elo for each match and control for it when estimating the win probability function (see Online Appendix OB.6 for additional analyses of opponent assignment). The progression system of the game incentivizes players to continually improve their skills to maintain or increase their win rates and progress through the levels, inducing dynamic and forward-looking behavior.

Figure 2 Game Points Design: Expected points by Win Probability by Level



Note: Expected points are calculated as $\mathrm{E}[\text{Points}] = p \cdot \psi_v^\ell + (1-p) \cdot \psi_d^\ell$, where p is the win probability, ψ_v^ℓ is the number of points awarded for a win, and ψ_d^ℓ is the number of points deducted for a loss for each level ℓ.

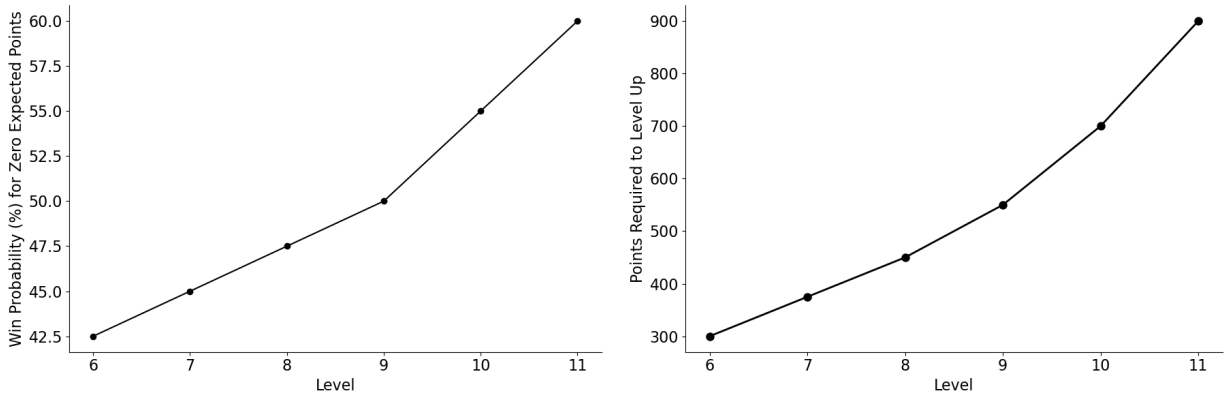
[5] Certain product offerings involve randomness in the type of tool received upon purchase, but players are guaranteed to receive a high-quality tool. Differences across randomly received tool types in subsequent win rates and in the probabilities of play and purchase are not statistically significant (see Online Appendix OB.16).

[6] While higher-priced offerings provide slightly improved tools, the difference between the \$9.99 and \$19.99 product offerings is not substantial; thus, we treat all durable tool purchases as a single product and estimate an average effect in the model (see Online Appendix OB.15).

[7] In the game, around 60% of gameplay matches involve bots, but players cannot distinguish whether their opponents are real players or computer-simulated.

Figure 3 Game Points and Level-Progression Design

(a) Win Probability for Zero Expected Points by Level **(b) Points Collection Requirement to Level-Up**



3.2. Data

We leverage a comprehensive dataset comprising every player action and state in the game environment, spanning each player's lifetime from initiation to exit. This includes detailed observations of players' decisions to play or quit, progression, rewards, and the gaming environments they encounter, including the opponents they face. Additionally, we have detailed records of players' in-game tool purchase transactions, enabling us to examine the relationship between the timing of in-game purchases, game environment conditions, and player performance.

Our analysis uses data from 4,163 players spanning a 15-month period, from October 2021 to January 2023. We construct the sample as follows. First, we take a random 10% sample of players who meet the criteria of having valid play records in the data period. Second, players who have completed at least level 5 (i.e., collected at least 15% of total points available) are retained. This cutoff is motivated by our estimation strategy: we use levels 1–5 to estimate persistent player ability and then estimate dynamic play, purchase, and exit decisions from level 6 onward. Our structural estimates therefore describe players who exhibit sufficient engagement to reach level 6, excluding less committed players who discontinue playing earlier. Finally, we exclude outliers with unusually long within-level gameplay – defined as those exceeding 1.5 times the interquartile range above the upper quartile of within-level number of plays. The final dataset for analysis includes around 750,000 match records.

Table 2 contains summary statistics of our player sample. A median player engages in 273 games throughout their lifetime, and the median exit level is 7.[8] The average game play duration, measured by the number of days from first play to exit, is 106.21 days. Finally, given our sample, which includes players with a minimum progression to level 6, 22.36% have purchased tools at least once. While the median player makes no purchases, there is significant heterogeneity in the player purchase behavior at the higher end of the distribution, with total transactions ranging between 14 and 41 for players above the 99th percentile.

Table 2 Descriptive Statistics of Players

	Mean	SD	Median	75th	90th	99th	Max
Play							
Total Number of Games Played	330.57	225.18	273	405	582	1202	3210
Maximum Level Reached	7.41	1.54	7	9	10	11	11
Game Duration (*Days from First Play to Exit*)	106.21	98.55	73	151	259	405	454
Purchase							
In-app Purchase Player Share	22.36%						
Total Number of Purchases	0.88	2.69	0	0	3	13	41

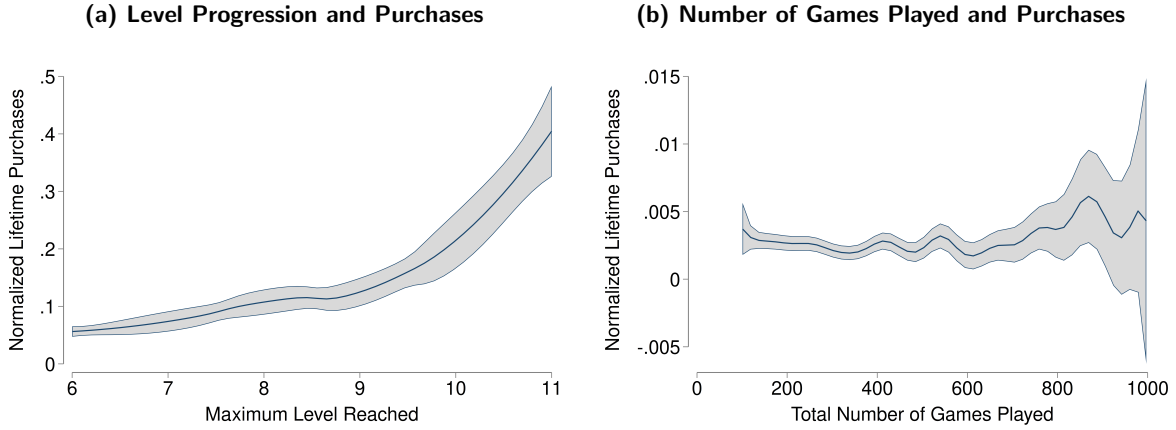
3.3. Model-Free Evidence

We begin this section by providing descriptive evidence on the relationship between players' play and purchase decisions. To account for the fact that players who play more games have more opportunities to make purchases, we normalize total purchases by the maximum level reached and total play counts. We plot players' final level progression state with their normalized lifetime number of purchases in Figure 4a. We show that players who have reached higher levels are also those who spend more, providing suggestive evidence consistent with the core dynamics of our play and purchase framework. We note that the slope becomes steeper at later levels, where the game becomes more difficult. In contrast, Figure 4b presents normalized purchases by game and total number of games played, but reveals no clear pattern – suggesting that purchases are driven more by progression through increasingly difficult levels than by the total number of games played

[8] The two-week condition applies to inactivity at the end of our observation period. Among players in our estimation sample who had not completed the game, 92.7% had no subsequent play for at least 30 days. We provide further details on sample selection in Online Appendix OB.1.

alone. We report versions of Figure 4 without normalization or smoothing and regressions of total purchases on level progression and games played in Online Appendix OB.5.

Figure 4 Relationship Between Player Retention and Monetization



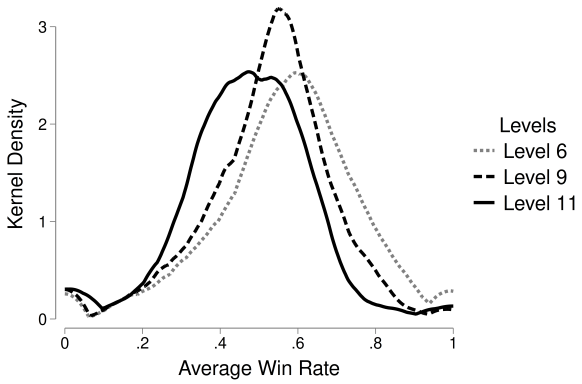
Note: The solid line is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

Next, we present three key features of the data that inform our model development. First, we provide evidence of substantial heterogeneity in player win probability within and across levels. Second, we show the relationship between players' win rates and their purchase and exit decisions. Third, we demonstrate the dynamics in the timing of players' purchase and exit decisions in relation to their level completion status.

In the gaming environment, there exists substantial heterogeneity in player win rates, reflecting their ability to progress through the game levels. We present the differences in average win rate across players for different levels in Figure 5. First, players have generally lower win rates at higher levels, consistent with the increased level difficulty designed into the game. Within each level, there exist significant differences in player win probabilities, providing suggestive evidence of the varied abilities of players even at the same level. We report the Gini indices of player average win rate within each level in Online Appendix OB.3.

We next examine the relationship between players' previous-level win rates and their purchase and exit decisions, as illustrated in Figure 6. To mitigate reverse causality (i.e., players purchasing tools which then increase their win probability), we use win rates from the previous level when

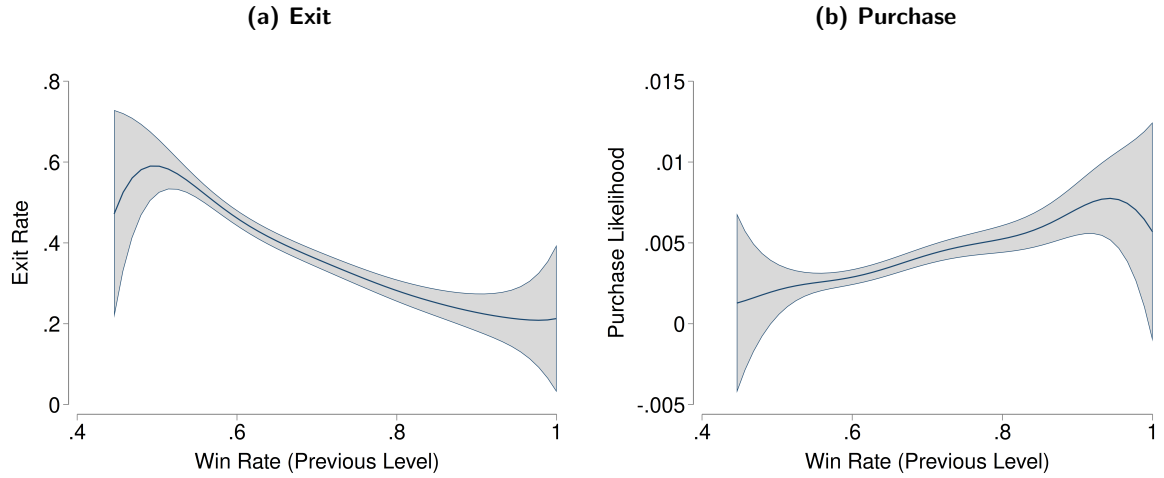
Figure 5 Average Win Rate Distribution of Players By Levels



Note: The lines represent kernel density estimates using the Epanechnikov kernel.

plotting players' purchase and exit decisions. Within each level, we find that players with higher previous-level win rates are more likely to continue playing and make in-game purchases. Conversely, players with lower win rates are more likely to exit the game and purchase less (see Online Appendix OB.3 for the full set of graphs by level). We present these findings as model-free descriptives, as unobserved player heterogeneity may influence both previous-level win rates and current purchase and exit decisions.

Figure 6 Player Exit and Purchase by Win Rate



Note: The solid line is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

Finally, Table 3 presents players' purchase and exit likelihood in relation to their level completion status. Players are more likely to exit earlier within a level, and exit probability decreases as players

progress through that level. Purchases are also more likely to occur earlier in the level. We note that these results are model-free descriptives of how exit and purchase probabilities vary with players' states. Players observed later within a level are those who continue playing and progress further, so the results do not by themselves establish the underlying dynamic behavioral mechanisms.

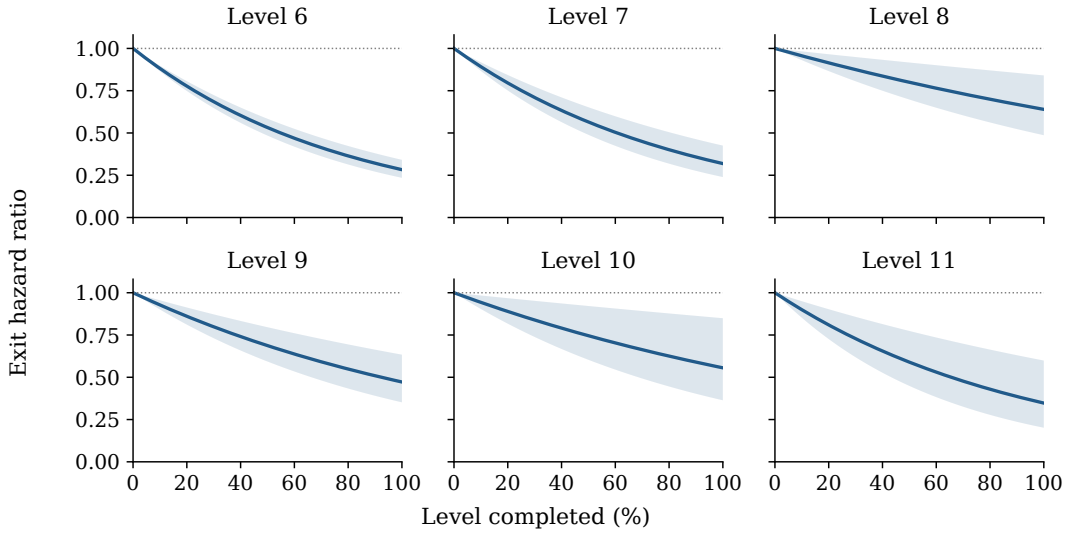
Table 3 Relationship Between Level Completion and Exit/Purchase Probability

	Pr(Exit)	Pr(Purchase)
% Level Complete	-0.934^{***}	-0.153^{**}
	(0.059)	(0.067)
Tool stock	-0.008	0.199^{***}
	(0.007)	(0.004)
Ability	-0.703^{***}	1.759^{***}
	(0.116)	(0.160)
Constant	Y	Y
Level FE	Y	Y
N	747,800	747,800

Note: Logistic regression estimates with robust standard errors in parentheses. Ability is the normalized player fixed effect estimated using games played at levels 1–5 (see Section 5.1 for details). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

We further illustrate the decline in exit as players approach level-up in Figure 7, which reports estimates from Cox proportional hazards models for each level. The declining exit hazard is consistent with players taking future level-up into account and motivates modeling players as forward-looking.

Figure 7 Exit Hazard by Percentage Level Complete: Estimates by Level



Note: Estimates using a separate Cox proportional hazards model for each level, including all players observed at that level. Time t is measured by the number of games. Controls include ability and tool stock. Hazard ratios are relative to 0% level completion, and the shaded regions show 95% confidence intervals using standard errors clustered by player.

4. Model

Based on the model-free evidence, we develop a dynamic model of player action in the gaming environment. Time indexed by t is discrete, denoting each game.[9] At the beginning of each period t, player i decides whether to (1) play, (2) make a purchase and play, or (3) exit the game permanently. To represent player states in the gaming environment, we parsimoniously track two key state variables $S_{it} = \{k_{it}, z_{it}\}$. First is player tool stock k_{it}, which tracks the number of in-app purchases a player incurred to improve their probability of winning. A higher tool stock corresponds to an upgrade in the quality of players' tools, which are durable and remain available for repeated use across games. Second is player progression points stock z_{it}. The points accumulation state (z_{it}) maps to the level (ℓ_{it}) according to the game's points and level progression design Ψ in Table 1. The points state therefore represents both the player's level and progress within that level, so a separate state variable for level is not needed.

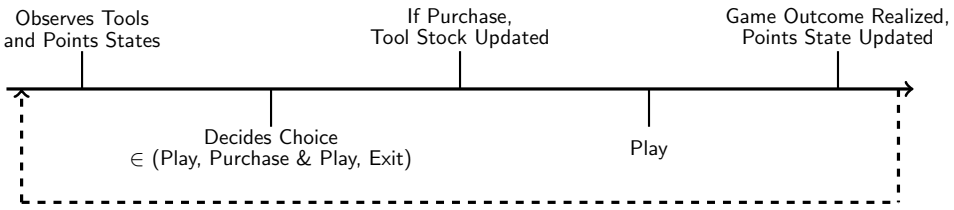
Figure 8 describes the timeline of the model. Before entering a game t, player i observes the current tool stock and points accumulation state. The player decides on the action, including playing, purchasing a tool and then playing, and exiting the game permanently. If the player's decision is to purchase a tool, the tool stock is updated before play. The player then enters the game, and an idiosyncratic game outcome shock is realized. The realized game outcome affects the player's progression points state in the next period. The model repeats every period over an infinite horizon.[10]

The game's point system and progression design induce dynamic, forward-looking behavior in players. A player choosing the action considers not only the current payoff but also the effect of that choice on all future costs and rewards.

[9] We define time t as a game, rather than calendar time – an approach consistent with other dynamic models of player behavior (e.g., Zhao et al. 2022, Amano and Simonov 2024). Tools increase win probabilities and can reduce the number of games needed to reach a given level. Our model focuses on this progression and players' decisions to continue playing and purchasing. Thus, players discount future utility based on the number of games played rather than the passage of real time. We abstract from when players return after a temporary break or how the length of that break affects purchases, and acknowledge this as a limitation of the game-based time formulation. We report the distribution of inter-play gaps and their relationship with purchase behavior in Online Appendix OB.4.

[10] Players can continue to play the unlocked stages after completing all levels. In our estimation sample, 1.2% of players complete all levels. Among these players, 97.9% continue playing afterward, and the average number of additional games is 230.5.

Figure 8 Model Timeline



4.1. Player Utility Generation Process

Player i in period t receives the following per-period utility based on his or her choice of actions $d_{it} \in \{1,2,3\}$: whether to play the game without making an in-app purchase ($d_{it}=1$), whether to make an in-app purchase and play the game ($d_{it}=2$), or whether to permanently exit the game platform ($d_{it}=3$).

$$u(S_{it}, d_{it}, \epsilon_{it}) = \begin{cases} c_p + \theta p_{it}^{d_{it}} + R \cdot \mathbf{1}_{\text{LEVEL UP}(z_{it})} \cdot L(z_{it}) + c_m \cdot \mathbf{1}_{\{d_{it}=2\}} + \epsilon_{dit}, & \text{if } d_{it} \in \{1,2\}, \\ \epsilon_{3it}, & \text{otherwise (exit).} \end{cases} \tag{1}$$

The parameters c_p, θ, R, and c_m vary across latent segments; we suppress segment subscripts here and introduce them in estimation. A purchased tool is available for the current game, so the win probability uses tool stock $k_{it} + \mathbf{1}_{\{d_{it}=2\}}$. The indicator $\mathbf{1}_{\text{LEVEL UP}(z_{it})}$ identifies a level-up state. At this state, $L(z_{it}) = L_\ell$, where ℓ is the level just completed and L_ℓ is its points requirement.

In each period, if the player decides to play, the player incurs a cost of play c_p that captures the net utility or disutility from engaging in gameplay, independent of the game outcome. For example, a negative c_p suggests that the time cost of play outweighs the enjoyment. In every play, the player also receives utility from winning the current game, where θ represents the player's utility from a win and p_{it}^d is player i's probability of winning under action $d \in \{1,2\}$.

In addition to the immediate utility from winning, the player receives a level-up reward every time the player reaches a new level, where R represents the intrinsic utility from reaching the next level. We multiply R by the level achievement criterion $L(z_{it})$ (the points requirement for the level just completed).[11] The level-up points criterion increases monotonically with levels in the current

[11] For example, the level-6 requirement is 300 points in the "Total Available Points" column of Table 1. On reaching level 7, $L(z_{it}) = 300$, so the level-up utility is $R \times 300$. The player receives the level-up utility only once for each new level reached.

game design. This allows us to parsimoniously model the increasing utility from leveling up at higher (more difficult) levels.[12]

If the player decides to purchase a tool ($d_{it} = 2$), he or she incurs a monetary cost c_m. The decision to purchase updates the player's current tool stock state, which increases the chance of winning, including the immediate period and all future play sequences of the game. The idiosyncratic shock ϵ_{it} follows a Type-I extreme value distribution. We normalize the exit value as 0.

4.1.1. Player Win Probability Function We model the player winning outcome $W_{it} \in \{0, 1\}$ to be a function of player level $\ell(z_{it})$, tools k_{it}, and player ability α_i, such that

$$p_{it}^d = \Lambda\left(\gamma_{\ell(z_{it})} + \tau\alpha_i + \delta_{\ell(z_{it})}\left(k_{it} + \mathbf{1}_{\{d=2\}}\right) + \lambda o_{it}\right), \qquad \Lambda(x) = \frac{\exp(x)}{1+\exp(x)}, \tag{2}$$

where $\Lambda(\cdot)$ is the logistic function.

The player's win probability function accommodates the following key characteristics to model the gameplay outcome.[13] First, the level-specific intercepts γ_ℓ allow win probabilities to differ across levels, with the player's level $\ell(z_{it})$ determined by the point accumulation state z_{it}. Second, the player's win probability increases with both the tool stock k_{it} and the player's innate ability α_i. We also allow for an interaction between tools and level difficulty to account for differences in tool effects across levels. Finally, we control for opponent ability o_{it}, measured by the opponent's Elo score observed for each match, with coefficient λ in the player win probability function. The logistic form allows the marginal effect of a tool to vary with player ability and level difficulty. Increasing marginal effects in its convex region can support complementarity through the immediate win payoff in flow utility, whereas diminishing marginal effects in its concave region can support substitution. Note that dynamic complementarity is also possible because an additional tool changes players' expected play and progression trajectories and, therefore, their continuation values. We discuss these mechanisms in Section 7.1.

[12] In our game setting, the points progression system of the game ensures that boredom is not a concern: players who win faster face the next challenge level sooner, and level 11 is difficult for all players.

[13] The win probability function gives the probability of winning, $W_{it} = 1$, conditional on player level, tool stock, and player ability, controlling for the effect of opponent ability. In our setting, player ability and tool stock are productive inputs into game performance, while match-level uncertainty makes the realized win outcome stochastic.

There are a few points to note about the win probability function. First, $\gamma_{\ell(z_{it})}$ captures not only level difficulty but also the learning that may occur across levels, reflecting how players improve through repeated play. If learning occurs, players should experience higher win probabilities as they progress. However, the game design that increases difficulty at higher levels counteracts this effect, and $\gamma_{\ell(z_{it})}$ captures this net effect of increased difficulty and player improvement in ability. In addition, $\delta_{\ell(z_{it})}k_{it}$ accounts not only for the effect of tools at different levels but also for the possibility that k_{it} might reduce the learning as net effects.[14]

4.2. State Transitions

The state variables tool stock k_{it} and points stock z_{it} evolve deterministically, conditional on the player's action and the game outcome, as follows:

$$k_{i,t+1} = \begin{cases} k_{it}+1, & \text{if } d_{it}=2, \\ k_{it}, & \text{otherwise.} \end{cases} \tag{3}$$

The player's tool stock k_{it} increments by one with each purchase decision. Players can make multiple purchases over their lifetime.

The per-period points reward r_{it} is determined by the game winning outcome $W_{it} \in \{0,1\}$ and the winning and losing points reward design of the game, $(\psi_v^\ell, \psi_d^\ell) \in \Psi$, which differs by each level ℓ. That is,

$$r_{it} = W_{it} \cdot \psi_v^\ell + (1 - W_{it}) \cdot \psi_d^\ell. \tag{4}$$

Here $\psi_v^\ell > 0$ and $\psi_d^\ell \leq 0$, with $\ell = \ell(z_{it})$.

$$z_{i,t+1} = \begin{cases} \min\left(z_{it} + \psi_v^{\ell(z_{it})}, \overline{\psi}^L\right), & W_{it}=1, \\ \max\left(z_{it} + \psi_d^{\ell(z_{it})}, \underline{\psi}^{\ell(z_{it})}\right), & W_{it}=0. \end{cases} \tag{5}$$

Upon realization of the game outcome, the player's points stock state z_{it} evolves deterministically following the game's points progression design Ψ. Here, $\overline{\psi}^L$ is the maximum points state of the

[14] After accounting for ability, level, and tool effects, we find that the incremental variance explained by player learning by doing is only around 5%. This is because much of the player learning is already absorbed as a net effect of level progression. Given that our current specification captures 94.6% of the incremental variance, incorporating learning by doing would dramatically increase the state space without significant explanatory gains. Therefore, we abstract away from it and keep our state space tractable (see Online Appendix OB.9 for details).

game at the final level, and $\underline{\psi}^{\ell}$ is the points criterion to unlock level ℓ. This ensures that once a player has reached a certain level by accumulating enough points, they cannot fall back below that level's threshold, ensuring that once a level is unlocked, it remains unlocked. In other words, z_{it} is incremented by ψ_v^{ℓ} when winning the game, but cannot exceed $\overline{\psi}^{L}$. Similarly, z_{it} is reduced by $|\psi_d^{\ell}|$ but not below the minimum threshold $\underline{\psi}^{\ell}$.

4.3. Bellman Equation

A player chooses action d_{it} that maximizes the expected discounted sum of utilities given the game design Ψ, the state variables and their transitions, and the idiosyncratic shock ϵ in each period. The Bellman equation can be written as

$$V(k,z,\alpha,\epsilon;\Theta,\Psi) = \max_{d\in\{1,2,3\}} \Big\{ u_d(k,z,\alpha,\epsilon;\Theta,\Psi) + \beta\, \mathbf{1}_{d\in\{1,2\}}\, \mathbb{E}\big[V(k',z',\alpha,\epsilon';\Theta,\Psi) \mid k,z,\alpha,d\big] \Big\}, \quad (6)$$

where $\Theta = (c_p, c_m, \theta, R)$ is the vector of utility parameters for the player's segment, the idiosyncratic shock ϵ follows a Type-I extreme value distribution, and the discount factor β is set to 0.9.[15] The player continues to play the game if the value of playing or purchasing and playing is greater than the value of exit.

5. Estimation

We estimate the model using two-step estimation (Hotz and Miller 1993). In the first step, we estimate player ability and the player win probability function, then the conditional choice probabilities (CCPs) of player action choice as a flexible function of state variables and ability. The key assumption in the two-step estimation is that the first-stage CCPs represent the agent's optimal action probability given the state variables and ability. In the second step, we estimate the structural parameters that rationalize the first stage policy estimates.

[15] $\beta = 0.9$ gives the best fit compared to $\beta = 0.95$, 0.99, and 0.995. In our context, setting higher discount factors greatly underestimates player exit probabilities and hence overestimates player purchase probabilities. The average number of inactive days between play dates is 1.4, and the median is zero. Average lifetime play duration is 106 days.

5.1. Step 1: Estimating CCPs

In the first stage CCP estimation, we estimate a flexible mapping between observable states and player action probability. The relevant state variables in our model are tool and points stock states, $S_{it} = \{k_{it}, z_{it}\}$. We estimate the two-step procedure using the player state space starting from level 6, using win outcomes at levels 1–5 to estimate player ability.[16] The early levels 1–5 consist of beginner level tutorials and short level length, with these five levels accounting for only 16% of total points collection available in the game. We discretize the points state for each level from 6 to 11 into 10 increments and additional level-up states to track the level-up reward (a total of 65), and tool stock state to evolve deterministically up to 25 transactions. This leaves us with a total of 1,690 state combinations for each ability bin and segment.

Typically, player ability is treated as an unobserved variable. Because our setting is a game, where we observe every player action, environment, and outcome, we can estimate player ability as an individual fixed effects parameter from a linear probability regression using levels 1–5, controlling for the effects of tools and game environments on player win rates. For the first stage estimation, we normalize player ability score as a continuous variable between 0 and 1, with 1 representing the highest ability level.[17] We then include the normalized ability index α_i in the win probability function in Equation 2 and the first stage policy function, both estimated using the main estimation data. This enables us to account for rich heterogeneity in player ability levels in predicting player action.

Given the state variables and estimated player ability, we estimate the player action policy using a flexible multinomial logistic regression. We account for player unobserved heterogeneity in the first stage CCP estimation through persistent latent segments and estimate heterogeneous policy functions using the EM algorithm (Arcidiacono and Miller 2011, Chung et al. 2014). We assume

[16] Our estimates and counterfactual results are qualitatively robust to the alternative ability and model estimation windows (see Online Appendices OB.8 and OB.13). For alternative win probability and utility specifications, see Online Appendices OB.14 and OB.17–OB.19.

[17] Ability is normalized as $\alpha_i = (a_i - a_{\min})/(a_{\max} - a_{\min})$. Normalization does not change the estimated win probabilities or CCPs (see Online Appendix OB.10).

that player i belongs to one of G segments $g \in \{1, \ldots, G\}$, with posterior segment probabilities $q_i = \{q_{i1}, \ldots, q_{iG}\}$. Let π_g denote the population probability of being in segment g. Player i's segment g_i represents a persistent structural preference type and is assumed to be policy-invariant. Posterior segment probabilities may update as additional behavior is observed. We clarify that changes in posterior classification do not imply that the underlying preference type changes. We iteratively maximize the log likelihood,

$$\log L = \sum_{i=1}^{N} \log \left[\sum_{g=1}^{G} \pi_g \prod_{t=1}^{T_i} \mathcal{L}(d_{it} \mid S_{it}, \alpha_i, g, \kappa_g) \right], \tag{7}$$

where T_i is the number of observed decisions for player i, and

$$q_{ig} = \frac{\pi_g \prod_{t=1}^{T_i} \mathcal{L}(d_{it} \mid S_{it}, \alpha_i, g, \kappa_g)}{\sum_{h=1}^{G} \pi_h \prod_{t=1}^{T_i} \mathcal{L}(d_{it} \mid S_{it}, \alpha_i, h, \kappa_h)}. \tag{8}$$

The choice probability of taking action $d_{it} = j$ for segment type g is

$$\mathcal{L}(d_{it} = j \mid S_{it}, \alpha_i, g, \kappa_g) = \frac{\exp\{X(S_{it}, \alpha_i)' \kappa_g^j\}}{\sum_{\tilde{j}=1}^{3} \exp\{X(S_{it}, \alpha_i)' \kappa_g^{\tilde{j}}\}}, \qquad \kappa_g^1 = 0. \tag{9}$$

Here, $X(S_{it}, \alpha_i)$ describes the player's tool stock, current points progression state, and ability, and κ_g^j is the corresponding coefficient vector for action j in segment g. We normalize the coefficients for playing without a purchase to zero.

The EM algorithm begins by setting initial values for κ_g and π_g.

(a) Compute $q_{ig}^{(m+1)}$ using Equation 8 with $\kappa_g^{(m)}$ and $\pi_g^{(m)}$.

(b) Update population shares:

$$\pi_g^{(m+1)} = \frac{1}{N} \sum_{i=1}^{N} q_{ig}^{(m+1)}.$$

(c) Update $\kappa_g^{(m+1)}$ for each segment g by maximizing the weighted log likelihood,

$$\sum_{i=1}^{N} q_{ig}^{(m+1)} \sum_{t=1}^{T_i} \log \mathcal{L}(d_{it} \mid S_{it}, \alpha_i, g, \kappa_g).$$

We iterate steps (a)–(c) until convergence. We initialize κ_g by randomly partitioning the players into G segments and maximizing the log-likelihood, and population shares to be $1/G$.

From this iterative estimation step, we obtain segment-specific policy function parameters, along with the population segment probability estimates π. We use the segment-level policy functions, together with posterior segment probabilities, to obtain structural parameters of each segment, which we describe in the next section. A caveat with the two-step estimation is that the first-stage policy function estimates can be biased if the state variables in the policy function are correlated with the first-stage errors. Our approach, which leverages player-level ability estimates and unobserved player heterogeneity, helps mitigate this issue.

5.2. Step 2: Structural Parameter Estimation

The key idea of the two-step estimation is to represent the value function in terms of the policy function estimated in the first stage, which reflects the player's optimal actions. For second-stage estimation, we discretize continuous player ability into 20 bins and solve the value function for each ability bin and segment. Given our discrete state space, we can solve the value function as a system of linear equations,

$$V^{\widehat{P}}(\Theta_g) = (I - \beta F^{\widehat{P}})^{-1} \sum_{d=1}^{3} \widehat{P}_d * [\bar{u}_{g,d}(\Theta_g) + E[\epsilon_d \mid d]], \tag{10}$$

where $*$ denotes element-by-element multiplication, $\widehat{P}_d$ is the vector of first-stage choice probabilities for action d, and $\bar{u}_{g,d}$ is flow utility excluding the idiosyncratic shock. The transition matrix $F^{\widehat{P}}$ weights the transition matrices for play and purchase by their first-stage choice probabilities (Pesendorfer and Schmidt-Dengler 2008). The Type I extreme value assumption of the error term allows us to solve the value function analytically, such that $E[\epsilon_d \mid d] = \gamma - \ln \widehat{P}_d$. γ is Euler's constant.

Furthermore, we can express the player's choice probability of action d_{it} under the structural parameters of our model in closed form using the distribution assumption of the errors as follows:

$$\widetilde{\Pr}(d \mid S, \alpha, g; \Theta_g, \Psi) = \frac{\exp(v_d(S, \alpha, g))}{\sum_{\tilde{d}=1}^{3} \exp(v_{\tilde{d}}(S, \alpha, g))}, \tag{11}$$

where $v_d = \bar{u}_{g,d} + \beta F_d V^{\widehat{P}}$ for play and purchase, and $v_3 = 0$ for exit.

We construct the minimum distance estimator, where $\widehat{\Pr}$ is the optimal policy estimated from the first stage, and $\widetilde{\Pr}$ is the policy informed by the model parameters (Hotz and Miller 1993). We minimize Equation 12, the distance between the optimal policy and the model choice probabilities, weighted by player posterior segment probabilities q_{ig} from the first stage.

$$\sum_{i=1}^{N}\sum_{t=1}^{T_i} q_{ig}\Big\{\Big[\widehat{\Pr}(d_{it}=2\mid S_{it},\alpha_i,g;\widehat{\kappa}_g)-\widetilde{\Pr}(d_{it}=2\mid S_{it},\alpha_i,g;\Theta_g,\Psi)\Big]^2 + \Big[\widehat{\Pr}(d_{it}=3\mid S_{it},\alpha_i,g;\widehat{\kappa}_g)-\widetilde{\Pr}(d_{it}=3\mid S_{it},\alpha_i,g;\Theta_g,\Psi)\Big]^2\Big\}. \tag{12}$$

We keep the first-stage estimates fixed during structural estimation.

To compute standard errors, we generated 500 bootstrap datasets following Bajari et al. (2007), sampling players with replacement and using all gameplay records for each selected player. For each bootstrapped dataset, we estimate player ability, the win probability function, and both the first and second stages to account for the estimation errors in ability and the first stage policy estimation.

5.3. Identification

There are a few challenges in identifying the dynamic structural model of game playing. First, the intrinsic player ability is not observed. Our long panel of gameplay records and the variations in environments, tool stock, and winning outcomes across and within players allow for the identification of player ability through player fixed effects estimated using levels 1–5, as described in Section 5.1. We follow Kasahara and Shimotsu (2009) for the identification of unobserved finite mixture heterogeneity. Under their conditions on state transitions and variation in choice probabilities across segments, three periods of panel data can be sufficient for identification. Our panel data can easily satisfy the condition.[18]

Second, the parameters in the flow utility function need to be identified by the revealed preference argument. We fix the discount factor to 0.9. The likelihood of playing and exit across states jointly identifies θ and c_p as win probabilities vary across states, while c_p is a constant. The disutility of

[18] Our free tool giving counterfactual findings are qualitatively robust to using fewer observations for segment classification (see Online Appendix OB.12).

purchase parameter c_m is identified from the play versus play and purchase decisions across states. Lastly, the intrinsic utility from level-up R is identified from player behavior (i.e., exit/purchase decision) at the threshold between the level-up points, as the level-up utility in non-level-up periods does not affect current utility, only future utility.

6. Results

We discuss our results in the following order: 1) player win probability function and ability estimates, 2) first stage policy estimates, and 3) second stage structural parameter estimates.

6.1. Player Win Probability Function and Ability Estimates

We estimate the player win probability function using a logit model of player ability, tool stock, and level states, controlling out the opponent effect.[19] The estimated player win probability function, shown in Table 4, accommodates the following three key features. First, win probability is higher for higher ability players. Second, player win probability increases with tools.[20] With the logistic specification, the marginal effect of an additional tool varies with player ability, tool stock, and level. Third, player win probabilities are heterogeneous across levels.

To demonstrate the reliability of our ability measure, we present the relationship between player average win probability and ability in Figure 9a. There is a clear positive relationship, with higher ability players exhibiting higher win rates. Figure 9b displays the distribution of the player ability estimates, highlighting significant heterogeneity in our player sample.

6.2. First Stage Policy Function Estimates

Given the ability estimates of players, we estimate the first stage policy function of player action probabilities using a flexible multinomial logistic regression, accounting for both observed and unobserved heterogeneity. We find that the two-segment model best fits the data based on the BIC criterion.[21] We report the first stage policy estimates for each segment in Table 5. Segment

[19] For the second stage structural parameter estimation, we control out the effect of opponents using the mean value of the opponent Elo scores – the rating system used by the company for calculating the relative skill levels of players – at each level.

[20] To provide additional evidence of the effect of tool purchase on player win probability, we conduct a more localized before-and-after analysis of player win rates in Online Appendix OB.7.

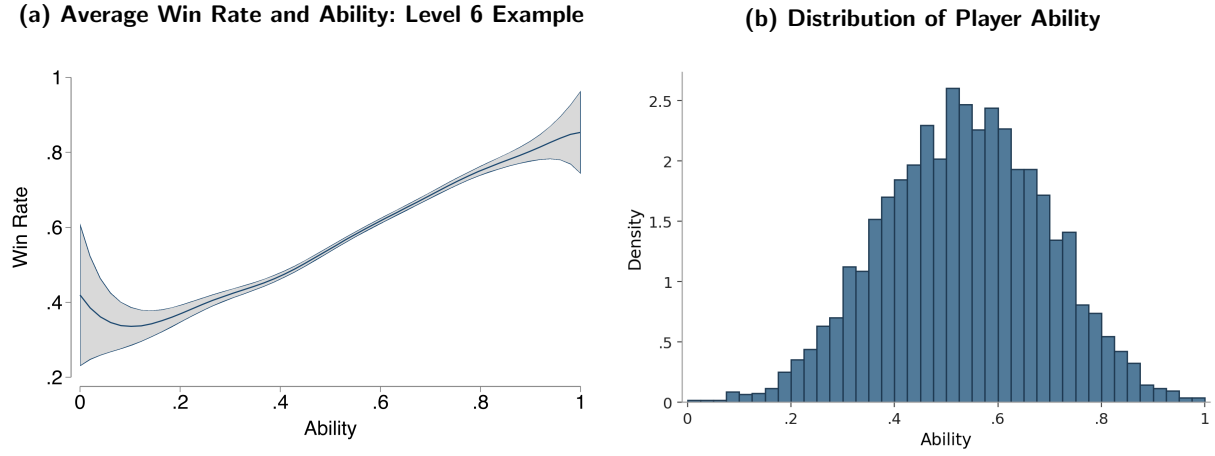
[21] The BIC values for the two- and three-segment models are 77,759.26 and 77,963.86, respectively.

Table 4 Win Probability Function Estimates

	Win
ability	1.68160***
	(0.01771)
tool stock	0.04798***
	(0.00370)
lvl:7	0.13421***
	(0.00745)
lvl:8	0.14843***
	(0.00763)
lvl:9	-0.19796***
	(0.00839)
lvl:10	-0.08103***
	(0.01067)
lvl:11	-0.16044***
	(0.01284)
lvl:7 × tool stock	-0.00185
	(0.00517)
lvl:8 × tool stock	-0.01127**
	(0.00472)
lvl:9 × tool stock	-0.01819***
	(0.00439)
lvl:10 × tool stock	-0.03392***
	(0.00437)
lvl:11 × tool stock	-0.02948***
	(0.00397)
opponent elo score	-0.88782***
	(0.01371)
Constant	-0.55691***
	(0.01017)
Observations	725,351
Log-likelihood	-484,589.1

Note: Logit model – level 6 used as baseline.
Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

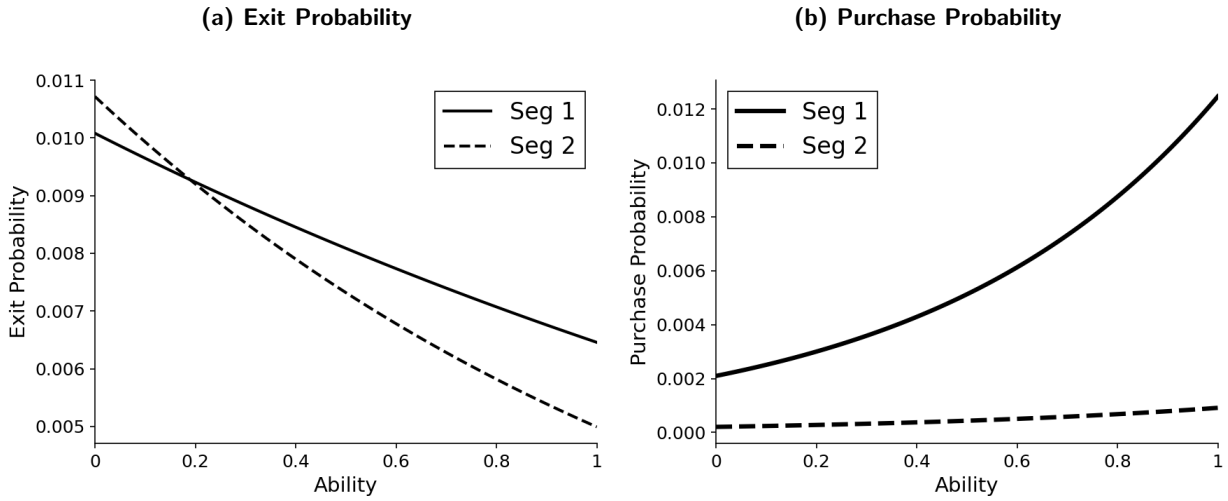
Figure 9 Player Ability Estimates



Note: The solid line in Figure 9a is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

1 comprises 17.6% of players, and Segment 2 comprises 82.4%. Figure 10 illustrates how purchase and exit choice probabilities change with player ability. Consistent with our model-free evidence on player win rates, higher ability players are more likely to purchase and less likely to exit in both segments. Segment 1 has higher purchase choice probabilities and a steeper increase with ability, whereas Segment 2 has a steeper decrease in exit choice probabilities with ability.

Figure 10 Player Action Policy by Ability



Note: Policy visualization based on example state: (Level 6, Level Completion 20-30%, and Item Stock 0)

To gain deeper insights on the segment characteristics, we report descriptive statistics in Table 6. The descriptive evidence further provides support that the smaller segment, Segment 1, comprises high-spending individuals who also stay longer in the game and progress to higher levels. Segment 2 has lower average player ability and spends less. We provide additional descriptive statistics on player segments in Online Appendix OB.11.

6.3. Structural Parameter Estimates

Table 7 reports the estimates of the structural parameters of the model, which further reveal important differences in the game play preferences among the two segments.

Segment 1 is the smaller of the two segments (17.6% of players) and has lower price sensitivity for tools (c_m).[22] These players receive greater utility from winning the current game (θ), while

[22] Since price is invariant, we normalized the price to 1; so c_m is the disutility for paying that unit price.

Table 5 First Stage Estimates: Player Action Policy Function

		Seg 1	Seg 2
Segment Probability		0.1759	0.8241
SE		0.0321	0.0321
Choice			
Purchase	ability	1.8722***	1.6080***
		(0.6249)	(0.5476)
	tool stock	0.2850***	1.9630***
		(0.0357)	(0.2039)
	tool stockˆ2	-0.0072***	-0.2169***
		(0.0015)	(0.0392)
	pct lvl complete	1.1275	1.4005*
		(0.7499)	(0.8261)
	ability × pct lvl complete	-0.3337	-0.5190
		(1.0629)	(1.1297)
	pct lvl completeˆ2	-1.5392***	-1.6964**
		(0.4830)	(0.7139)
	lvl:7	-0.1040	-0.3107**
		(0.1179)	(0.1311)
	lvl:8	-0.2114	-0.4936***
		(0.1438)	(0.1553)
	lvl:9	-0.5264***	-0.6988***
		(0.1344)	(0.1395)
	lvl:10	-0.9376***	-0.7159***
		(0.1809)	(0.1906)
	lvl:11	-1.4036***	-0.5340*
		(0.1790)	(0.2765)
	cons	-6.3406***	-8.7143***
		(0.4108)	(0.3274)
Exit	ability	-0.6687	-1.1339***
		(0.5992)	(0.1898)
	tool stock	-0.1364***	-0.0530
		(0.0355)	(0.0495)
	tool stockˆ2	0.0061***	0.0254***
		(0.0014)	(0.0092)
	pct lvl complete	-1.9978**	-1.6965***
		(0.9606)	(0.3046)
	ability × pct lvl complete	0.9212	1.4600***
		(1.7892)	(0.5285)
	pct lvl completeˆ2	0.8928	-0.0410
		(0.9965)	(0.3063)
	lvl:7	-0.6686***	-0.2947***
		(0.2297)	(0.0545)
	lvl:8	-0.5583***	-0.2860***
		(0.1796)	(0.0507)
	lvl:9	-0.3957**	-0.1285**
		(0.1708)	(0.0529)
	lvl:10	-0.4433**	-0.2641***
		(0.1927)	(0.0788)
	lvl:11	-0.6397***	-0.6026***
		(0.2358)	(0.1177)
	cons	-4.1413***	-4.0983***
		(0.2702)	(0.0861)

Note: Bootstrap standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table 6 Descriptive Characteristics of Segment

Seg 1 (17.6%); N=285	Mean	Median	75Q	90Q	99Q
Avg Total No. Purchases	4.9368	3	7	12	25
Avg Level Reached	9.1123	9	10	11	11
Avg Total No. of Games Played	408.59	312	540	797	1847
Ability Score	0.5814	0.5858	0.6743	0.7628	0.8966
In-app Purchase Player Share			98.95%		
Seg 2 (82.4%); N=3,829	Mean	Median	75Q	90Q	99Q
Avg Total No. Purchases	0.2332	0	0	1	5
Avg Level Reached	7.3003	7	8	9	11
Avg Total No. of Games Played	164.85	110	224	385	879
Ability Score	0.5308	0.5319	0.6439	0.7354	0.8768
In-app Purchase Player Share			10.08%		

Note: The descriptive statistics are based on deterministic segment assignment (i.e., the maximum of probabilistic segment probabilities estimated in the first stage).

their utility from level-up (R) is lower than that of Segment 2. We henceforth label this group as *win-seekers*.

Table 7 Structural Parameter Estimates

	Segment 1 *Win-Seekers* (17.6%)	Segment 2 *Progress-Seekers* (82.4%)
c_p	-0.1772	-0.1633
	(0.0642)	(0.0203)
c_m	-4.7218	-6.5900
	(0.1259)	(0.1039)
θ	0.3194	0.2519
	(0.1090)	(0.0367)
R	0.0017	0.0041
	(0.0010)	(0.0007)

Note: standard errors are shown in parentheses.

Segment 2 constitutes the largest group with 82.4% of the players, with higher price sensitivity for tools (c_m) than *win-seekers*. These players receive higher utility from level-up rewards (R) but lower utility from winning the current game (θ) than *win-seekers*. We label them as *progress-seekers*.

To assess model fit, we generate a representative sample of 50,000 individuals from the segment, ability, and initial tool stock state distribution, simulating a play trajectory for each player (about 190 games on average). We report the model fit in Table 8, averaged across five seeds. We show that the model performs reasonably well in matching total play count and total purchases, both in aggregate and by segment.

Table 8 Model Fit

	Real Data			Model Simulation		
	All players	Seg 1	Seg 2	All players	Seg 1	Seg 2
Average Lifetime Play Count	181.73	225.21	172.45	190.66	233.12	181.64
Average Level Reached	7.43	7.70	7.37	7.86	8.28	7.77
Average Total Purchases	0.56	2.04	0.24	0.58	2.11	0.26
Average Purchase Rate (per 100 plays)	0.27	0.72	0.18	0.27	0.87	0.14

Note: Segment averages in the real data are weighted by players' estimated segment probabilities.

7. Counterfactuals: Dynamic Complementarity and Substitution

We use the estimated model to run counterfactual simulations to understand the dynamic interlinkages between play and purchase decisions. We examine how game design interventions, such as free tool giving and difficulty adjustment, change retention and monetization, and why these effects differ by player ability and segment. For each counterfactual, we generate a representative sample of 50,000 players using the empirical distributions of segment, ability, and initial tool stock observed in the data,[23] and average results across five simulation runs with different random seeds.

7.1. Free Tool Giving: Understanding Dynamic Complementarity and Substitution

In our first counterfactual simulation, we examine a free tool-giving policy in which the firm offers one free tool to each player to understand dynamic complementarity and substitution between current tools and future purchases. Free tool giving can be seen as an extreme case of providing discounts on tools. On one hand, a free tool may act as a substitute for future purchases: by increasing win probability, the tool can reduce the need to buy additional tools. On the other hand, the same intervention may produce dynamic complementarity: the tool can help players continue playing and reach more difficult levels, where they may make purchases that would not occur under no giving. Hence, whether the net effect is substitution or complementarity is an empirical question.[24]

[23] We also examine zero initial tool stock as a mechanism illustration and robustness check (see Online Appendix OC.1.1).

[24] We acknowledge that free tool giving can change the composition and tool stock of human opponents. Our baseline counterfactuals hold the opponent environment fixed. Our findings remain qualitatively robust when the composition and tool stock of human opponents change under free tool giving (see Online Appendix OC.1.3).

7.1.1. Aggregate Effects of Free Tool Giving. We begin by assessing whether substitution or complementarity dominates in aggregate. Table 9 compares giving each player one free tool at the beginning of one of levels 6–11, with no free tool giving as the baseline. The optimal uniform policy maximizes total purchases across all players. We find that giving at the beginning of level 6 generates the largest increase in total purchases. Giving at early levels (6–9) increases average progression, games played, and total purchases, whereas giving at later levels (10–11) reduces both gameplay and purchases.

Table 9 Aggregate Effects of Uniform Free Tool Giving by Level

	% Change (Baseline: No Giving)		
Giving level	Avg. level progression	Total purchases	Avg. games played
6	1.29	1.69	1.95
7	0.64	0.69	0.94
8	0.33	0.27	0.48
9	0.17	0.13	0.28
10	0.02	-0.08	-0.01
11	–	-0.06	-0.04

Note: Five-seed averages relative to no giving, with empirical initial tool stock. Level progression is the highest level reached.

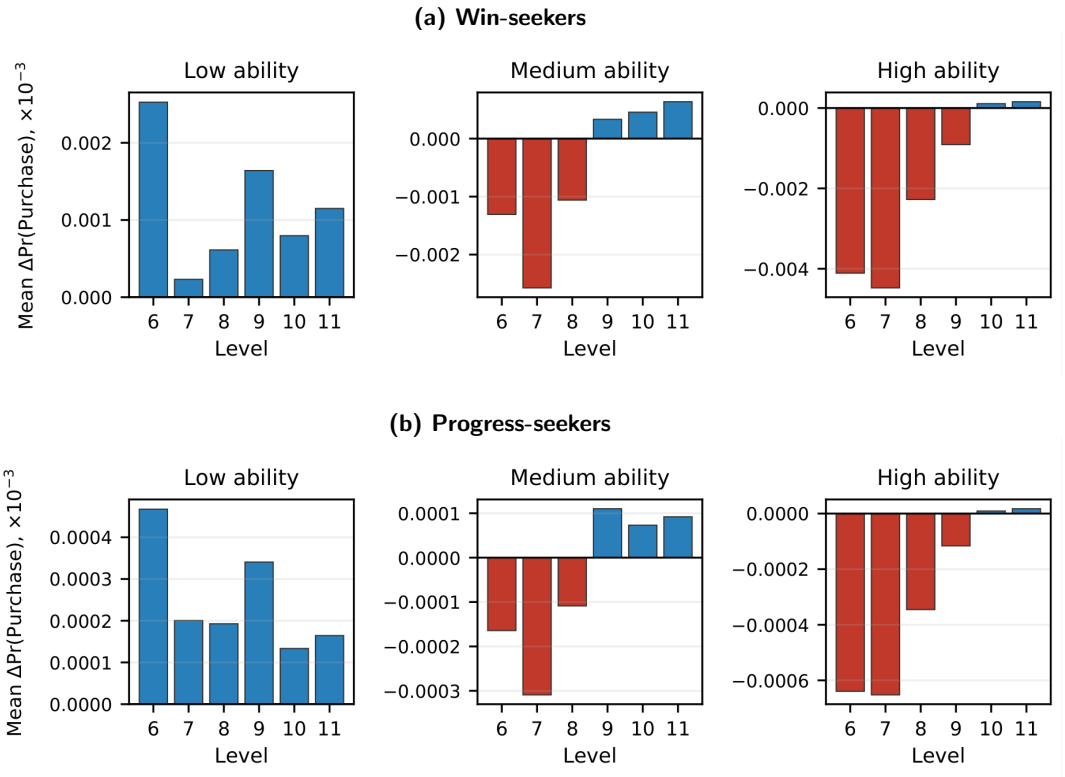
7.1.2. When Tools Complement and Substitute. Next, we examine the mechanism behind these aggregate timing effects and try to understand when and for whom a free tool complements or substitutes for purchases at a given state. When the marginal effect of tool stock on winning is increasing, an additional tool can raise the value of acquiring another tool, generating within-state complementarity. At the same time, a free tool may substitute for a purchase at the moment it is given while increasing subsequent purchases by raising the value of continued play and helping players reach more difficult levels. Online Appendix OA presents a stylized model that illustrates these mechanisms of dynamic complementarity and substitution in greater detail.

Figure 11 shows the change in the conditional choice probability (CCP) of purchase with one additional tool, by level, segment, and ability group.[25] Low-ability players show net complementarity at both early and late levels. For medium-ability players, an additional tool is a net substitute

[25] Within each segment, we group abilities using average CCP changes over early levels (6–8) and late levels (9–11). Low-ability groups have positive average changes in purchase CCPs at both early and late levels; medium-ability groups have negative average changes at early levels and positive average changes at late levels; high-ability groups have negative average changes at both.

for purchases at early levels and a net complement at late levels. High-ability players show net substitution in both ranges. However, an additional tool can increase purchase probabilities toward the final levels, where the game becomes more difficult even for high-ability players. Thus, ability alone does not determine whether a tool complements or substitutes for purchases; the player's progression also matters.

Figure 11 Changes in Purchase Probability by Level, Segment, and Ability Group



Note: We use the average tool stock from the model simulation, with ability bins weighted by original-sample player counts. Low, medium, and high-ability groups comprise bins 1–8, 9–13, and 14–20 for win-seekers and 1–10, 11–16, and 17–20 for progress-seekers, respectively.

7.1.3. Continued Play, Progression, and Total Purchases. The purchase probabilities from CCP we use in Section 7.1.2 are conditional on a given state and therefore capture only the within-state effect of receiving an additional tool. A free tool can also change the states that players subsequently reach by encouraging continued play and progression. To capture this second channel, we use counterfactual simulations of players' play and progression trajectories and measure the resulting effect on total purchases. This allows us to distinguish purchases that are cannibalized

at states players would have reached anyway from purchases generated because the free tool helps players reach additional states.

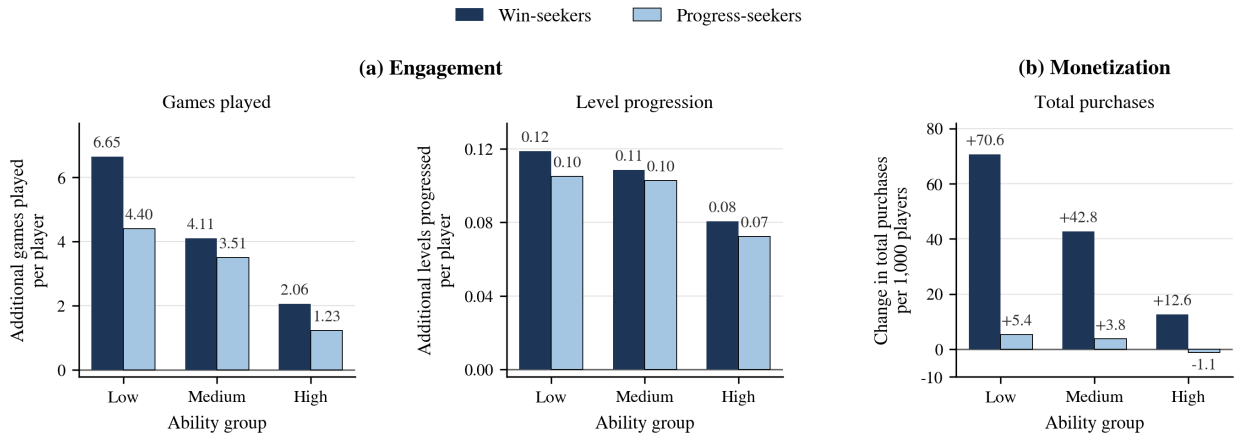
The importance of this progression channel depends on when the tool is given. When the tool is given later at higher levels rather than at early levels, fewer players remain in the game to receive it, leaving less scope for the intervention to create additional progression and future purchase opportunities. In Online Appendix OC.1, we compare early versus late tool giving and show that the cannibalization effect dominates under late giving for high-ability players and medium-ability progress-seekers.

To understand why the monetization effects of free tool giving differ across players, we use giving at level 6 as an example. Win-seekers receive higher utility from winning the current game and have lower price sensitivity, whereas progress-seekers place relatively greater value on level-up rewards and are more price sensitive. For a given increase in win probability, a free tool therefore raises the immediate value of play more for win-seekers relative to the disutility of purchasing another tool. For progress-seekers, by contrast, the free tool can help them reach level-up rewards without another purchase, while their greater price sensitivity reduces their willingness to buy additional tools.

Figure 12 shows that giving a free tool at level 6 increases games played for every segment–ability group, but the effect on total purchases is heterogeneous. Among high-ability players, win-seekers purchase more in total, whereas progress-seekers purchase less despite playing more. Low-ability players experience the largest increases in both games played and purchases. Because they reach fewer levels under the no-giving baseline, the free tool creates greater scope for additional progression and future purchase opportunities. High-ability players also play more after receiving the tool, but the increase is smaller because they already progress relatively far without the intervention.

Lastly, we decompose the change in total purchases into purchases at levels reached with and without free tool giving and purchases generated by reaching additional levels. A free tool allows players to reach later levels, where additional purchases can replace those cannibalized early in

Figure 12 Free tool giving at Level 6: Changes in Engagement and Monetization



Note: Five-seed averages relative to no giving, with empirical initial tool stock. Level progression is the highest level reached. Ability groups are those in Figure 11.

the game (see Online Appendix OC.1). Table 10 shows that free tool giving increases purchases through additional progression for every group. At the same time, it cannibalizes purchases at levels players would have reached even without the free tool. For low- and medium-ability players in both segments, as well as high-ability win-seekers, purchases generated at additional levels exceed the cannibalized purchases, so total purchases increase. For high-ability progress-seekers, the opposite holds, and total purchases decrease.

Table 10 Purchase Changes by Levels Reached with Free Tool Giving at Level 6

Segment	Ability	Net purchases from differences in levels reached (%)	Purchase change at levels reached with and without giving (%)	Total purchase change (%)
All players		+3.17	−1.48	+1.69
Win-seekers	Low	+5.14	−0.88	+4.26
	Medium	+3.06	−1.01	+2.05
	High	+1.92	−1.37	+0.55
Progress-seekers	Low	+4.49	−1.98	+2.51
	Medium	+3.52	−2.12	+1.40
	High	+2.16	−2.54	−0.38

Note: Five-seed averages relative to no giving, with empirical initial tool stock. Ability groups are those in Figure 11.

These results connect our mechanism to prior research in CRM on intervention timing and heterogeneous customer responses (e.g., Ascarza et al. 2018). The stylized model in Online Appendix OA and the conditional purchase probabilities in Figure 11 characterize substitution and complementarity at a given state. The counterfactual simulations add the dynamic effect operating

through retention and progression. For example, high-ability win-seekers exhibit net substitution in their conditional purchase probabilities, yet giving at level 6 increases their total purchases because the intervention keeps more players active and creates purchase opportunities at levels they would not otherwise reach. For high-ability progress-seekers, these additional purchase opportunities are not large enough to offset cannibalization, so they purchase less despite playing more. Thus, greater retention need not lead to greater monetization, as the net effect of free tool giving depends on whether dynamic complementarity or substitution dominates.[26]

7.2. Free Tool Giving with Difficulty Adjustment

The previous subsection showed that free tool giving can increase play while either complementing or substituting for future purchases, depending on the player. We next consider difficulty adjustment as an alternative policy lever that directly changes players' win probabilities without changing their tool stock.[27] Easing the game may encourage continued play but reduce the incentive to purchase tools, whereas making it more difficult may increase purchase incentives but also increase the risk of exit. We examine these trade-offs across player ability and preference segments by combining free tool giving with difficulty adjustment. To illustrate, we hold free tool giving fixed at entry to level 6 and adjust difficulty separately across early levels (6–8) and late levels (9–11).[28] For each segment–ability group, we search over difficulty adjustments that change win probabilities by up to five percentage points in either direction in each level range.[29]

Table 11 compares engagement and monetization under the purchase-maximizing difficulty adjustment combined with free tool giving against free tool giving alone by segment and ability group. The joint policy makes the early levels less difficult for every segment–ability group and the

[26] We report the purchase-maximizing free tool giving policy for each segment and ability group and counterfactual results for varying the timing of free tool giving within a level in Online Appendices OC.1 and OC.1.2.

[27] The firm can personalize difficulty by assigning bots of different abilities to adjust each player's average win probability at each level.

[28] We report the results for difficulty adjustment without free tool giving and for policies that also vary the timing of free tool giving in Online Appendices OC.2 and OC.2.1.

[29] With giving fixed at level 6, we choose the difficulty adjustments that maximize average total purchases across five seeds for each group. We adjust players' win probabilities at early levels (6–8) and late levels (9–11) within a range of ± 5 percentage points, including no change, in one-percentage-point increments (121 combinations).

Table 11 Free Tool Giving at Level 6 with and without Difficulty Adjustment

			Engagement				Monetization	
			Players reaching level 11 (%)		Avg. games played (%)		Total purchases (%)	
Segment	Ability	Difficulty adjustment early, late (p.p.)	Free tool	Joint	Free tool	Joint	Free tool	Joint
All players			+10.46	+98.07	+1.95	+10.58	+1.69	+9.32
Win-seekers	Low	(+5, +5)	+28.32	+797.75	+3.68	+18.76	+4.26	+19.55
	Medium	(+5, +4)	+12.41	+114.92	+1.78	+10.10	+2.05	+10.53
	High	(+5, 0)	+6.58	+17.73	+0.81	+2.92	+0.55	+2.16
Progress-seekers	Low	(+5, +5)	+31.53	+616.01	+2.88	+16.74	+2.51	+16.14
	Medium	(+5, +5)	+11.89	+128.26	+1.83	+9.53	+1.40	+9.15
	High	(+5, −3)	+5.55	+0.11	+0.57	+4.06	-0.38	+3.46

Note: Five-seed average percentage changes relative to no giving and no difficulty adjustment, with empirical initial tool stock. The joint policy includes the early/late win-probability adjustments shown, in percentage points (p.p.).

later levels less difficult for low- and medium-ability players. For high-ability progress-seekers, however, the optimal policy makes the later levels more difficult, increasing the value of performance improvements from tools. Compared with making only the early levels less difficult, making the later levels more difficult increases purchases among players who reach those levels, although fewer players ultimately reach level 11. With this adjustment, free tool giving increases total purchases relative to no giving, reversing its negative effect on purchases under the current game difficulty design. Firms can therefore use levers such as free tool giving and difficulty adjustment to manage retention and monetization over time, making current tools complement future purchases through continued play and progression.

8. Conclusion

This paper develops a dynamic structural model of consumer behavior in gaming environments, focusing on the joint and interdependent decisions to continue playing and purchase durable in-game tools. By modeling these decisions together, we show when and why current tools complement or substitute for future purchases. These dynamics imply that the same free tool can produce opposing purchase effects across players and within the same player over time. Depending on the timing of giving, a free tool can cannibalize purchases or help players continue playing and reach later levels where they make additional purchases. Whether complementarity or substitution dominates depends on player ability, preferences, and game design.

Our counterfactual simulations show that higher retention does not necessarily lead to higher monetization. Free tool giving at early levels can increase retention but reduce monetization for high-ability players who place greater value on level progression. The firm can induce dynamic complementarity between the free tool and future purchases by jointly adjusting free tool giving and game difficulty, increasing both retention and monetization for these players. While we focus on tool-giving and difficulty design, the framework can be extended to other levers such as individualized targeted discounts and game progression pacing. Importantly, our model can accommodate a range of firm objectives, including maximizing monetization, retention, or a balance of both.

Our work opens several avenues for future research. We focus on durable performance-enhancing tools, leaving out consumables and non-performance purchases (e.g., cosmetic items like clothes or identity-based items like avatars) that may affect behavior through different psychological and social channels. Exploring these dimensions could deepen our understanding of monetization strategies in virtual environments.

More broadly, as VR/AR technologies and metaverse platforms expand, firms are increasingly able to dynamically tailor product experiences in real time. Our model offers a scalable and theory-driven framework for designing retention and monetization strategies in such settings. By modeling how durable tools affect continued play and purchase decisions, we show when and why current tools complement or substitute for future purchases and how firms can induce dynamic complementarity by designing the game environment.

Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Amano, T. and Simonov, A. (2024). What Makes Players Pay? An Empirical Investigation of In-Game Lotteries. *Available at SSRN 4355019*.

Appel, G., Libai, B., Muller, E., and Shachar, R. (2020). On the monetization of mobile apps. *International Journal of research in Marketing*, 37(1):93–107. Publisher: Elsevier.

Arcidiacono, P. and Miller, R. A. (2011). Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity. *Econometrica*, 79(6):1823–1867.

Ascarza, E., Neslin, S. A., Netzer, O., Anderson, Z., Fader, P. S., Gupta, S., Hardie, B. G. S., Lemmens, A., Libai, B., Neal, D., Provost, F., and Schrift, R. (2018). In Pursuit of Enhanced Customer Retention Management: Review, Key Issues, and Future Directions. *Customer Needs and Solutions*, 5(1-2):65–81.

Ascarza, E., Netzer, O., and Runge, J. (2025). Personalized game design for improved user retention and monetization in freemium games. *International Journal of Research in Marketing*. Publisher: Elsevier.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica*, 75(5):1331–1370.

Castelo-Branco, B. and Manchanda, P. (2023). Is Video Gaming Addictive?: An Empirical Analysis. *An Empirical Analysis (June 1, 2023)*.

Chen, J. (2023). Learning and skill set formation: A structural examination of version upgrades, user visibility, and AI strategies. *Production and Operations Management*, 32(12):3856–3872.

Chen, N., Elmachtoub, A. N., Hamilton, M. L., and Lei, X. (2021). Loot Box Pricing and Design. *Management Science*, 67(8):4809–4825.

Chung, D. J., Kim, B., and Park, B. G. (2021). The Comprehensive Effects of Sales Force Management: A Dynamic Structural Analysis of Selection, Compensation, and Training. *Management Science*, 67(11):7046–7074.

Chung, D. J., Steenburgh, T., and Sudhir, K. (2014). Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans. *Marketing Science*, 33(2):165–187.

Guo, H., Hao, L., Mukhopadhyay, T., and Sun, D. (2019). Selling Virtual Currency in Digital Games: Implications for Gameplay and Social Welfare. *Information Systems Research*, 30(2):430–446.

Haenlein, M., Libai, B., and Muller, E. (2023). Satiation and cross promotion: Selling and swapping users in mobile games. *International Journal of Research in Marketing*, 40(2):342–361. Publisher: Elsevier.

Harutyunyan, M. and Koca, E. (2024). Pricing and Design of Pay-to-Win Add-ons.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529. Publisher: Wiley-Blackwell.

Huang, N., Wang, L., Hong, Y., Lin, L., Guo, X., and Chen, G. (2023). When the Clock Strikes: A Multi-method Investigation of On-the-Hour Effects in Online Learning. *Information Systems Research*, page isre.2023.1234.

Huang, N., Zhang, J., Burtch, G., Li, X., and Chen, P. (2021). Combating Procrastination on Massive Online Open Courses via Optimal Calls to Action. *Information Systems Research*, 32(2):301–317.

Huang, Y., Jasin, S., and Manchanda, P. (2019). "Level Up": Leveraging Skill and Engagement to Maximize Player Game-Play in Online Video Games. *Information Systems Research*, 30(3):927–947.

Hunicke, R. (2005). The case for dynamic difficulty adjustment in games. In *Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology*, pages 429–433, Valencia Spain. ACM.

Kang, E. H., Yoganarasimhan, H., and Jain, L. (2025). An Empirical Risk Minimization Approach for Offline Inverse RL and Dynamic Discrete Choice Model. arXiv:2502.14131.

Kasahara, H. and Shimotsu, K. (2009). Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices. *Econometrica*, 77(1):135–175.

Lee, P. S., Sudhir, K., and Wang, T. (2026). Consumer Engagement with Sequential Content: A Content-Aware Dynamic Choice Model. Working paper.

Leung, A. C. M., Santhanam, R., Kwok, R. C.-W., and Yue, W. T. (2023). Could Gamification Designs Enhance Online Learning Through Personalization? Lessons from a Field Experiment. *Information Systems Research*, 34(1):27–49.

Ma, L., Huang, T.-W., Ascarza, E., and Israeli, A. (2026). Dynamic Personalization with Multiple Customer Signals: Multi-Response State Representation in Reinforcement Learning. Working paper.

Meng, Z., Hao, L., and Tan, Y. (2021). Freemium Pricing in Digital Games with Virtual Currency. *Information Systems Research*, 32(2):481–496.

Miao, J. and Jain, S. (2024). Designing Loot Boxes: Implications for Profits and Welfare. *Marketing Science*, page mksc.2023.0007.

Nevskaya, Y. and Albuquerque, P. (2019). How Should Firms Manage Excessive Product Use? A Continuous-Time Demand Model to Test Reward Schedules, Notifications, and Time Limits. *Journal of Marketing Research*, 56(3):379–400.

Orhun, A. Y., Guo, T., and Hagemann, A. (2022). Reaching for Gold: Frequent-Flyer Status Incentives and Moral Hazard. *Marketing Science*, 41(3):548–574.

Pesendorfer, M. and Schmidt-Dengler, P. (2008). Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies*, 75(3):901–928. Publisher: Wiley-Blackwell.

Runge, J., Levav, J., and Nair, H. S. (2022). Price promotions and "freemium" app monetization. *Quantitative Marketing and Economics*, 20(2):101–139.

Santhanam, R., Liu, D., and Shen, W.-C. M. (2016). Research Note—Gamification of Technology-Mediated Training: Not All Competitions Are the Same. *Information Systems Research*, 27(2):453–465.

Sheng, L., Zhao, X., and Thomas Ryan, C. (2024). Selling Bonus Actions in Video Games. *Management Science*, page mnsc.2022.02348.

Steenburgh, T. J. (2008). Effort or timing: The effect of lump-sum bonuses. *Quantitative Marketing and Economics*, 6(3):235–256.

Wang, L., Lowry, P. B., Luo, X. R., and Li, H. (2023). Moving Consumers from Free to Fee in Platform-Based Markets: An Empirical Study of Multiplayer Online Battle Arena Games. *Information Systems Research*, 34(1):275–296.

Xue, S., Wu, M., Kolen, J., Aghdaie, N., and Zaman, K. A. (2017). Dynamic Difficulty Adjustment for Maximized Engagement in Digital Games. In *Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion*, pages 465–471, Perth, Australia. ACM Press.

Zhao, Y., Yang, S., Shum, M., and Dutta, S. (2022). A Dynamic Model of Player Level-Progression Decisions in Online Gaming. *Management Science*, 68(11):8062–8082.

Zohaib, M. (2018). Dynamic Difficulty Adjustment (DDA) in Computer Games: A Review. *Advances in Human-Computer Interaction*, 2018:1–12.