Time–Money Choices in Virtual Environments: A Structural Approach to Player Monetization

Seung Yoon Lee, K. Sudhir, Kosuke Uetake*
Yale School of Management
seungyoon.lee@yale.edu, k.sudhir@yale.edu, kosuke.uetake@yale.edu

April 2025

Monetizing virtual environments (e.g., games, metaverse platforms) requires understanding how players make interlinked decisions about spending time and money. We develop a dynamic structural model of player behavior that captures these time–money interactions: players choose when and whether to play and/or purchase in-game tools, with purchases dynamically altering the utility of gameplay in the present and future. Our framework generalizes durable goods purchase models (focused on purchases) and incentive compensation models (focused on effort/time), and is applicable to gamified virtual contexts such as online learning, health, and loyalty programs. The model clarifies how purchases interact with the environment and produce dynamic complementarities or substitutions in future play and purchase decisions. Using data from a single-player golf game, we uncover three latent segments: *premium enthusiasts*, who enjoy gameplay and are most willing to purchase tools; *win-seekers* and *progress-seekers*, who both find gameplay costly and are more price sensitive—the former primarily values immediate rewards, while the latter values level-up rewards. Counterfactual simulations of monetization strategies—such as whom to give free tools and when, or dynamic difficulty adjustments—reveal opposing effects across segments and game states. Our findings offer actionable guidance on designing personalized in-game environments that align player incentives and maximize monetization.

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 gain in importance both in terms of share of consumer time and money. As of March 2024, an estimated 59% of consumers spend more than an hour a day on video games, with 32% claiming to spend more than 5 hours a day.[1] Worldwide revenues for the video games market in 2023 is estimated at $250 billion and projected to grow over 8% a year. Revenue per user for mobile games is estimated at $53 in 2023.[2] Advancements in VR/AR technologies and the growth of metaverse platforms are expected to further accelerate consumer spending and time spent in virtual environments.

Virtual environments possess ideal properties for studying how consumers make joint decisions involving time and money in response to dynamic features of the environment—such as play 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 real-time observability of both behavior and contextual rewards makes virtual environments uniquely well-suited for empirically studying how consumer utility endogenously arise as a function of time and money inputs. Further, purchases are not merely goods that provide immediate utility; when players acquire durable tools—items that remain available for repeated use across levels and enhance gameplay performance—they increase their ability and raise their chances of success in both current and future gameplay. In this way, consumer decisions to spend time and money are dynamically interlinked: money (through tool purchases) alters the utility derived from gameplay time, which in turn influences subsequent play and purchase behavior.

[1] https://www.statista.com/forecasts/997154/hours-spent-on-playing-video-games-per-week-in-the-us

[2] https://www.statista.com/outlook/dmo/digital-media/video-games/worldwide

This observability of the choices and the "rules of the game" in virtual environments allows us to raise novel questions around managing the dynamic interaction between player decisions to play and purchase in virtual environments. For 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 later on their own. Furthermore, because tools increase players' win probability, they make progression easier and may inadvertently expedite player exit. On the other hand, since tools enable players to progress to higher difficulty levels, 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 the impact of tools intervention, the same framework can be applied to other aspects of game design (e.g., dynamic difficulty adjustment, level progression speed) to gain insights on the dynamic interlinkage of game design interventions.

To answer those questions above, we develop a dynamic structural model that captures the active process of consumer utility generation in virtual environments, where players dynamically and heterogeneously optimize their choices of time (play or quit decisions) and monetary inputs (purchasing tools) in response to the game environment. The model captures key aspects that are common in such environments: (i) the level progression design that increases in difficulty, resulting in player attrition and retaining only the high-performing players, (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 generalizes existing frameworks that have typically focused on either monetary purchases, as in dynamic durable goods models, or time and effort response, as in incentive compensation models. This makes the model applicable not only to gaming but also to other virtual environments such as digital learning platforms and digital health applications that require active user participation in effort and monetary investment that reduces the cost of effort.

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 allows us to examine the relationship between the timing of tool purchases and their impact on player performance. In most settings, individual ability is considered an unobserved variable. Our game environment context and detailed data enable us to account for wide varying heterogeneity in player ability by estimating it directly from lifetime gameplay records.

Our estimation strategy extends and adapts the two-step estimation framework in Hotz and Miller (1993). First, we estimate the player win probability function to obtain player ability estimates and incorporate player ability heterogeneity in 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 estimating the value function for each ability and latent segment type. Estimates reveal three latent segments of players: (i) *premium enthusiasts*, the smallest share of players who are the least price sensitive and do not find playing the game costly but rather enjoy spending time in the game, (ii) *win-seekers* who have the second lowest price sensitivity but find playing the game itself costly, and primarily values immediate wins more than long-term level-up rewards, and (iii) *progress-seekers*, the largest share of players who have the highest price sensitivity and find playing the game itself costly to play, but receive higher utility from level-up rewards than *win-seekers*.

We use the estimated model to run counterfactual simulations where we examine the dynamic impact of game design interventions and generate personalized policies that improve player retention and monetization. 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. We show that giving tools for free can, counterintuitively, increase both firm profit and player retention. That is, in the aggregate, the effect of dynamic complementarity dominates at the optimal level of tool-giving.

At the individual level, however, we uncover substantial heterogeneity in responses to interventions: the same intervention can act as a complement for some players and a substitute for others, depending on their ability, price sensitivity, and preferences, as well as the evolving game state. We show that even within the same player, the timing of the intervention (e.g., early vs. late levels) can produce opposing effects, underscoring the value of a structural model to navigate the complex, multidimensional policy space – not only in terms of who to target, but also when. The findings highlight the potential value of real-time personalization that maximizes monetization through the design of dynamic interventions. 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 studies 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 investigate 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 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 time (play) and money (purchase) inputs in gaming environments. Unlike traditional models that separately examine either durable goods purchases or effort/time responses, our framework captures interactions between both, enabling a comprehensive understanding of consumer behavior in virtual environments. This makes our model broadly applicable to settings that involve both time and money inputs, such as gaming, digital learning, and digital health habits. We illustrate our conceptual framework in Figure 1. Purchases improve players' ability to win, changing the value of time spent playing. This creates dynamic interlinkages between play and purchase decisions, where purchases change the utility of time spent, which in turn impacts players' future play and purchase behavior.

Thus, our framework extends and generalizes previous literature that have examined effort/time 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 empirical setting is its ability to enhance observability of both effort/time inputs and purchases

[3] There is also a stream of analytical papers studying 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)

[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

Figure 1 The Play and Purchase Loop in the Gaming Environment

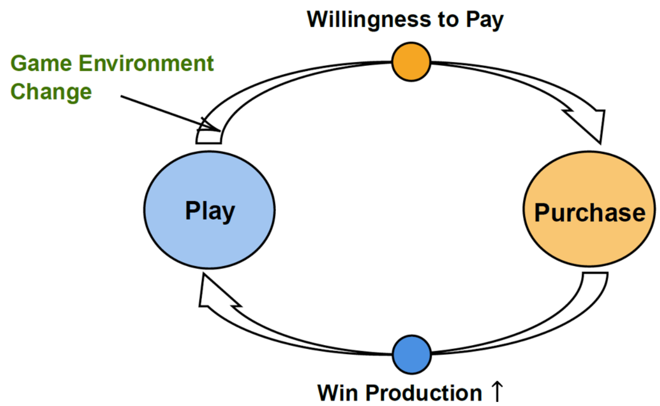

in real-time. By integrating consumer decisions of both time and money, our paper presents a comprehensive framework for modeling novel virtual environments.

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

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).

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.

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 illustrates 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 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] 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 a dynamic and forward-looking behavior.

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 focuses on a random sample of 4163 players spanning a 15-month period, from October 2021 to January 2023. We construct the sample as follows. First, we take a random 10%

[5] Although certain product offerings feature randomized rewards (e.g., receiving random items from a predefined set), the randomness is limited and does not resemble gambling-like loot box products (e.g., Amano and Simonov 2024). Players are guaranteed to receive a high-quality club upon purchase.

[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.

[7] In the game, approximately 40% of gameplay matches involve bots, but players cannot distinguish whether their opponents are real players or computer-simulated. The competitor's ability increases by each level due to selection, and players are randomly matched with opponents of average ability corresponding to their level.

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

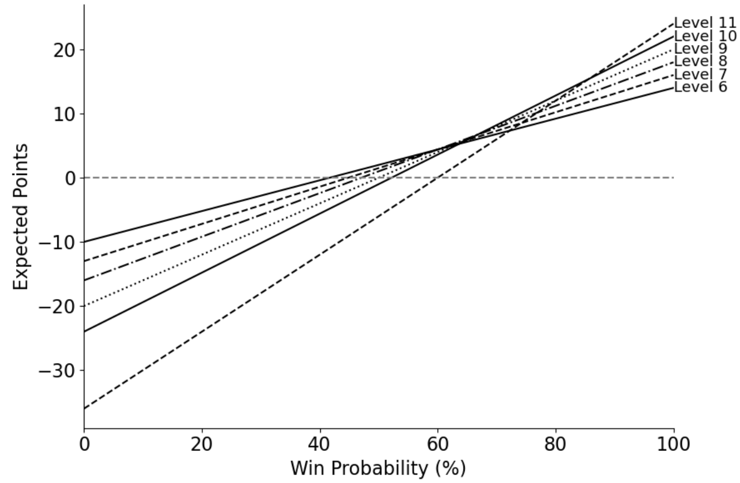

Note: Expected points are calculated as $\mathrm{E}[\text{Points}] = p \cdot \psi_v^\ell + (1-p) \cdot \psi_d^\ell$, where p is the win probability, ψ_v^ℓ is the number of points awarded for a win, and ψ_d^ℓ is the number of points deducted for a loss for each level ℓ.

Figure 3 Game Points and Level-Progression Design

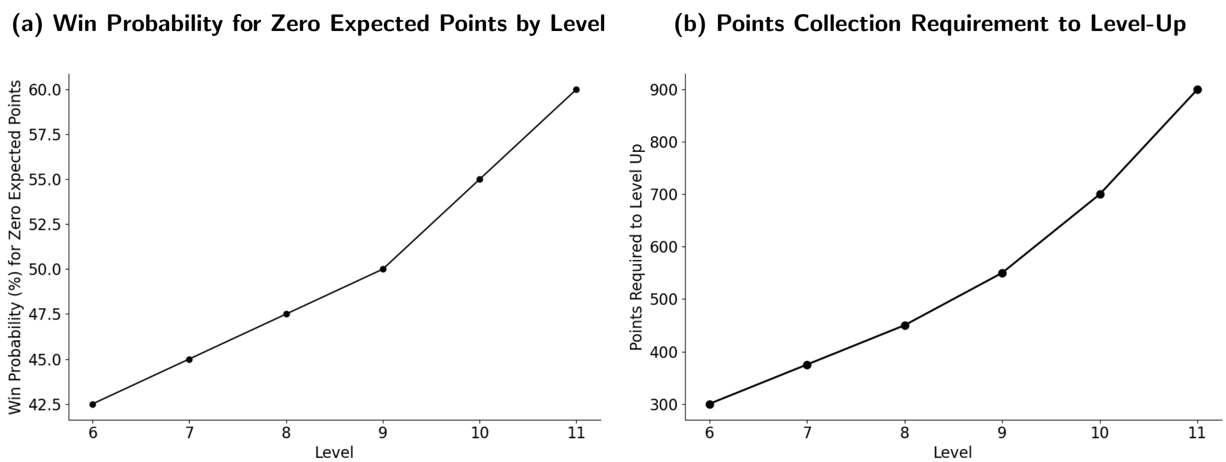

sample of players who meet the criteria of having valid play records in the data period. Second, players who have completed at least level 5 (i.e., collected at least 15% of total points available) are retained. This criterion ensures that our analysis focuses on individuals who exhibit a sufficient level of engagement within the game environment, while excluding less committed players who frequently download the game but discontinue usage shortly thereafter. Finally, we exclude outliers with unusually long within-level gameplay – defined as those exceeding 1.5 times the interquartile range 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 274 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 in the 99th percentile and above.

Table 2 Descriptive Statistics of Players

	Mean	SD	Median	75th	90th	99th	Max
Time							
Total Number of Games Played	331.57	225.18	274	406	583	1203	3211
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
Money							
In-app Purchase Player Share	22.36%						
Total Number of Purchase	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 player retention (time) and monetization (money). To account for the fact that players who spend more time in the game 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 total playtime alone.

[8] To define player exit beyond our data period, we apply the two-week churn condition.

Figure 4 Relationship Between Player Retention and Monetization

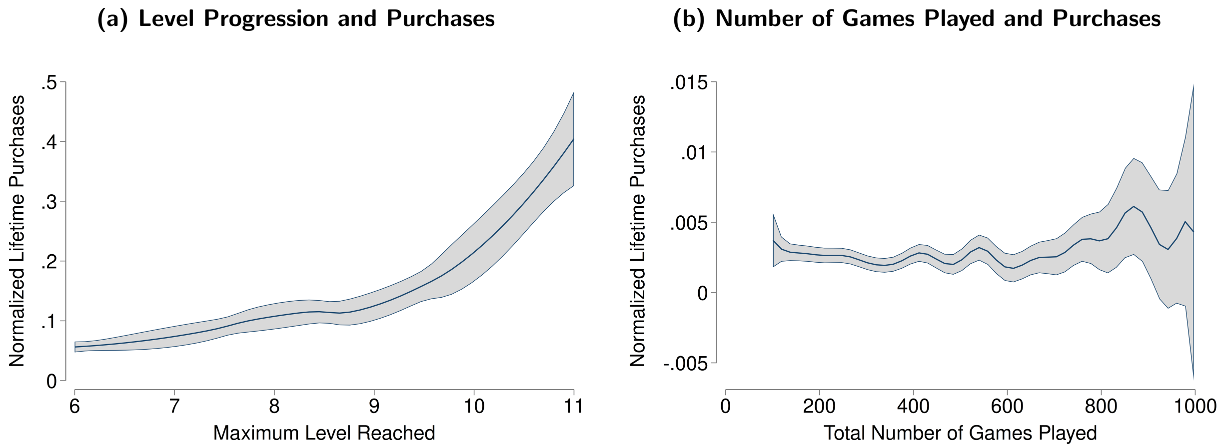

Note: The solid line is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

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 Table A.1 in Appendix A.

We next examine the relationship between players' win probability and their purchase and exit decisions, as illustrated in Figure 6. To avoid confounding the direction of the relationship (i.e., players purchasing tools which then increase their win probability), we use win rates from the previous level when plotting players' purchase and exit decisions. Within each level, we find that players with higher win probabilities are more likely to continue playing and make in-game purchases as they gain positive continuation value in the gaming experience. Conversely, players with

Figure 5 **Average Win Rate Distribution of Players By Levels**

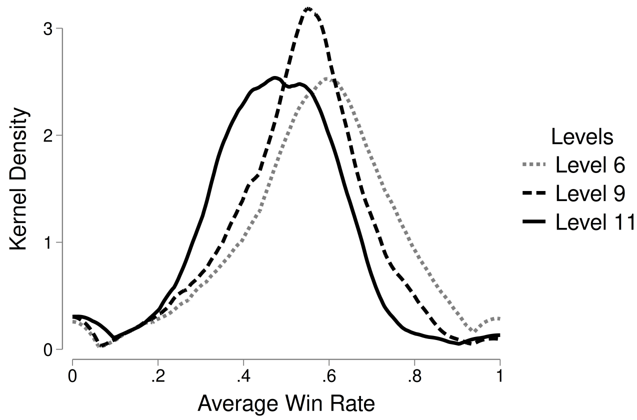

Note: The lines represent kernel density estimates using the Epanechnikov kernel.

lower win rates are more likely to exit the game and purchase less, as utility decreases with lower win rates, reducing players' perceived value of the game (see Figure A.1 in Appendix A for the full set of graphs by level).

Figure 6 **Player Exit and Purchase by Win Rate**

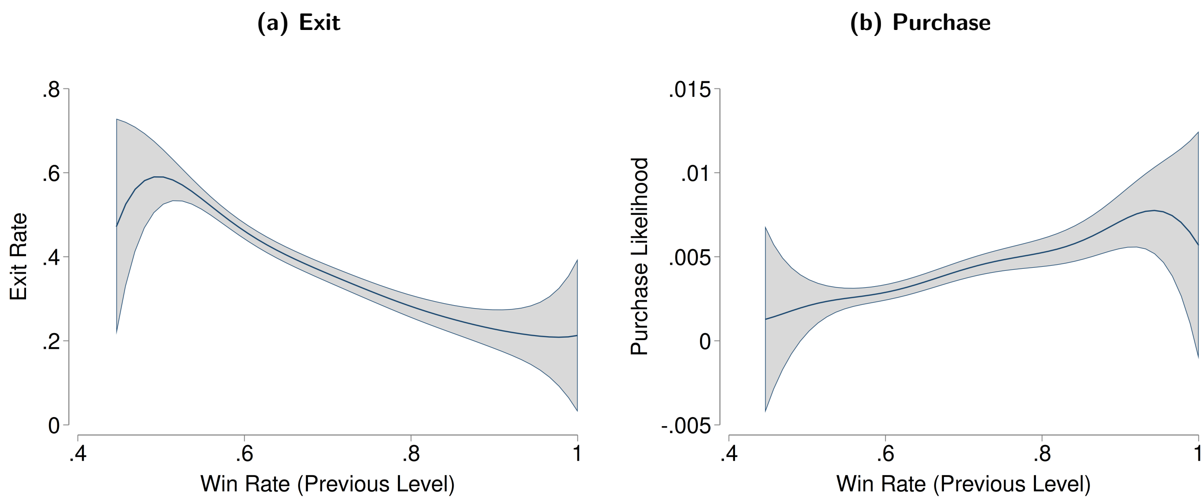

Note: The solid line is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

Finally, Table 3 presents players' purchase and exit likelihood in relation to their level completion status. Players are more likely to exit immediately upon reaching a higher level and falls as they progress through that level. Purchases are also more likely to occur earlier in the level, suggesting that players purchase the tool as soon as they face higher levels of difficulty in the new level; it also makes sense to buy early, as these tools are durable in that they help improve future gameplay.

Table 3 Relationship Between Level Completion and Exit/Purchase Probability

	Pr(Exit)	Pr(Purchase)
% Level Complete	-0.709***	-0.316***
	(0.061)	(0.072)
Constant	-3.786***	-7.140***
	(0.056)	(0.088)
Controls	Y	Y
N	747,800	747,800

***p <0.01; **p <0.05; *p <0.1; Controls: tool stock, level FE, ability FE

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 ability. 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}) has a one-to-one mapping to the level (ℓ_{it}), determined by the game's points and level progression design Ψ, and reflects the player's current progression state in the game.

Figure 7 describes the timeline of the model. Before entering a game t, player i observes the current tools and progression (i.e., accumulated points) state in the game. 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

[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). Thus, players discount future utility based on the number of games played rather than the passage of real time.

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. After the player completes all levels, the player can continue to play the unlocked stages, which is a pattern we observe for the majority of players in the game.

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.

Figure 7 Model Timeline

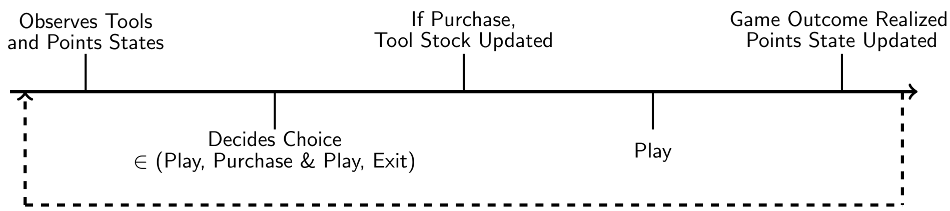

4.1. Player Utility Generation Process

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 incurring in-app purchase ($d_{it}=1$), whether to make in-app purchase and play the game ($d_{it}=2$), or whether to permanently exit the game platform ($d_{it}=3$).

$$u(S_{it}, d_{it}) = \begin{cases} c_p + \theta E[r_{it}(W_{it};\Psi)] + R \cdot \mathbf{1}_{\text{LEVEL UP}(z_{it})} \cdot L_{\ell(z_{it})} + c_m \cdot \mathbf{1}_{\{d_{it}=2\}} + \epsilon_{dit}, & \text{if } d_{it} \in \{1,2\}, \\ \epsilon_{3it} & \text{otherwise (exit).} \end{cases}$$

In each period, if the player decides to play, the player incurs a cost of play c_p that capture the net utility or disutility from engaging in gameplay, independent of the game outcome. A positive c_p indicates that the player receives intrinsic enjoyment from play, while a negative c_p suggests that the time or effort costs outweigh the enjoyment. In every play, the player also receives utility from immediate rewards, where θ represents the player's utility derived from point reward r_{it} (i.e., expected number of points gained from the current period play). 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,

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

and

$$\mathrm{E}[r_{it}] = \Pr(W_{it} = 1) \cdot \psi_v^\ell + \Pr(W_{it} = 0) \cdot \psi_d^\ell \tag{2}$$

In addition to these immediate rewards, the player receives a level-up reward R every time the player reaches a new level ℓ, where R represents the intrinsic utility from reaching the next level. To parsimoniously model the increasing utility from leveling up at higher levels, we adjust R with respect to the level achievement criterion (i.e., the number of points required to level up), $L_{\ell(z_{it})}$. Since the level-up points criterion increases monotonically with levels in the current game design, the player receives greater rewards for achieving higher levels.[10]

If the player decides to purchase tools ($d_{it} = 2$), 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 an extreme value distribution. We normalize the exit value as 0.

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

$$\Pr(W_{it} = 1) = \ell(z_{it}) + \alpha_i + \delta_1 k_{it} + \delta_2 k_{it} \cdot \ell(z_{it}) + \lambda e_{it} + \xi_{it}. \tag{3}$$

The player's win production function accommodates the following key characteristics to model the gameplay outcome. First, the win probability decreases as the level difficulty ℓ_{it} increases, which is a one-to-one mapping from the point accumulation state z_{it}. Second, the player's win

[10] 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.

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 the diminishing effect of tools at higher levels. Finally, the exogenous opponent ability e_{it} and idiosyncratic shock ξ_{it} affects the outcome of the game.[11]

There are a few points to note about the win production function. First, ℓ_{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 ℓ_{it} captures this net effect of increased difficulty and player improvement in ability. In addition, $k_{it} \cdot \ell(z_{it})$ 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.[12]

4.2. State Transitions

The state variables tool stock k_{it} and points stock z_{it} evolves deterministically as follows:

$$k_{it} = \begin{cases} \min(k_{it-1} + 1, \overline{K}) & \text{if } d_{it} = 2 \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

The player's tool stock k_{it} increments by one with each purchase decision. Players can make multiple purchases over their lifetime. $\overline{K}$ is the maximum tool stock, representing the maximum tool quality upgrade level.

$$z_{it} = \begin{cases} \min(z_{it-1} + \psi_v^{\ell}, \overline{\psi}^{L}) & \text{if } W_{it} = 1 \\ \max(z_{it-1} - \psi_d^{\ell}, \underline{\psi}^{\ell}) & \text{otherwise} \end{cases} \tag{5}$$

[11] Note that players are randomly matched with opponents (can be either bots or real players) of average ability corresponding to their level. In our model estimation, we control for the effect of the opponent in the player win production function.

[12] After accounting for ability and level effects, we find that the incremental variance explained by player learning by doing is only around 3%. This is because much of the player learning is already absorbed as a net effect of level progression. Given that our current specification captures 97% 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 Appendix B for details).

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 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 ψ_d^ℓ 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

$$V(k,z,\alpha,\epsilon;\Theta,\Psi) = \max\left\{u_d(k,z,\alpha,\epsilon;\Theta,\Psi) + \beta\,\mathbf{1}_{d\in\{1,2\}}\,\mathbb{E}\Big[V(k',z',\alpha,\epsilon';\Theta,\Psi)\,\Big|\,k,z,\alpha,d\Big]\right\}, \tag{6}$$

where the idiosyncratic shock ϵ follows a Type-I extreme value, and the discount factor β is set to 0.9.[13] The player continues to play the game if the expected continuation value is greater than 0, the normalized outside value of exit.

5. Estimation

We estimate the model using two-step estimation (Hotz and Miller 1993). In the first step, we estimate the player win probability function and the conditional choice probabilities (CCPs) of player action choice as a flexible function of state variables. 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. In the second step, we estimate the structural parameters that rationalize the first stage policy estimates.

[13] $\beta = 0.9$ gives the best fit compared to β= 0.95, 0.99, and 0.995. In our context, setting lower discount factors greatly underestimates player exit probabilities and hence overestimates player purchase probabilities. Average inter-play time gap (in day) is around 2 days, and 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, both for computational efficiency in the state space and practical reasons. The early levels 1-5 consist of beginner level tutorials and short level length, with these five levels accounting for only 15% 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 transition states to track level-up bonus (a total of 65), and tool stock state to evolve deterministically up to 25 transactions. This leaves us with a total of 1690 state combinations.

Typically, player ability is treated as an unobserved variable. In order to incorporate rich player heterogeneity in win rates as shown in the model-free evidence, we instead treat player ability as an observed variable by estimating it from player lifetime gameplay records. Because our setting is a game, where we observe every player action, environment, and outcome, we can estimate player ability α_i as an individual fixed effects parameter from Equation 3, the player win probability function, controlling for the effects of tools and game environments on player win rates. By incorporating player ability as an observed variable, we can directly include it in our first stage policy estimation. This enables us to account for rich observed heterogeneity in player ability levels in predicting player action. 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.

Given the state variables and observed ability heterogeneity of the players, 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 that player i belongs to one of G segments $g \in \{1, ..., G\}$ with segment probabilities $q_i = \{q_{i1}, ...q_{iG}\}$. Let π_g denote the population probability of being in segment g. We iteratively maximize the log likelihood in Equation 7,

$$\sum_{i=1}^{N}\sum_{g=1}^{G}\sum_{t=1}^{T} q_{ig} \ln[\mathcal{L}(d_{it}|S_{it}, \alpha_i, g, \beta_g)], \tag{7}$$

where

$$q_{ig} = \frac{\pi_g \prod_{t=1}^{T} \mathcal{L}(d_{it}|S_{it}, \alpha_i, g, \beta_g)}{\sum_{g=1}^{G} \pi_g \prod_{t=1}^{T} \mathcal{L}(d_{it}|S_{it}, \alpha_i, g, \beta_g)}, \tag{8}$$

and $\mathcal{L}(d_{it}|S_{it}, \alpha_i, g, \beta_g)$ is the choice probability of taking action $d_{it} = j$ for segment type g,

$$\mathcal{L}(d_{it}|S_{it}, \alpha_i, g, \beta_g) = \frac{e^{\beta_g^j[S_{it},\alpha_i]}}{\sum_j e^{\beta_g^j[S_{it},\alpha_i]}}. \tag{9}$$

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

(a) Compute $q_{ig}^{(m+1)}$ using Equation (2) with $\beta_g^{(m)}$ and $\pi^{(m)}$.

(b) Update population shares

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

(c) Update $\beta_g^{(m+1)}$ for each segment g by maximizing Equation (1) with $q_{ig}^{(m+1)}$ and $\pi_g^{(m+1)}$.

We iterate steps (a)-(c) until convergence. We initialize β_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 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 rich observed 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. Given our discrete state space, we can solve the value function as a system of linear equations,

$$V(S_{it}, \alpha_i, g; \Theta, \Psi) = (I - \beta F)^{-1} \left\{ \sum_{d_{it} \in \{1,2,3\}} P(S_{it}, \alpha_i, g; \Theta, \Psi) \cdot [u(S_{it}, \alpha_i, g, d_{it}; \Theta, \Psi) + E[\epsilon | d_{it}] \right\} \tag{10}$$

where F is the matrices of transition probabilities corresponding to action d_{it} (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 $E[\epsilon | d_{it}] = \gamma - \ln(P(\cdot; \Theta))$. γ is the 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:

$$\Pr(d_{it} | S'_{it}, \alpha_i, \epsilon; \Theta, \Psi) = \frac{\exp\left(u(d_{it}, S'_{it}, \alpha_i; \Theta, \Psi) + \beta E_{S''_{it}, \alpha_i, \epsilon'; \Theta, \Psi | d_{it}, \epsilon} V(S''_{it}, \alpha_i, \epsilon')\right)}{\sum_{\tilde{d}_{it} \in \{1,2,3\}} \exp\left(u(\tilde{d}_{it}, S'_{it}, \alpha_i; \Theta, \Psi) + \beta E_{S''_{it}, \alpha_i, \epsilon'; \Theta, \Psi | \tilde{d}_{it}, \epsilon} V(S''_{it}, \alpha_i, \epsilon')\right)}. \tag{11}$$

We construct the minimum distance estimator, where $\hat{Pr}$ is the optimal policy estimated from the first stage, and $\tilde{Pr}$ 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 segment probabilities q_{ig}.

$$\begin{aligned} \sum_{i=1}^{N} \sum_{t=1}^{T} q_{ig} \Big[& (\hat{Pr}(d_{it} = 2 | S_{it}, \alpha_i, g, \epsilon; \beta_g) - \tilde{Pr}(d_{it} = 2 | S_{it}, \alpha_i, g, \epsilon; \Theta_g, \Psi))^2 \\ & + (\hat{Pr}(d_{it} = 3 | S_{it}, \alpha_i, g, \epsilon; \beta_g) - \tilde{Pr}(d_{it} = 3 | S_{it}, \alpha_i, g, \epsilon; \Theta_g, \Psi))^2 \Big] \end{aligned} \tag{12}$$

For the second stage model estimation, we discretize the player ability type variable into 20 bins and estimate the value function for each ability type and segment. To compute standard errors, we generated 500 bootstrap datasets following Bajari et al. (2007). For each bootstrapped dataset, we estimate both the first and the second stage to account for the estimation errors from 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 in equation 3. The average gameplay records used for estimating the win probability function for each player is 171. While it is theoretically possible to identify tools-ability substitutability or complementarity (i.e., the interaction between tools and ability in the win probability function), doing so is practically infeasible given our individual-specific measure of ability. The trade-off we make for capturing this level of granularity in individual heterogeneity is the inability to separately identify the effects of player learning and tools-ability interactions. We follow Kasahara and Shimotsu (2009) for the identification of unobserved finite mixture heterogeneity, which show that at least three periods of panel data are necessary for the identification. Our panel data can easily satisfy the condition.

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 expected points rewards $\mathrm{E}[r_{it}]$ varies by states, while c_p is a constant. The disutility of purchase parameter c_m is identified from the play versus play and purchase decisions across states. Lastly, the intrinsic reward 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 reward 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 Production Function and Ability Estimates

We estimate the player win production function using a linear probability model of player tool and level states with individual fixed effects, controlling out the opponent effect.[14] We use the estimated individual fixed effects as the measure of player ability. 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. Specifically, one additional tool purchase at level 6 increases player win probability by around 2.6 percentage points [15] – with the marginal effect diminishing at higher levels. Third, player win probability decreases with level. Notably, the game becomes significantly more difficult at level 9, with win probability dropping by around 10 percentage points compared to the previous level. This increasing difficulty at higher levels is a common feature in gaming environments.

To demonstrate the reliability of our ability measure, we present the relationship between player average win probability and ability in Figure 8a. There is a clear positive relationship, with higher ability players exhibiting higher win rates. Figure 8b 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 estimation of player action probabilities using a flexible multinomial logistic regression, accounting for both observed and unobserved heterogeneity. We find that the three-segment model best fits the data based on the AIC and BIC criteria.[16] We report the first stage policy estimates for each segment in Table 5.

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

[15] 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, comparing the five games before and after the purchase incident in Appendix C. The measured effect size largely aligns with the incremental effect of the tool measured in our win probability function.

[16] The AIC for the two- and three-segment models are 76764.81 and 76299.59, respectively; the BIC values are 77041.41 and 76714.49, respectively.

Table 4 Production Function Estimates

	Win
tool stock	0.02623***
	(0.00157)
lvl:6	-0.08207***
	(0.00146)
lvl:7	-0.06441***
	(0.00187)
lvl:8	-0.07087***
	(0.00199)
lvl:9	-0.17077***
	(0.00228)
lvl:10	-0.17392***
	(0.00296)
lvl:11	-0.23899***
	(0.00384)
lvl:6 × tool stock	-0.00473***
	(0.00136)
lvl:7 × tool stock	-0.00699***
	(0.00141)
lvl:8 × tool stock	-0.01082***
	(0.00140)
lvl:9 × tool stock	-0.01151***
	(0.00143)
lvl:10 × tool stock	-0.01430***
	(0.00148)
lvl:11 × tool stock	-0.01505***
	(0.00147)
opponent elo score	-0.24202***
	(0.00277)
Observations	1,168,880
Individual FE	Y
Adjusted R^2	0.049

Note: Linear probability model – level 5 used as baseline.
Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Segment 1 is the smallest, comprising 7% of the players, followed by Segment 2 with 23%, and Segment 3 with 70%. We report some illustrative features of the policy with respect to player ability in Figure 9. Consistent with our model-free evidence, the probability of exit decreases with player ability, while the probability of purchase increases with ability across all segments. Segment 1 has the highest purchase rates and generally the lowest exit rates, indicating that these players are the most engaged and likely to spend money in the game. Segment 3, the largest segment, has the lowest purchase rates. While Segment 3 shows a relatively flat decrease in exit probability across abilities, Segment 3 exhibits a steeper decline among lower ability players.

To gain deeper insights on the segment characteristics, we report descriptive statistics in Table 6. The descriptive evidence further provides support that the smallest share of players, Segment 1

Table 5 First Stage Estimates: Player Action Policy Function

		Seg 1	Seg 2	Seg 3
Segment Probability		0.0692	0.2287	0.7021
SE		0.0046	0.0048	0.0046
Choice				
Purchase	ability	1.0078**	1.8555***	1.9965***
		(0.3696)	(0.3389)	(0.3636)
	tool stock	0.2119***	0.4950***	1.8689***
		(0.0185)	(0.0163)	(0.0478)
	tool stockˆ2	-0.0069***	-0.0128***	-0.1999***
		(0.0010)	(0.0007)	(0.0086)
	pct lvl complete	2.6695***	-0.9826**	1.2923***
		(0.4361)	(0.4249)	(0.4825)
	ability × pct lvl complete	-0.7362	0.1418	-0.3467
		(0.7061)	(0.6362)	(0.7756)
	pct lvl completeˆ2	-2.7620***	0.2979	-1.8515***
		(0.4350)	(0.4454)	(0.4203)
	lvl:7	-0.0348	-0.2065**	-0.2914***
		(0.1038)	(0.1016)	(0.0872)
	lvl:8	0.1223	-0.4376***	-0.5923***
		(0.1009)	(0.1042)	(0.0908)
	lvl:9	-0.3773***	-0.6725***	-0.7774***
		(0.1139)	(0.1089)	(0.0983)
	lvl:10	-0.6498***	-1.3472***	-0.8373***
		(0.1268)	(0.1435)	(0.1201)
	lvl:11	-0.9534***	-1.5476***	-1.0228***
		(0.1431)	(0.1598)	(0.1534)
	cons	-5.5698***	-6.8071***	-8.5193***
		(0.1687)	(0.1457)	(0.1699)
Exit	ability	0.7717**	-3.0300***	-3.6848***
		(0.3469)	(0.2113)	(0.1987)
	tool stock	-0.1393***	-0.2053***	-0.0152
		(0.0286)	(0.0219)	(0.0262)
	tool stockˆ2	0.0056***	0.0087***	0.0089***
		(0.0013)	(0.0009)	(0.0034)
	pct lvl complete	-0.9030*	-1.7382***	-1.3912***
		(0.5159)	(0.2657)	(0.2639)
	ability × pct lvl complete	-5.8553***	7.3085***	1.4232***
		(0.9142)	(0.5172)	(0.4862)
	pct lvl completeˆ2	2.8489***	-2.1943***	-0.0362
		(0.4691)	(0.2832)	(0.2415)
	lvl:7	-0.8087***	-0.3470***	-0.0847*
		(0.0877)	(0.0532)	(0.0500)
	lvl:8	-0.6385***	-0.3837***	0.0199
		(0.1128)	(0.0568)	(0.0520)
	lvl:9	-0.5340***	-0.1310**	0.3191***
		(0.1140)	(0.0640)	(0.0549)
	lvl:10	-0.9494***	0.0801	0.3248***
		(0.1757)	(0.1075)	(0.0785)
	lvl:11	-0.7838***	-0.6067***	0.3303***
		(0.1840)	(0.1916)	(0.1036)
	cons	-4.3423***	-3.4950***	-3.4308***
		(0.1257)	(0.0748)	(0.0701)

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

Figure 8 Player Ability Estimates

(a) Average Win Rate and Ability: Level 6 Example **(b) Distribution of Player Ability**

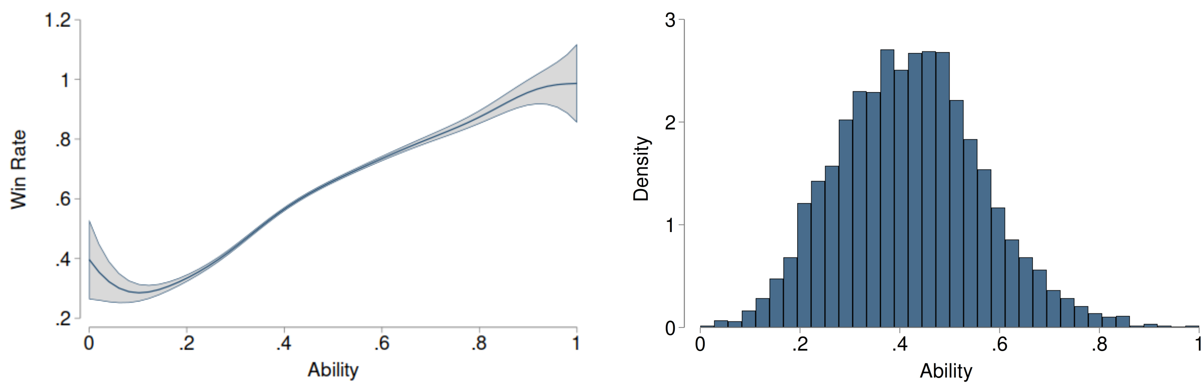

Note: The solid line in Figure 8a is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

Figure 9 Player Action Policy by Ability Type

(a) Exit Probability **(b) Purchase Probability**

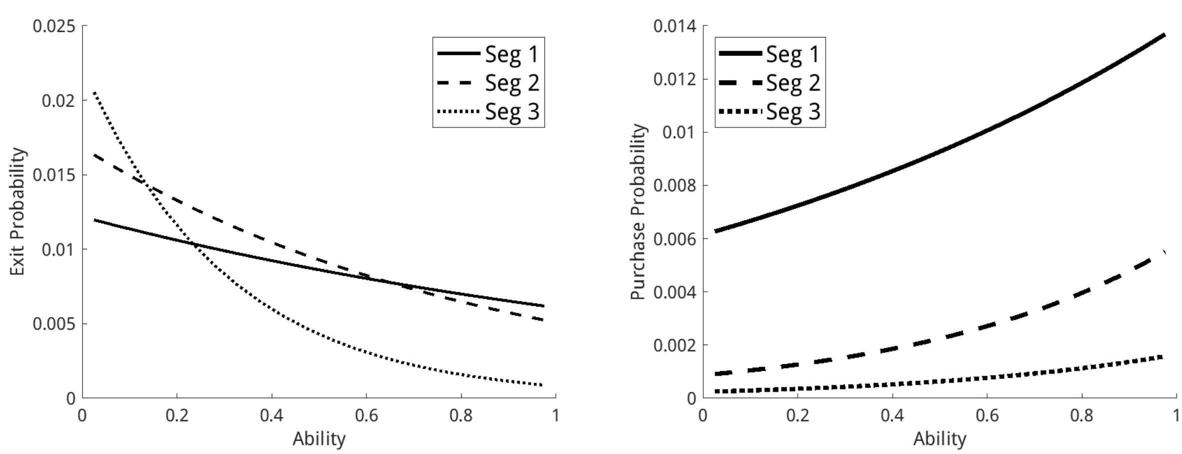

Note: Policy visualization based on example state: (Level 6, Level Completion 20-30%, and Item Stock 0)

comprises of high-spending individuals who also stays longer in the game and progress to higher levels, followed by Segment 2. Segment 3 has the lowest average player ability and spends the least. Because this segment has a greater share of low-ability players, it also has higher exit rates, especially at the early levels. We provide additional descriptive statistics on player segments in Appendix D.

Table 6 Descriptive Characteristics of Segment

Seg 1 (7.0%); N=107	Mean	Median	75Q	90Q	99Q
Avg Total No. Purchase	6.8785	5	9	14	24
Avg Level Reached	9.3364	9	11	11	11
Avg Total No. of Games Played	470.66	314	639	1021	1980
Ability Score	0.4823	0.4860	0.5912	0.6861	0.8398
Pay-to-Win Player Share			100.00%		
Seg 2 (23.0%); N=218	**Mean**	**Median**	**75Q**	**90Q**	**99Q**
Avg Total No. Purchase	2.7661	1	3	7	25
Avg Level Reached	7.9725	8	9	10	11
Avg Total No. of Games Played	263.68	164	444	673	916
Ability Score	0.5330	0.5193	0.6684	0.7812	0.9282
Pay-to-Win Player Share			68.35%		
Seg 3 (70.0%); N=3,838	**Mean**	**Median**	**75Q**	**90Q**	**99Q**
Avg Total No. Purchase	0.2507	0	0	1	5
Avg Level Reached	7.3233	7	8	9	11
Avg Total No. of Games Played	166.74	112	227	387	887
Ability Score	0.4129	0.4117	0.5056	0.5882	0.7556
Pay-to-Win Player Share			10.76%		

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

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 three segments.

Segment 1 represents the smallest share of players (7%) but those who spend the most, which is reflected in their lower price sensitivity for ability enhancers (c_m).[17] Despite their low sensitivity to immediate point rewards (θ), they derive significant utility from progressing through game levels, as indicated by their high level-up reward (R). This segment's positive (near-zero) cost of play (c_p) suggests that these players do not find playing the game costly, but rather enjoy spending time in the game. We henceforth label this group as *premium enthusiasts*.

Segment 2 players, representing the share of 23% of players, receive a greater utility from immediate points reward (θ) and have the second lowest price sensitivity for ability enhancers (c_m), below the *premium enthusiasts*. They however have the highest cost of play (c_p), indicating that they find playing the game more of a chore. This suggests that these players are less likely to continue playing without high enough points reward from the current game, and their utility from level-up (R) is the lowest among the segments. We label this group as *win-seekers*.

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

Table 7 Structural Parameter Estimates

	Segment 1 *Premium Enthusiasts* (7.0%)	Segment 2 *Win-Seekers* (23.0%)	Segment 3 *Progress-Seekers* (70.0%)
c_p	0.0002	-0.0395	-0.0288
	(0.0020)	(0.0036)	(0.0026)
c_m	-4.4089	-5.2775	-6.4997
	(0.0333)	(0.0386)	(0.0448)
θ	0.0003	0.0050	0.0050
	(0.0002)	(0.0004)	(0.0003)
R	0.0039	0.0016	0.0036
	(0.0006)	(0.0009)	(0.0014)

Note: standard errors are shown in parentheses.

Segment 3 constitutes the largest group with 70% of the players. These players have the same utility from immediate points reward (θ) as *win-seekers*, but their price sensitivity for enhancers (c_m) is the highest. They also find playing the game costly (c_p), but receive higher utility from level-up rewards (R) than *win-seekers*. We label them as *progress-seekers*.

To assess model fit, we generate a representative sample of 20,000 individuals from the segment, ability, and initial tool stock state distribution, simulating a play sequence for each player until they exit or finish the game; the sequences average about 200 games per player. We report the model fit in Table 8. In general, the model performs reasonably in matching the target moments, time (total play count) and money (total purchase), and in accounting for the heterogeneity across the three latent segments. Hence, we conclude that the model can reasonably match player action and play behavior observed in data, especially with respect to different insights and predictions across the latent segments.

Table 8 Model Fit

	Real Data			Model Simulation		
	Seg 1	Seg 2	Seg 3	Seg 1	Seg 2	Seg 3
Average Lifetime Play Count	231.35	147.07	185.14	243.15	166.73	196.51
Average Level Reached	7.69	7.15	7.47	8.92	7.79	7.97
Average Total Purchase	2.77	0.70	0.29	2.95	0.89	0.31
Average Purchase Rate (per 100 plays)	0.97	0.33	0.19	1.17	0.53	0.15

7. Counterfactuals

We use the estimated model to run a series of counterfactual simulations, where we examine the dynamic effects of game design interventions on players' purchase and play decisions. To evaluate the effectiveness of various policy interventions, we generate a representative sample of 50,000 players that mirrors the segment-ability distribution observed in the data.[18] We also simulate personalized intervention policies, identifying the optimal targets and timing that maximize the net dynamic complementarities in future purchases.

7.1. Dynamic Interlinkages in Purchases with Free Tools

In our first counterfactual simulation, we examine the implications of a free tool-giving policy, which can be seen as an extreme case of discounting. This allows us to cleanly isolate how tool acquisitions affect future play and purchase behavior. On one hand, a free tool may act as a substitute for future purchases: by increasing player ability and accelerating progression, the tool can reduce the perceived need to buy additional tools. Faster progression may also lead to earlier game exit, reducing long-term engagement and monetization. On the other hand, the same intervention may produce dynamic complementarity: enabling players to reach more challenging levels increases difficulty, prompting subsequent tool purchases to sustain win probability and progression—thereby reinforcing both play and purchase. Hence, whether the net effect is substitution or complementarity is ultimately an empirical question.

Given these opposing effects, we now address a key managerial question: If a firm were to offer one free tool to each player, when (i.e., at which level) and to whom should it be offered? We begin by assessing whether substitution or complementarity dominates in aggregate. We then examine whether the same tool-giving policy produces opposing effects (i) across players and (ii) within the same player over time. Leveraging these insights, we design personalized real-time policies aimed at maximizing the net effect of dynamic complementarities.

[18] For counterfactual simulations, initial tool stock is set to zero to isolate the effects of the interventions from pre-existing tool ownership.

Uniform Tool-Giving Policies. We first consider uniform tool-giving policies as a benchmark. Table 9 reports the results of offering a free tool at the start of each level. As expected, tool-giving increases players' win probability and enables progression to higher levels across all timing strategies. However, the impact on future purchases and play behavior varies significantly by timing. Depending on the level at which the tool is given, the intervention can, on average, act as either a dynamic substitute or complement to future purchases.

Table 9 Aggregate Effects: Uniform Free Tool Giving Policy by Level

Level	% Change (Baseline: No Tool-Giving)		
	Avg. Level Progression	Profit	Avg. Play Count
6	**1.80**	**0.64**	**1.33**
7	0.99	-0.17	0.72
8	0.53	-0.30	0.37
9	0.32	0.20	0.24
10	0.08	-0.54	-0.11
11	0.00	0.05	-0.08

Heterogeneous Effects Across Players. We next explore the heterogeneous effect across players; specifically, we ask who should receive the tool, holding the optimal tool-giving level fixed. Figure 10 decomposes the effect of the free tool across player segments and ability levels. Even when the timing is fixed,we find that providing a free tool leads to opposing effects across player types, substituting future purchases for some, while complementing them for others.

Timing Effects Within Players. Beyond heterogeneity across players, we also investigate whether the timing of tool-giving can produce opposing effects even within the same player. As an illustrative example, Figure 11 compares outcomes from giving a tool at different levels, showing that the same free tool can either substitute or complement future purchases for the same player depending on its timing.

While prior work in CRM has examined the importance of timing in interventions (e.g., Ascarza et al. 2018) and found heterogeneous impact, they do not find opposite effects within an individual at different points in time. Our result is novel in that even within a player, the same intervention can act as a substitute or a complement depending on when it is introduced. This insight highlights the potential value of considering dynamics in the timing of an intervention.

Figure 10 **Player Monetization and Retention Under Optimal Uniform Tool Giving Policy**

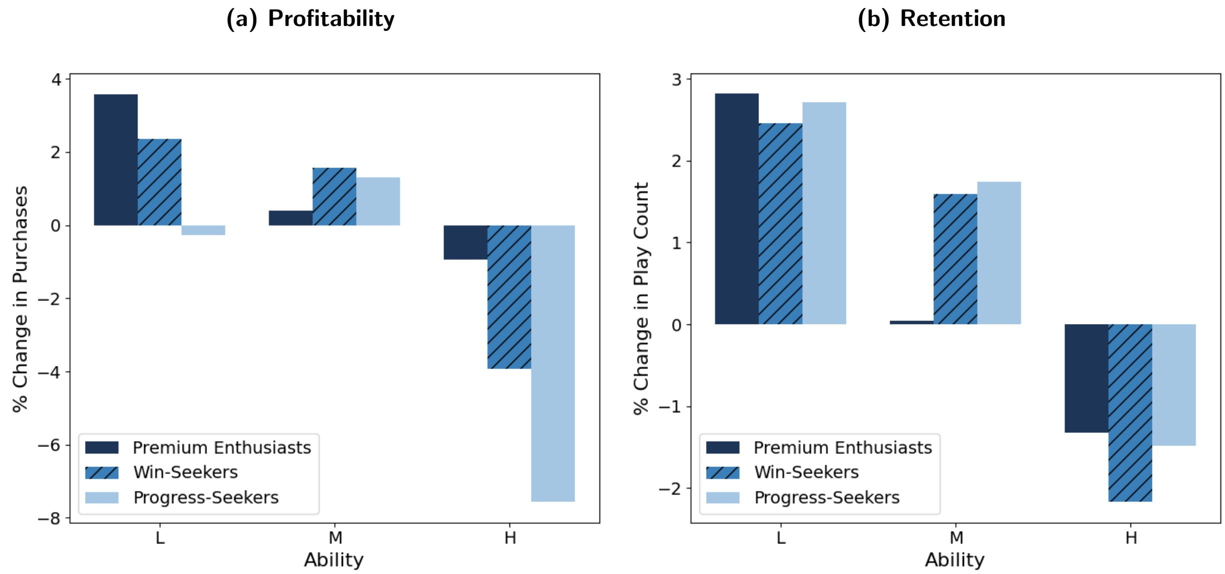

Figure 11 **Effect of Tool Giving Timing on Player Purchases**

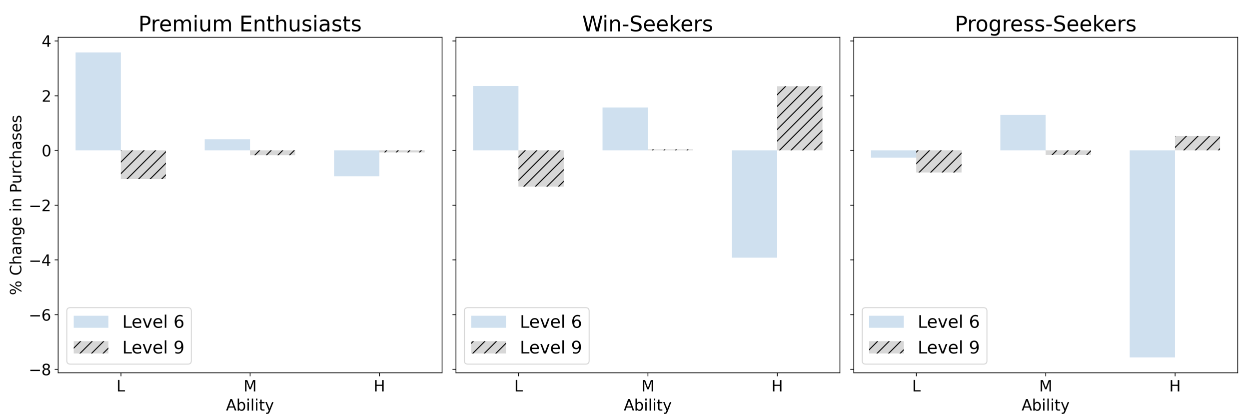

Personalized Tool-Giving Policies. The evidence that tool interventions can produce opposing effects across players and over time highlights the value of personalization. Using our estimated model, we seek to quantify the value of personalized policies that optimize both the target and timing of tool-giving to maximize net dynamic complementarities.

Figure 12 shows that the personalized policy increases incremental profits by over three-fold—from 0.64% under the best uniform policy to 2.09%. Personalization is also more resource-efficient, requiring 34% fewer tools. Figure 13 further shows that across all segments and ability levels, personalization consistently outperforms the uniform benchmark.

Figure 12 **Player Monetization and Retention Across Free Tool-Giving Policies**

(a) Monetization **(b) Retention**

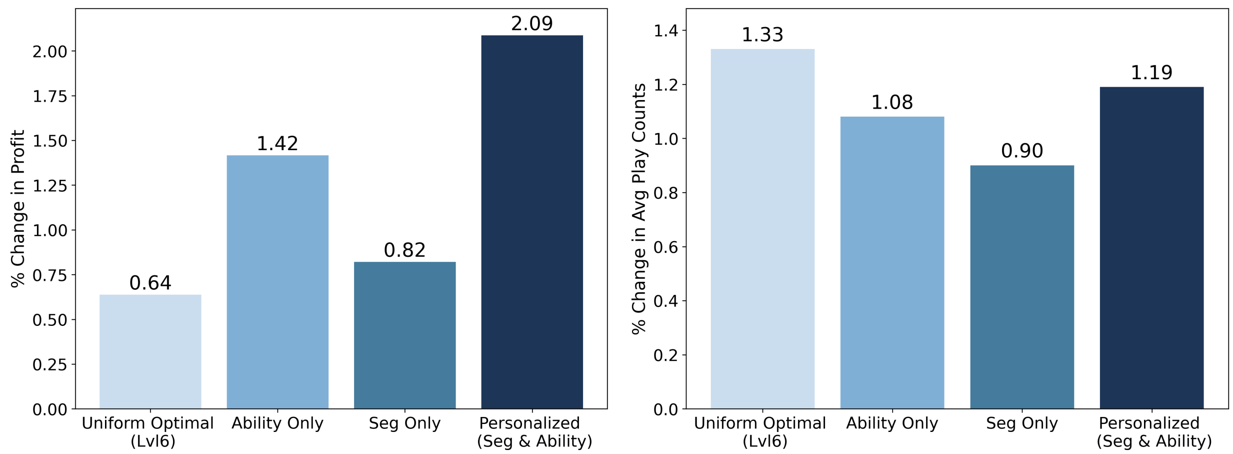

Figure 13 **Player Monetization and Retention Under Personalized Tool Giving Policy**

(a) Monetization **(b) Retention**

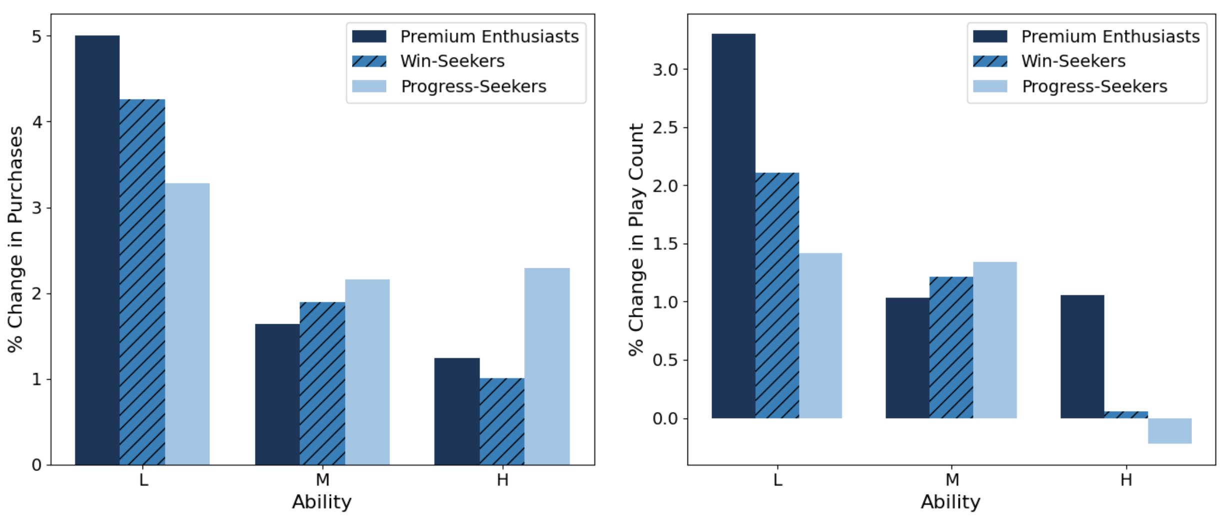

7.2. Personalized Game Progression Difficulty Design

Our first set of counterfactuals examined the dynamic interlinkages between purchases via tool-giving interventions that increased players' win probabilities. We next turn to game difficulty personalization, a more flexible design for managing both retention and monetization. Unlike tool-giving, difficulty adjustments can make the game either easier or more challenging, allowing the environment to adapt dynamically to player segments, abilities, and progression stages.

To operationalize personalized difficulty design, we divide game levels into two groups—low levels (6–8) and high levels (9–11)—and conduct a grid search over difficulty adjustments within a ±5% range of expected win probability (including no change) to maximize profits. These small, controlled changes preserve the integrity of the original game structure while enabling nuanced adaptation across player types.

We use the uniform profit-maximizing policy as a benchmark to evaluate the benefits of personalization in Table 10. Again, we show that personalization improves profit considerably over the uniform policy of easing the game. The second column reflects a policy where difficulty adjustments are limited to easing the game. Even under this constraint, personalization still improves profits compared to uniform, suggesting that the current game design is, overall, too difficult for many players. Finally, in the most flexible case, where difficulty adjustments include making the game more challenging when appropriate, the firm captures further profit improvements.

Table 10 Player Monetization and Retention Across Game Difficulty Adjustment Policies

	Difficulty Adjustment Policy		
	(1)	(2)	(3)
	Uniform Best	Personalized (- Difficulty)	Personalized (+/- Difficulty)
Monetization (% Change in Profit)	+ 1.66	+ 4.92	+ 5.92
Retention (% Change in Avg Play Count)	+ 2.78	+ 3.68	+ 3.95
% Change in Avg Level Reached	+ 3.73	+ 3.60	+ 3.45

Note: The uniform best policy: easing the game by 4% at lower levels and by 3% at higher levels.

Together, these findings highlight the value of our structural model in capturing the nonlinear and dynamic interplay between player heterogeneity and game design. Our counterfactual results align with recent experimental findings suggesting that easing the game can increase long-term monetization (Ascarza et al. 2025). While one possible outcome, our paper demonstrates that the effectiveness of such strategies depends not only on who receives the intervention, but also when, and how much difficulty is adjusted– underscoring the need for a model to clarify such underlying dynamics. The model helps firms to design scalable, real-time personalization strategies to optimize engagement and profitability.

8. Conclusion

This paper develops a dynamic structural model of consumer behavior in gaming environments, focusing on the joint and interdependent decisions of time spent playing and money spent on durable in-game tools. By modeling these decisions together, we uncover a central feature of virtual environments: dynamic complementarities and substitutions between play and purchase. These dynamics imply that even giving tools away for free can, under the right conditions, increase long-term engagement and firm profitability.

Our analysis reveals that the same intervention—such as a free tool or difficulty adjustment—can act as a complement or substitute depending on the player's ability, preferences, and game state. Counterfactual simulations show that personalized policies, which tailor both the timing and target of interventions, significantly outperform uniform strategies—delivering higher profits with fewer resources. 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. First, we abstract from potential interactions between player ability and tool effectiveness; this could further inform personalization and item design. Second, 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 personalized engagement and monetization strategies in such settings. By integrating the logic of dynamic effort response and durable goods choice, it provides a foundation for future work on real-time personalization where consumers co-produce utility through both time and monetary inputs.

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.

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

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.

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.

Appendix

A. Additional Model-Free Evidence

We report the Gini indices of player average win rate within each level in Table A.1. We note that even small variations in win rates result in considerable differences in expected points rewards and player progression speed, as shown previously in Figure 2.

Table A.1 Gini Indices of Player Average Win Rate by Level

Level	Gini Index
6	0.18
7	0.18
8	0.13
9	0.18
10	0.17
11	0.20

We next provide the full set of graphs showing the relationship between players' win probabilities and their purchase and exit decisions by level in Figure A.1. To address potential reverse causality—where purchasing tools could influence win rates—we plot players' purchase and exit decisions against their win rates from the previous level. Consistent with the aggregate pattern described in Section 3.3, within each level, players with higher win probabilities are more likely to continue playing and make in-game purchases, while those with lower win probabilities are more likely to exit and purchase less. In other words, higher win rates enhance players' continuation value and perceived utility from the game, whereas lower win rates diminish it, leading to higher exit rates and lower monetization.

Figure A.1 Player Exit and Purchase by Win Rate

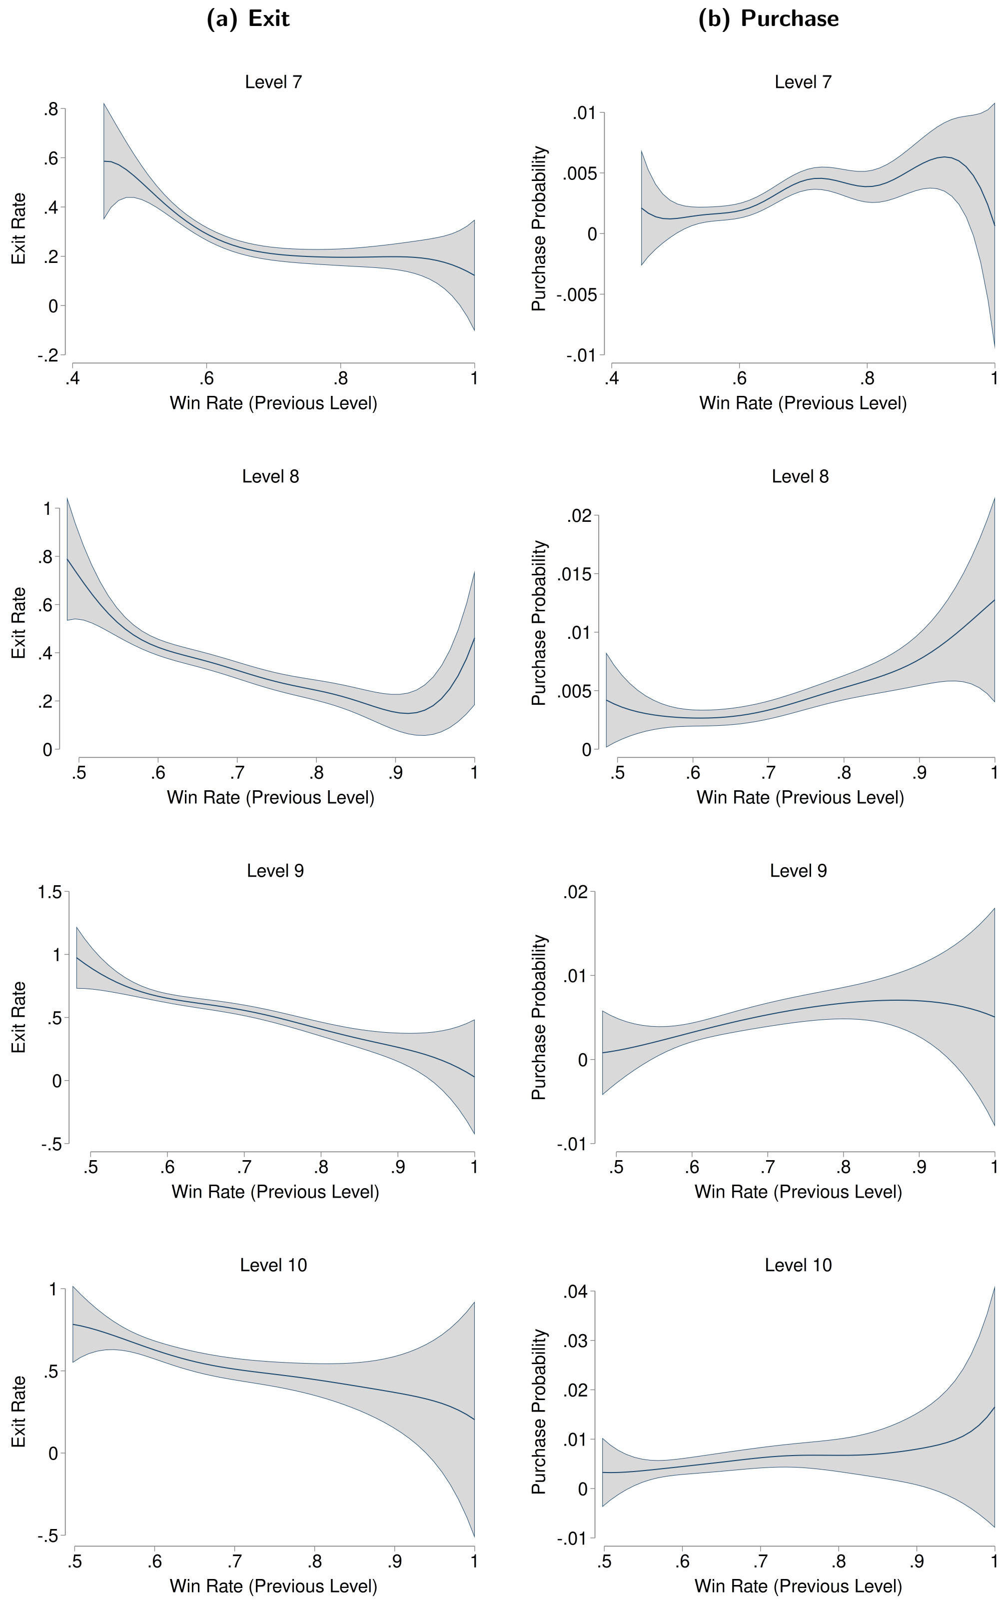

Note: The solid line is a kernel-weighted second-degree polynomial regression using a Gaussian kernel, and the shaded area is the 95% confidence interval.

B. Player Win Production Function Specification: Effect of Learning

If learning occurs and player ability improves over time, win probabilities should increase across levels. However, in our game setting, the difficulty of the level also increases, which counteracts these improvements. As a result, what we capture in the level effects parameter is the net effect of both learning and increasing difficulty. Our current model implicitly absorbs player learning as part of this net effect.

To provide support, we perform a variance decomposition analysis by removing variables from the model and comparing the reduction in mean squared error (MSE) relative to our baseline model. We present the results of the variance decomposition of our win production function specification in Table B.1. We find that ability explains approximately 79.3% of the total incremental variance, while level effects accounts for 17.5%. The remaining effects of learning by doing, as captured by the cumulative number of games played, contribute only 3.2% to the incremental variance explained. Given that 97% of the total incremental variance is already captured by ability and level progression, explicitly incorporating learning – such as tracking the number of games played – would increase our state space dramatically without significant explanatory gains. Note that this interpretation relies on the assumption that the rate of improvement is parallel across players, and that differences in learning rates do not systematically bias the estimated level effects. Given the small residual variance remaining, and the fact that much of the learning effect is implicitly absorbed through the net effect of levels, we abstract away from such effects in our main model.

Table B.1 Variance Decomposition of Player Performance

Variable	Variance Explained (MSE)	
	Incremental Variance Explained By	Share of Total Incremental Variance
Ability FE	0.0089	79.3%
Level FE	0.0020	17.5%
Number of Games Played	0.00036	3.2%
Total	0.011	100%

As an alternative test, we examine whether ability fixed effect estimates improve from low to high levels. If learning plays a substantial role, we would expect to see a large increase in ability fixed effects at higher levels, reflecting improved player skill over time. To ensure comparability between the ability estimates at low (Level: 6,7,8) and high levels (Level:9,10,11), we focus on a subsample of players who reached the final level.

Table B.2 reports summary statistics for players' ability fixed effects across low and high levels. While the results suggest that some learning occurs on average ($p<0.01$), it does not lead to substantial improvements in estimated ability (i.e., less than a 2.5% increase on average). This finding is consistent with our earlier variance decomposition analysis, which showed that much of the learning effect is already captured as net effects of players' level states.

Table B.2 Summary Statistics of Ability Fixed Effects (Low vs High Levels)

	Average	Median	75th	SD
Ability FE - Low Levels	0.9639	0.9650	1.0120	0.0934
Ability FE - High Levels	0.9877	0.9890	1.0310	0.0809
Difference (H-L)	0.0238	0.0240	0.0570	0.0670

C. Effect of Tool Purchase on Player Win Probability

We provide additional evidence in estimating the effect of tools purchase on player win probability by conducting a localized before-and-after analysis of player win rates. This analysis focuses on player performance immediately before and after a tool purchase. Specifically, we compare the win rates of players over the course of five games before and five games after they make an in-game purchase.

The after coefficient in Table C.1 captures the immediate effects of the purchased tool on player win probability, controlling for ability, level, and opponent effects. We find that one tool purchase translates to around 3.6 percentage points increase in win probability. This effect size is largely consistent with the effects of a tool measured in our win probability function, which quantifies the overall impact of tool purchase on player win probability across their entire match records.

Table C.1 Linear Probability Model: Before and After Purchase (5 Games)

	(1)
	Win
ability	0.52063***
	(0.02648)
after	0.03554***
	(0.00679)
lvl=7	0.02375**
	(0.01136)
lvl=8	0.02674**
	(0.01129)
lvl=9	-0.07839***
	(0.01148)
lvl=10	-0.05936***
	(0.01279)
lvl=11	-0.10529***
	(0.01271)
opponent elo score	-0.14020***
	(0.01811)
Constant	0.33278***
	(0.01627)
Observations	20241
Adjusted R^2	0.029

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

D. Additional Descriptive Statistics of Estimated Player Segments

This section provides additional descriptive statistics of the estimated player segments. Note that for ease of interpretation, we refer to the player segments by their descriptive names (1:Premium Enthusiasts, 2: Win-Seekers, and 3: Progress-Seekers) throughout this appendix. The formal introduction of these segment names is provided in Section 6.3.

First, we visualize the survival rate and the ability distribution of each segment in Figure D.1. Premium enthusiasts, which comprise the smallest share of players, exhibits the highest survival rates, followed by win-seekers, while progress-seekers, the largest segment, shows lower survival rates and ability distributions. Although win-seekers have slightly higher average ability than premium enthusiasts, premium enthusiasts persist longer in the game, consistent with their greater enjoyment of gameplay and lower price sensitivity.

Figure D.1 Descriptive Characteristics of Segment: Survival Rate and Ability Distribution

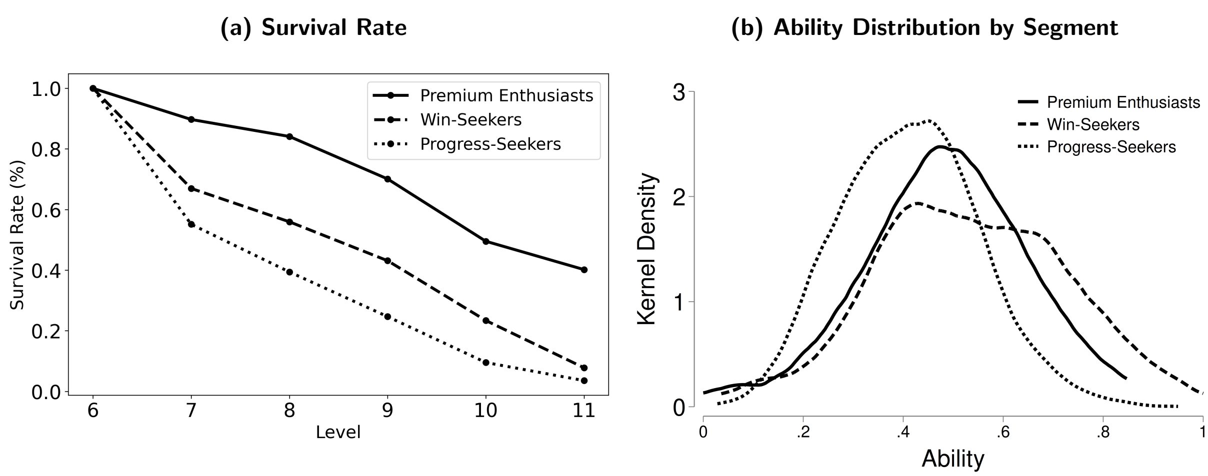

Next, we report Table D.1, which summarizes the average total purchases across segments. A key takeaway from these statistics is the presence of high spenders in every segment, underscoring the importance of accounting for both unobserved heterogeneity and ability.

We further examine how purchase behavior varies with ability by analyzing the relationship between total purchases and player ability within each segment in Figure D.2. Premium enthusiasts make purchases across all ability levels, whereas among win-seekers, lower-ability players are more likely to make purchases given

Table D.1 Average Total Purchases by Segment

Segments	Mean	Median	75Q	90Q	**99Q**
Premium Enthusiasts	6.88	5	9	11	**24**
Win-Seekers	2.77	1	3	7	**25**
Progress-Seekers	0.25	0	0	1	**5**

their sensitivity to immediate wins. Progress-seekers tend to spend the least, with purchases primarily coming from a small subset of high-ability players who are motivated by leveling up and completing the game.

Figure D.2 Total Purchases Across Ability by Segment

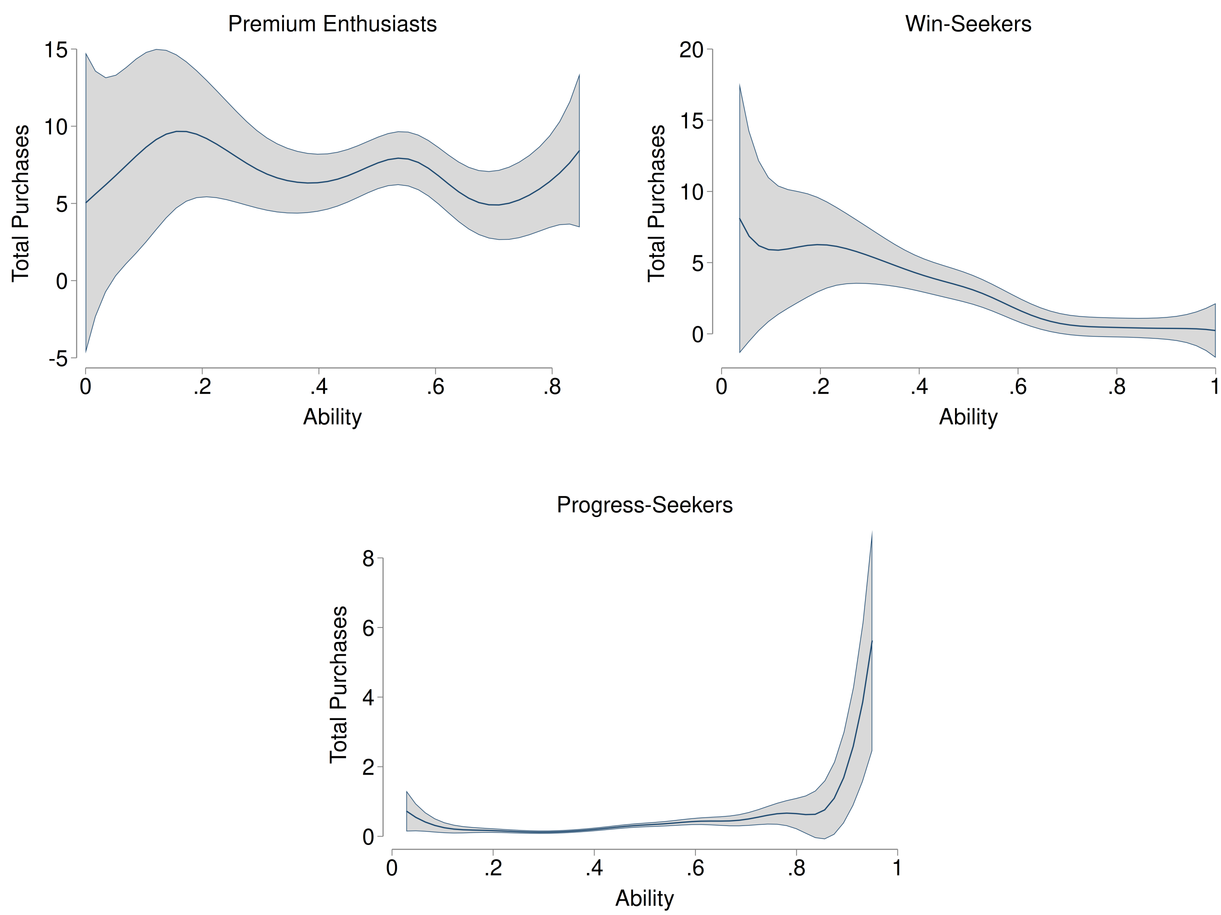

Figure D further illustrates how the timing of purchases varies across player segments and ability levels. We show that premium enthusiasts make purchases throughout their gameplay, rather than at specific levels. For win-seekers, purchases are more likely to be made among low-ability players, who purchase more as they progress to higher levels. Conversely, among progress-seekers, high-ability players are more likely to make purchases, who also purchase more as they reach higher levels.