informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–14
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Behavioral Market Design for Online Gaming Platforms

Ala Avoyan,[a,*] Robizon Khubulashvili,[b] Giorgi Mekerishvili[c]

[a] Department of Economics, Indiana University, Bloomington, Indiana 47405; [b] Department of Economics, University of San Francisco, San Francisco, California 94117; [c] AlixPartners LLP, Boston, Massachusetts 02110

*Corresponding author

Contact: aavoyan@iu.edu, https://orcid.org/0000-0002-7597-2045 (AA); rkhubulashvili@usfca.edu, https://orcid.org/0000-0002-3077-8679 (RK); gmekerishvili@alixpartners.com (GM)

Received: November 2, 2021
Revised: January 30, 2023; October 24, 2023
Accepted: March 3, 2024
Published Online in Articles in Advance: September 16, 2024

https://doi.org/10.1287/mnsc.2021.03628

Abstract. In this paper, we investigate market design for online gaming platforms. We ask what motivates people to continue participation—success or failure. Using data from an online chess platform, we find strong evidence of heterogeneous history-dependent stopping behavior. We identify two behavioral types of people: those who are more likely to stop playing after a loss and those who are more likely to stop playing after a win. We propose a behavioral dynamic choice model in which the utility from playing another game is directly affected by the previous game's outcome. We estimate this time nonseparable preference model and conduct counterfactual analyses to study alternative market designs. A matching algorithm designed to leverage stopping behavior can substantially alter the length of play.

History: Accepted by Yan Chen, behavioral economics and decision analysis.
Supplemental Material: The online appendix and data files are available at https://doi.org/10.1287/mnsc.2021.03628.

Keywords: time nonseparable preferences • platform design • online gaming • history-dependent stopping behavior • chess.com

1. Introduction

What determines our decision of when to stop a given endeavor? Does our past success motivate the stopping decision, or is a failure the primary determining factor? This paper focuses on the online gaming industry. Specifically, we explore the motivation behind stopping behavior using data from an online chess platform. The online gaming industry generated $162.3 billion in revenue in 2020 and is predicted to reach an annual gross revenue of $295.6 billion by 2026.[1] In this context, we investigate whether wins or losses influence people to play another game. Utilizing the identified behavioral patterns, we develop a theory that offers insights into how to encourage or discourage users on the platform from playing additional games.

We collect data from chess.com, the leading online chess platform boasting over 77 million users, where an average of 11 million chess games are played daily.[2] We select a random sample of users and scrape the entire history of their play for the years 2017 and 2018. Using the 2017 data and based on their stopping behavior, we identify 79% of the players as behavioral types and the remaining 21% of the users as nonbehavioral types. Among the behavioral group, about 30% are *win-stoppers* (players who are substantially more likely to stop playing after a win), and 70% are *loss-stoppers* (players who are substantially more likely to stop playing after a loss).[3] When classifying the same players using the 2018 data, we observe that their classifications remain stable over time for the vast majority of individuals. That is, 76.4% of users are identified as being the same type in 2017 and 2018. Because the user's type seems consistent over time, the following pattern might be relevant for various interventions; loss-stoppers play more when they win, whereas win-stoppers play more when they lose. Consequently, by increasing or decreasing the user's chances of winning a game, the platform can alter the likelihood of the user playing another game.

We develop a theoretical framework to further study and quantify the impact of changing the likelihood of winning for different types. Our model allows for time nonseparable preferences, in which *future* game utility can depend on the history of play.[4] The structural estimates from the model are consistent with the above-mentioned reduced-form evidence. For some people, a loss in a given game decreases the utility of playing another game, whereas for others, it increases the utility from playing another game. We show that matching win-stoppers with, on average, more challenging opponents increases the average number of games played.

We use the structural estimates to conduct counterfactual analyses, exploring the outcomes of alternative matching algorithms and quantifying the effects of such alterations. The platform currently prioritizes matching similarly rated players. Changing the matching algorithm, which results in changing the winning chances, impacts users' continuation likelihood. To illustrate, modifying a pairing that decreases a win-stopper's winning percentage from 50% to 45% (or 40%) results in a 4% (or 6%) increase in the average number of games played during a session. Similarly, a pairing that increases a loss-stopper's winning percentage from 50% to 60% (65%) can increase the average number of games played by a loss-stopper during a session by 1% (8%). To put these numbers in context, consider that over the course of a year, a 5% increase in session duration translates to the average user playing an additional 45 games, amounting to an extra 6 hours and 37 minutes spent on the platform.[5]

Gaming platforms have several key objectives, including gaining and retaining user popularity while generating profits through various channels, such as in-app advertising, subscriptions, and sponsorship. How could the platform use the information about the users' behavioral types to achieve these goals? First, increased user engagement in gaming sessions presents more opportunities for the platform to display advertisements.[6] Second, an essential aspect of the online chess experience is the speed at which players are matched with opponents. The platform's ability to efficiently match players within their skill level significantly impacts the user experience. The findings of our study can help platforms increase market thickness by motivating players to play more, which is particularly important during periods of low user activity online.

Our methodology is adaptable to other online platforms provided that two conditions are met. First, there needs to be an environment in which a person repeatedly makes a decision. Second, one needs definitions of what constitutes success and failure in a given environment. Under these two conditions, the methodology developed in the paper can be applied to new data from other settings. More generally, our findings could be applied in various other environments beyond online gaming. Consider the advantages of identifying a student's behavioral type, enabling educators or tutors to tailor the curriculum for improved learning outcomes. For instance, students categorized as loss-stoppers might benefit from a gradual introduction to new concepts, whereas those classified as win-stoppers might thrive when presented with more significant challenges to sustain their interest. Observing and identifying behavioral types in children could enable parents to frame problems in ways that cater to their child's personality, ultimately enhancing their chances of success. Our approach takes the types as given and creates an environment that could benefit all types of individuals.

2. Literature Review

Fundamentally, this paper presents and estimates a dynamic discrete choice model in which the agent may have time nonseparable preferences over the stochastic outcomes of their actions. In that sense, the application is analogous to the optimal stopping problems faced by, for example, taxi drivers, whose decisions to end their shifts may be influenced by their recent fares (see Camerer et al. 1997).[7] Recent empirical research on this topic is complicated by spatial search frictions and is limited by the imperfect observability of both decision makers' identities and decision makers' histories of the outcome. In contrast, in the current paper, the data allow us to observe the stopping decisions, outcomes, and independent realizations of each agent's decision problem. We take advantage of the rich data to demonstrate that an agent's decisions cannot be reconciled in a model without time nonseparable preferences and that there is substantial heterogeneity in preferences across players. By structurally estimating a model with heterogeneous time nonseparable preferences, this paper contributes to a growing body of literature on structural behavioral economics (see DellaVigna 2018 for a review of the studies on structural estimation of behavioral models).

This paper also contributes to the literature on the source of motivation, particularly the effects of wins and losses on future behavior. The existing findings in this literature are mixed. For example, Haenni (2019) and Cai et al. (2018) show that past failure has a discouraging effect on amateur tennis players and workers, respectively. In contrast, in a study of National Basketball Association and National Collegiate Athletic Association basketball players, Berger and Pope (2011) find an encouraging effect of being slightly behind at halftime. We deviate from this literature by focusing on heterogeneity among players rather than an overall effect. We find that losses have encouragement effects for some individuals and discouragement effects for others.

This paper is related to the literature on reference dependence—the effect that has been documented in various settings. For example, researchers have explored reference dependence for cab drivers' labor supply (Crawford and Meng 2011), professional golf players' effort choice (Pope and Schweitzer 2011), risky choices in the Deal or No Deal game (Post et al. 2008), domestic violence (Card and Dahl 2011), and police performance after a lower than expected pay raise (Mas 2006). We examine two types of reference dependence in the paper. First, we assume that the reference point is a player's rating at the start of a session. Second, we assume that the reference point is the expectation of winning based on the opponent's rating. That is, if the opponent has a higher rating, the player is more likely to expect to lose and vice versa. We calculate the magnitudes of these effects in our data, and we find that they are fairly limited; the reference dependence effect magnitudes are roughly 17–70 times smaller compared with the impact of the last game outcome.

Finally, the paper is related to studies using chess data. Researchers have used data from chess games to study risk, time, and other behavioral preferences for different age and gender groups.[8] The closest parallel to the current study is a paper by Anderson and Green (2018), in which the authors use data on blitz games played on the Free Internet Chess Server between 2000 and 2015.[9] The authors show that players are more likely to stop playing after they set a new personal best rating. This is an interesting result; however, players rarely set such records.[10] Anderson and Green (2018) show that players, on average, achieve a new personal best rating only twice every 15 years. In contrast, the current study focuses on the impact of the previous game, which affects a user's decision after every game.

3. Data and Descriptive Results

In this section, we offer an overview of the chess.com platform, describe our data collection process, provide some definitions, and present descriptive results. We highlight patterns that suggest history dependence and heterogeneity in stopping behavior. We conclude by providing potential explanations for the observed behavioral types.

3.1. About chess.com

We scraped the data from chess.com, the world's most popular online chess platform, which caters to a diverse user base spanning from amateur enthusiasts to elite professionals. Notably, Magnus Carlsen, the reigning World Chess Champion from 2013 to 2023, is among the platform's users. The platform chess.com offers free registration, enabling anyone to engage in matches against human or computer opponents via the website or the mobile app. Beyond gameplay, users can access other resources, including chess lessons and puzzles.

Upon registration on chess.com, a player is assigned an initial rating. During the data collection period, the default starting rating was 1,200.[11] Subsequently, a player's rating adjusts following each rated game, considering the game's outcome and the opponent's rating. Consequently, a player's current rating serves as a reflection of the player's current proficiency in chess; a higher rating signifies greater skill.[12] We recover the rating updating rule from the data. Additionally, we recover from the data the matching mechanism, which closely follows the rules stated on chess.com: "When you choose to play a rated game with a specific time control (like 5 [minutes]), we try to find you an opponent who is closest to your current rating."

3.2. Data Collection

We conducted our data collection in two phases. Utilizing Python's Selenium package and the chess.com application programming interface (API), we compiled a list comprising 1,793,473 usernames. In the next step, we focused on collecting users' game histories against human players. To prevent potential issues with web page access, we limited our analysis to a subset of users. Our approach involved randomly selecting 1,000 usernames at a time and extracting their year 2017 history of play using the chess.com API. We repeated this procedure 41 times. We then repeated the data collection process to gather the game histories of the same users for the year 2018.

Each observation within the data set contains information pertaining to the user and game characteristics, including the username, the user's self-identified country of association, the user's platform rating, the game's duration, the game type, which user had white pieces, the game's start and end times, and its ultimate outcome. For a summary of the data set, please refer to Table 1.

During the data cleaning phase, we excluded users who had not participated in any games during 2017. Given our focus on relatively quick decision making, we omitted "Daily" games from the sample because they are long and can extend over several days. Additionally, we removed unrated games, accounting for 0.3% of the data. A game was designated as unrated if any result of the game did not impact the users' ratings. For the analysis presented in Section 3.4, we did not impose any further restrictions on the data. However, in instances where we introduced additional constraints during the analysis, we have detailed those specifics within the respective sections.

Table 1. Data Description

Games	50,165,970	Average number of sessions	630
Sessions[a]	13,237,558	Average session length	5.11
Users	20,997	Average rating[b]	1,218
Rated games	99.7%	$Pr(Win \mid White\ Pieces)$	50.9
Game types		$Pr(Win \mid Black\ Pieces)$	47.0
Blitz	71.9%	$Pr(Win)$	48.9
Bullet	21.7%	$Pr(Loss)$	47.9
Daily	2.2%	$Pr(Draw)$	3.2

Notes. The top left quadrant presents the number of games, sessions, and users in the sample. The top right quadrant presents per-user information on the number of sessions played, session length, and user rating. The bottom left quadrant presents the characteristics of games in the data set: the fraction of rated vs. unrated games and the fraction of the top three common game types. Finally, the bottom right quadrant presents information on the outcomes of games and highlights the small percentage (probability) of drawn games.

[a]In Section 3.3, we provide a formal definition of a session; alternative definitions and corresponding results are in Online Appendix D.1.

[b]The average number of sessions was calculated in two steps. First, we calculated the number of sessions for each user in the data. Second, we averaged across all users. Similarly, the average session length and the average rating were calculated in two steps.

3.3. Definitions

A *game g* is a single game played against a human opponent. A collection of games ordered by time stamp, $(g_1, g_2, \ldots, g_n)$, is called a *session* if no game was played T minutes before g_1 or after g_n, and for any $i \in \{1, \ldots, n-1\}$, the time between g_i and g_{i+1} is less than T.[13] We call sessions that contain only one game ($n = 1$) *only game*. For sessions with $n \geq 2$, g_1 is the *first game*, g_n is the *last game*, and any game between the first game and the last game is referred to as the *middle game*. Based on the terms defined above, we categorize games into four mutually exclusive groups: only (O), first (F), middle (M), and last (L) games.[14]

Let $f_W(\cdot)$ be a function that calculates the winning percentage in a particular type of game; for example, $f_W(L)$ is a user's winning percentage in the last games. In some cases, when the context is clear, instead of writing $f_W(O), f_W(F), f_W(M), f_W(L)$, we write O, F, M, and L to indicate the winning percentage in only, first, middle, and last games, respectively.

3.4. Descriptive Results

We first establish that session-stopping behavior is history dependent. We then provide evidence of heterogeneity in stopping behavior, define behavioral types, generate predictions for extreme type definition, and examine time stability of behavioral types.

3.4.1. History Dependence.

Consider a null hypothesis that a user decides to stop the game randomly; in other words, stopping behavior is history independent.

Hypothesis 1. *Users' stopping behavior is independent of the outcome of the previous game.*

In this case, the winning percentage in the last games should be similar to the winning percentage in any other type of game.

For *each user*, we calculated the winning percentages in the first, middle, and last games, as defined in Section 3.3. Figure 1 illustrates the relationship between the winning percentages in last games and middle games, with each point depicting one user (the solid line represents the linear regression line).

The null Hypothesis 1 posits that the correlation between the winning percentages in last games and middle games will be close to one. We find it to be -0.49 and statistically different from 1 with $p < 0.001$. Thus, at the aggregate level, the decision to stop is not random, and we reject Hypothesis 1.

3.4.2. Behavioral Types.

We indeed reject the null hypothesis of history independence; however, the alternative hypothesis does not elucidate the precise relationship between the outcome of the preceding game and the decision to engage in another one. Further exploration is essential to discern whether users exhibit a propensity to stop a session following a win or a loss.

Figure 1. (Color online) Winning Percentage by Game Category

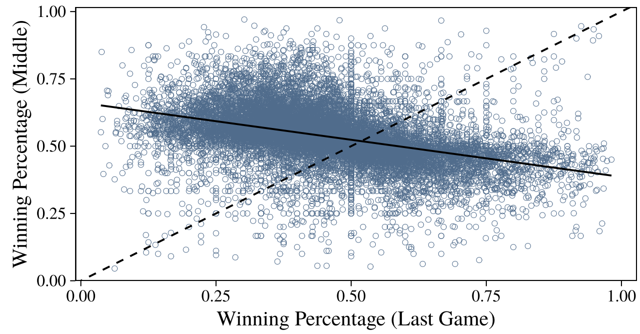

A closer examination of Figure 1 reveals an intriguing pattern; certain individuals demonstrate a considerably higher winning percentage in last games compared with middle games, whereas others exhibit a lower winning percentage in last games than in middle games. To unveil this heterogeneity and categorize users into distinct and exclusive types, we introduce the following definition.

Definition 1. A user is a behavioral type at the tolerance level of τ and referred as

- a win-stopper if $f_W(L) > f_W(M) + \tau$ and
- a loss-stopper if $f_W(L) < f_W(M) - \tau$.

A user is a nonbehavioral type and referred as a neutral type if $f_W(L) \in [f_W(M) - \tau, f_W(M) + \tau]$.

We describe how we calculate $f_W(\cdot)$ to unpack the above definition of behavioral types. For illustrative purposes, we focus on $f_W(F)$ for some user A. We take this single user's playing history for the year 2017 and look at every session that this user has played that lasted at least two games. For all these sessions, we examine the outcomes of only the first games that the player played and calculate the winning percentage by counting the number of wins. Say player A played 500 sessions that lasted at least two games and won 225 of the first games of each session, so $f_W(F) = 225/500 = .45$. In a similar fashion, we calculate $f_W(M)$, $f_W(L)$, and $f_W(O)$ for each user.

Recall that $f_W(M)$ tells us a player's winning probability in middle games, which can be thought of as the player's most typical games. If $f_W(M) = .5$, roughly speaking, the player wins 50% of the middle games that she plays. If the decision to stop the game is random and is not dependent on the outcome of the previous game, then there should be no difference between the winning probabilities in the middle and last games. Hence, we should have $f_W(M) \approx f_W(L) \approx .5$. We say $\approx$ to emphasize that we allow for some tolerance τ in Definition 1.

What does it mean if $f_W(M) = .5$ but $f_W(L) = .25$? Despite the ability to win 50% of typical games, the player is more likely to stop when she loses, leading to $f_W(L) < f_W(M)$. Definition 1 would classify this player as a loss-stopper as long as $\tau < 25\%$.

Finally, let us look at the decomposition of users into types using Definition 1. We classify users into types using data from sessions that lasted two or more games. At a tolerance level of $\tau = 7\%$, we find that 79% of users are behavioral types.[15] That is, for 79% of users in our data, the difference between their typical winning percentage and the winning percentage in the last game is at least 7%. Within this group of behavioral types, about 30% are win-stoppers, and 70% are loss-stoppers.

Using a moderately conservative threshold (equivalent to one and a half standard deviations of the winning probability distribution), a substantial portion of users exhibit history-dependent stopping behavior. To demonstrate that the tolerance level is large enough and that the results are not driven by chance, we simulated data with a random stopping rule. In simulated data, each player plays the number of games that the player plays in our actual data. The decision to stop or not stop is decided randomly with an equal chance. Simulated data show that if the stopping decisions were random, we would have classified 26% of the users as behavioral types instead of 79%. This strong evidence warrants further investigations into the existence and stability of such behavioral types.

3.4.3. Predictions. To further explore the existence of the behavioral types, we take Definition 1 to the extreme. Let us assume that there are behavioral types such that a win-stopper always stops after a win and a loss-stopper always stops after a loss. This extreme definition has a number of implications. That is, if there exist these behavioral types, we expect to see several patterns in the data. We formulate them as predictions.

Prediction 1. The correlation between the winning percentages in last games and only games is positive.

Prediction 1 follows from two observations. First, the winning percentage for win-stoppers in both only games and last games must be 100. Second, the winning percentage for loss-stoppers in both only games and last games must be zero. This is because if a win-stopper wins the initial game, she ends the session, and the game is classified as an only game. On the other hand, if this user loses the first game, she will start another game, making this session at least two games long; hence, the first game will be classified as the first game of a session and not as the only game. Therefore, a win-stopper's only games are always wins. In addition, whenever the extreme win-stopper wins, she ends the session, and we classify that game as the last game if the session is at least two games long. Therefore, the winning percentage in the last games is 100. Following similar logic for loss-stoppers, we get that the winning percentage for loss-stoppers in both only games and last games must be zero. The combination of these two observations across types and players leads to Prediction 1.

Figure 2(a) presents a scatterplot of the winning percentages in last games and only games. A strong positive relationship between the two winning percentages implies that individuals who are more likely to stop playing on a win (loss)—in other words, those who have a high winning (losing) percentage for last games—also have a higher winning (losing) percentage in only games, providing support for Prediction 1. Furthermore, we find that the winning percentage for only games is more than 25 percentage points higher for win-stoppers (64.0%) than for loss-stoppers (38.8%). Given that the average winning percentage in all games is 50.9% for win-stoppers and 50.4% for loss-stoppers, we can rule out the possibility that win-stoppers are simply better chess players.[16]

Figure 2. (Color online) Winning Percentage in Different Game Types

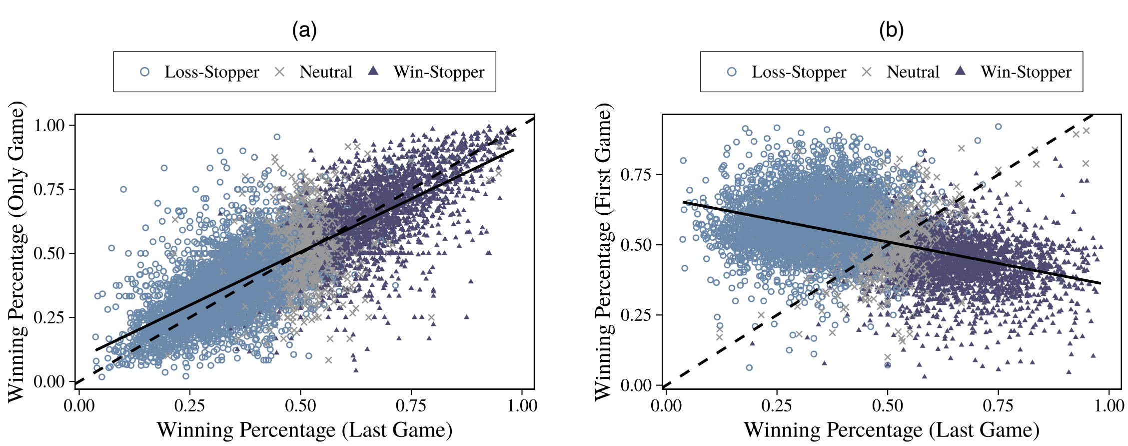

Notes. (a) Last and only games. (b) Last and first games.

Prediction 2. The correlation between the winning percentages for first games and last games is negative.

Let us examine the logic behind Prediction 2. If a loss-stopper wins the initial game, she plays another one, and thus, the initial game is classified as a first game. In contrast, if a loss-stopper loses the initial game, she stops playing, and thus, the initial game is classified as an only game. Therefore, using the extreme types, a loss-stopper's winning percentage for the first games must be 100, and by definition, the winning percentage for the last games must be 0. Similarly, for win-stoppers, the winning percentage in the first games must be 0, whereas the winning percentage in the last games must be 100. The combination of these two observations across types and players results in Prediction 2.

Figure 2(b) presents a scatterplot of winning percentages for the first games and last games. A strong negative relationship implies that individuals who are more likely to stop playing on a win (loss) have a lower winning (losing) percentage in the first games, supporting Prediction 2. Following a similar intuition as in Prediction 1 and Prediction 2, in Online Appendix B, we formulate and evaluate four other predictions about the relationships between the winning percentages in different types of games. Similar to Prediction 1 and Prediction 2, we find strong evidence in support of the four additional predictions, further highlighting the core findings on heterogeneous behavioral types.

3.4.4. Time Stability of Behavioral Types. We have classified users into types using the 2017 data. Here, we use the data for the same users from the year 2018 and classify them again using Definition 1. For each user, we have two labels: classification from 2017 and classification from 2018. We compare these classifications and calculate the fraction of users for whom the classifications match. We find a 76.4% match. Thus, 76.4% of users are identified as having the same type in 2017 and 2018. Furthermore, from the users identified as behavioral types in 2017, 84.5% of them are classified as the same type using 2018 data. Figure 3(a) presents the transition matrix between types from 2017 to 2018. Neutral types are most likely to experience a shift in classification. This result is not surprising because the definition of types is based on a threshold level, and most movement happens near this threshold.

When collecting the data, we did not place any restrictions on users' history. Some users played numerous games in 2017 and only a few in 2018, implying that the user's behavioral classification in 2017 is more accurate than the one in 2018 (because of the number of observations for this user). In addition, some users started playing late in 2017 (and therefore, played few games) but played many games in 2018. We show that this data limitation explains some of the movement between types as observed in the transition matrix. We redo the above analysis on a subsample of users who have played at least 300 games in both year 2017 and year 2018. Figure 3(b) presents the transition matrix between types from 2017 to 2018. As expected, the more information that we have on a user (more observations per user for each year), the more accurate the classification is; hence, there are fewer transitions that we find between the categories. See Online Appendix D.5, where we examine the time consistency for more active users.

As highlighted in Section 3.4.2, to ensure that our results are not driven by chance and a low threshold level, we simulated data for the year 2018 as well with a random stopping rule. We find that if the stopping rule was random, we would expect 74% of win-stoppers and loss-stoppers to change their type from 2017 to 2018. Instead, we see that 25% of win-stoppers and 12% loss-stoppers transition between types.

Figure 3. (Color online) Time Stability of Behavioral Types

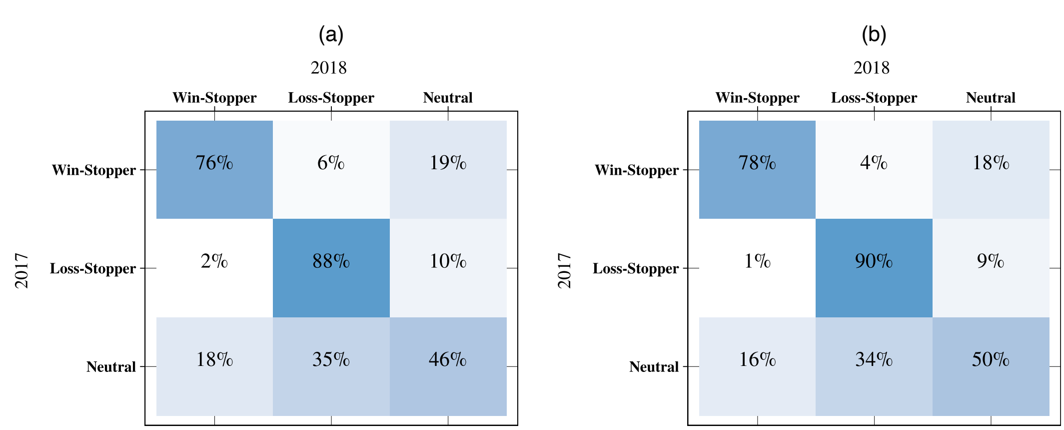

Notes. (a) Entire data set. (b) Subsample.

3.5. Theories Behind Behavioral Types

Before presenting our model, we explore potential explanations for the observed patterns of the two behavioral types in our data. Can reference dependence, fatigue, the gambler's fallacy, the hot-hand fallacy, or learning account for these patterns? We begin with reference dependence. One plausible explanation is that a user's personal best rating serves as a reference point; a user concludes a session when achieving a new personal best rating but continues playing otherwise. Although reference dependence can predict one type of behavior—ending a session after a win—it falls short in explaining loss-stoppers' behavior. Loss-stoppers' patterns do not align with similar reference-dependent reasoning.

Now, consider the idea of users' fatigue as they engage in successive games. Fatigue might lead to a decline in performance. We observe lower last-game scores among loss-stoppers, but the opposite holds for win-stoppers, who tend to achieve higher scores in their last game.

Next, let us explore two belief-based explanations: the gambler's fallacy and the hot-hand fallacy.[17] The gambler's fallacy implies the regression of events to the mean; if something occurs more frequently than usual during a given period, it will happen less often in the future. This fallacy suggests that if a player wins several games in a row, they might believe that their chances of winning again are reduced, leading them to stop after a win. Although the gambler's fallacy can explain patterns observed among win-stoppers, it contradicts loss-stoppers' behavior.

As for the hot-hand fallacy, some athletes (and their fans) believe that after succeeding several times in a row, they have a "hot hand," meaning that they are more likely to succeed in their next attempt. According to this belief, a player should continue playing after a win and stop after a loss because it indicates the end of their "hot hand." This reasoning can explain the last-game results of loss-stoppers but not win-stoppers.

Finally, let us turn to the concept of learning in games literature, which draws from the reinforcement learning literature in psychology (a meta-analysis of the learning literature on public good games by Cotla 2015 examines other possible learning models and suggests that learning aligns more closely with reinforcement learning as opposed to belief-based or regret-based learning). Research has shown that in repeated games, choices that yield higher payoffs in the past are more likely to be chosen in the future (e.g., Roth and Erev 1995, Chen and Tang 1998, Erev and Roth 1998, Haruvy and Stahl 2012). If a player plays online chess primarily for the pleasure of winning, the player is more likely to continue after a win and stop after a loss. This is similar to the evidence from the reinforcement learning literature in psychology, where the "win-stay, lose-shift" strategy is documented in many environments, such as repeated games (Posch 1999), sports (Tamura and Masuda 2015), and even among dogs (Byrne et al. 2020). This perspective could explain the behavior of loss-stopper types.

To summarize the discussion above, it can very well be that the entire population is a mixture of people who follow or fall into different theories or principles. In this paper, we refrain from taking a stance on the specific underlying psychological forces driving such behavior. Instead, we propose an approach that accommodates win-stoppers, loss-stoppers, and neutral types within the same model, focusing on outcomes and accounting for heterogeneity.

4. The Model

In this section, we initially present the model featuring three distinct player types. Subsequently, we provide an overview of our identification strategy.

4.1. Description

Here, we outline a chess player's dynamic decision-making process. Let y_t denote the player's rating at time t, which is observable to the player, the player's opponent, and the econometrician. We assume that the player's rating y_t belongs to a finite space denoted as Y. There are three distinct player types: win-stopper (θ_W), loss-stopper (θ_L), and neutral (θ_N). Let $\Theta = \{\theta_W, \theta_L, \theta_N\}$ be the set of all types, and let θ be an element of this set. Importantly, a player's type remains fixed over time.

A player's type profile at time t, denoted as (y_t, θ), consists of the player's time-variable characteristics, y_t, and a static, unobservable type θ. To denote current states, we use variables without time subscripts, whereas we use "prime" superscripts to represent states in the subsequent period.

Each period, a player faces the following decision; considering the outcome of the previous game, the player's type, and the player's current rating, the player must decide whether to engage in another game or opt for the outside option by going offline. Prior to making this decision, the player evaluates the expected utility from participating in an additional game, which is calculated as follows:

$$U(\theta, y, \chi) = u(y) + (1 - \chi) l_\theta, \tag{1}$$

where θ is the player's type, y is the player's current rating, χ is the outcome of the just-concluded game, and l_θ is the magnitude of the effect of the previous game outcome.[18] Note that l_θ can vary based on type $\theta \in \{\theta_W, \theta_L, \theta_N\}$, allowing for asymmetric effects (not restricting $l_{\theta_W} = -l_{\theta_L}$). If a player won a just-concluded game ($\chi = 1$), the utility from playing another game is $u(y)$. This term quantifies how much the player enjoys playing chess independently of the player's type. In the event of a loss in the previous game, the player's utility from playing another game is contingent on the player's type.

Definition 2. A player is a behavioral type if $l_\theta > 0$ or $l_\theta < 0$. She is

- a win-stopper if $l_\theta > 0$ and
- a loss-stopper if $l_\theta < 0$.

A player is a neutral type if $l_\theta = 0$.

There is an outside option, c, which is independently drawn from a distribution with density $f(c)$ in every period. If a player ends a session, she takes the outside option c. If the player does not end the session, her utility is $U(\theta, y, \chi)$ from playing a new game, and she moves to the next period. At this point, the player faces the same decision with an updated game history incorporating the result of the just-concluded game that she played (χ'). In each period, following the conclusion of a game, a player's decision problem gives rise to the following Bellman equation:

$$V(\theta, y, \chi, c) = \max\left\{c,\ u(y) + (1-\chi)l_\theta + \delta \sum_{\substack{y', \chi' \in \\ Y \times \{0,1\}}} p(y', \chi' | y) V(\theta, y', \chi')\right\}, \quad (2)$$

where δ is the discount factor and $p(y', \chi' | y)$ is the joint probability of the player receiving (transitioning to) the rating y' and the outcome of the next game being χ', conditional on the player's current rating y. We have

$$p(y', \chi' | y) = \sum_{y_{-i} \in Y} p(y, y_{-i}) p(y' | y, y_{-i}, \chi') p(\chi' | y, y_{-i}), \quad (3)$$

where $p(y, y_{-i})$ is the probability that a player with rating y is matched with a player with rating y_{-i}; $p(y' | y, y_{-i}, \chi')$ is the probability of receiving (transitioning to) rating y' given that the player's current rating is y; that in the next game, she is matched with a player with rating y_{-i}; and that the outcome of the next game is χ'. Note that we recover $p(y, y_{-i})$, $p(y' | y, y_{-i}, \chi')$, and $p(\chi' | y, y_{-i})$ from the data. In our counterfactual analysis, a player-to-player matching mechanism, $p(y, y_{-i})$, is a lever that market designers can use to influence a player's decision to start a new game.

4.2. Identification

The identification and estimation of the theoretical model follow the tradition of Hotz and Miller (1993). We show that we can forgo numerical dynamic programming to compute the value functions for every parameter vector, and we propose an estimation procedure that is simple to implement and computationally efficient. More details and most of the proofs are relegated to Online Appendix A.

We begin by identifying behavioral types. First, we establish that an optimal stopping rule is a threshold rule. Subsequently, we demonstrate that these thresholds exhibit the following characteristics. (i) They are higher after a loss than after a win for win-stoppers, (ii) they are lower after a loss than after a win for loss-stoppers, and (iii) they are equal after a loss and after a win for neutral types. This outcome implies that for a given player, comparing the *probability of ending a session after a win with the probability of ending it after a loss* allows us to determine the player's behavioral type.

Claim 1. The optimal stopping rule is a threshold rule in c.

Proof. Note that in Equation (2), continuation values do not depend on the current realization of c. Hence, fixing the continuation values and current period utility from playing another game, the second term under the max operator is lower than the outside option, c, for sufficiently high c. Thus, we have a threshold, $\bar{c}(\theta, y, \chi)$, above which the player stops playing and takes the outside option. □

Therefore, $\bar{c}(\theta, y, \chi)$ is a threshold such that a player with type profile (θ, y) who has an outcome χ in the last game ends a session if and only if the realized c is at least as large as $\bar{c}(\theta, y, \chi)$. Recalling Equation (2), we have

$$\bar{c}(\theta, y, \chi) = u(y) + (1-\chi)l_\theta + \delta \sum_{\substack{y', \chi' \in \\ Y \times \{0,1\}}} p(y', \chi' | y) V(\theta, y', \chi'). \quad (4)$$

The following proposition leads to the identification of behavioral types.

Proposition 1.

i. $\bar{c}(\theta_W, y, 0) > \bar{c}(\theta_W, y, 1)$;
ii. $\bar{c}(\theta_L, y, 0) < \bar{c}(\theta_L, y, 1)$; *and*
iii. $\bar{c}(\theta_N, y, 0) = \bar{c}(\theta_N, y, 1)$.

Proof. The proof follows from Equation (4) and Definition 2. □

Proposition 1 suggests that the probability of win-stoppers choosing to play another game is greater when they lost the previous game compared with when they won and conversely for loss-stoppers. Meanwhile, for neutral types, the probability of continuing the session remains consistent, regardless of the last game's outcome. Building on Proposition 1, we can determine a player's behavioral type from the data by examining the player's stopping probabilities following wins and losses. We use the data to recover winning and matching probabilities.

To identify the remaining parameters of the model, we make an assumption regarding the parametric distribution of the outside option, opting for an exponential distribution for the estimation process.[19] With this assumption in place, we identify both l_θ and the value associated with continuing a session based on the stopping probabilities following wins and losses.

A player's value function relies on both the values associated with continuing a session and the parameter of the outside option distribution. Provided that we identify the value from continuing a session and we normalize the distribution parameter, we identify the value functions. Finally, we show that δ and the utilities from playing a game are identified using the player's value from continuing a session and her value function. For more comprehensive information on the identification process, along with relevant claims and their corresponding proofs, please refer to Online Appendix A.

5. Estimation and Counterfactual Analysis

In this section, we first provide further details on the restrictions imposed on the data for structural estimation. Subsequently, we present the outcomes of the structural estimation and conduct a counterfactual analysis.

5.1. Preliminaries

We impose two restrictions on the sample.[20] First, we consider blitz games to ensure a more uniform time spent per game. Second, we focus on games where the users' pregame rating falls within the range of 1,000–1,600. This range selection is informed by the average and the standard deviation of the blitz ratings in the data.[21] The second condition is imposed to ensure that users ratings are reasonably close. This serves two purposes. First, we avoid matching users with significantly different ratings, and second, it helps mitigate missing values in the rating transition matrix.[22] After applying these restrictions, our data comprise 9,192,795 observations from 10,395 unique users.

The next step in the structural estimation analysis involves partitioning the rating range into grids. To achieve this, we utilize the average rating change as a guideline. In the primary data, the average rating increase following a win is 8.02 points, whereas the average decrease after a loss is 7.97 points. Consequently, we segment the rating space [1,000, 1,600] into eight-point intervals, yielding a total of 75 grids.

Proposition 1 suggests that for neutral types, the stopping probability is the same after both wins and losses. However, for practical empirical analysis, we redefine neutral types as users whose stopping probabilities after wins and losses are κ-close: that is, $|Pr(Stop|Win) - Pr(Stop|Loss)| \leq \kappa$. Similarly, we modify the win-stopper and loss-stopper definitions such that a user is a win-stopper if $Pr(Stop|Win) - Pr(Stop|Loss) > \kappa$ and a loss-stopper if $Pr(Stop|Win) - Pr(Stop|Loss) < -\kappa$. In this section, we use $\kappa = 0.07$. Online Appendix D presents the estimation results using $\kappa = 0.05$ and $\kappa = 0.09$, and the model type decomposition as κ is varied between 0 and 0.2.

5.2. Structural Estimates and Counterfactual Analysis

Our estimation strategy parallels the identification proof outlined in Online Appendix A. Recall Equation (1), which is the expected utility from playing an additional game. Let us focus on the parameters: l_θ for $\theta \in \{\theta_W, \theta_L, \theta_N\}$, which represents the additional utility associated with the outcome of the previous game. Table 2 presents the estimates of l_θ for each type. To ensure stability of the results, we bootstrap the data 300 times (see Online Appendix D.4 for the distribution of the point estimates).

Table 2 reveals notable differences. For win-stoppers, the utility from playing another game is 0.678 higher after a loss compared with after a win, as expected. This boost in expected utility for win-stoppers in the subsequent games after a loss ($\chi = 0$) is in contrast to the lower expected utility after a win ($\chi = 1$). On the other hand, for loss-stoppers, the expected utility from playing another game is 0.610 lower after a loss compared with after a win. For neutral types, the result of the last game has no sizable effect on their utility.

Now, we move to address the question of how the expected session length is affected when we modify the probability of winning by adjusting the matching algorithm on the platform. Figure 4 illustrates the percentage change in the average session length (the x axis) in response to alterations in the winning percentage (the y axis) resulting from changes in the matching algorithm. The solid line in Figure 4 represents a winning percentage of around 50%. According to the definition of behavioral types, a decrease in the winning percentage for win-stoppers and an increase in the winning percentage for loss-stoppers should lead to longer average session lengths. Figure 4 validates this intuition and quantifies the effects. When we alter the matching algorithm to pair win-stoppers (triangles in Figure 4) with increasingly higher-rated opponents on average, their winning percentage decreases, but the average session length increases.

For win-stoppers, using a matching process that decreases the winning percentage from 50% to 45% increases the average session length by 3.75%. Using a

Table 2. Bootstrapped Values for l_θ

Parameter	Mean	Standard deviation
l_{θ_W}	0.678	0.005
l_{θ_N}	−0.014	0.003
l_{θ_L}	−0.610	0.002

Figure 4. (Color online) Winning Percentage and Percentage Change in Session Length

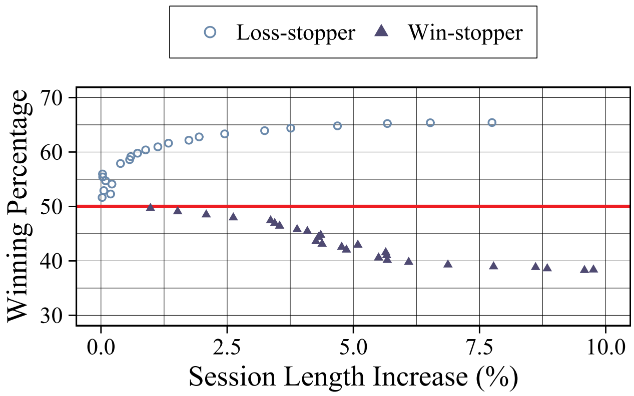

Notes. The *x* axis presents the percentage change in average session length, and the *y* axis depicts change in the winning percentage, which in turn, is a result of changing the matching algorithm. The winning percentage against a similarly rated player is around 50% (the solid line).

matching algorithm that drops the winning percentage to 40% increases the average session length by around 6%. Similarly, for loss-stoppers, using a matching algorithm that increases the chances of winning from 50% to 60% increases the average session length by 1%. Using a matching process that increases the winning percentage from 50% to 65% increases the average session length by more than 7.5%.

In the sample with sessions with only blitz games, an average user played 274 sessions per year. An average session lasted about 3.29 games, and the average blitz game lasted 7 minutes and 29 seconds. Thus, over one year, a 5% increase in session length results in an average user playing 45 more games or spending 6 hours and 37 minutes longer on the platform.

It is important to highlight that more games do not have to translate into extended time on the platform. One could further argue that asymmetric matching in ratings could even reduce the length of a game because a strong player could win the game faster against a weaker player. We explore these concerns in detail in Online Appendix E. In particular, Online Appendix E.1 shows that the median correlation between the minutes spent on a session and the number of games played during the session is 0.98 across users. Online Appendix E.2 displays that the correlation between opponents' rating difference and how long the game lasts is close to zero. We conjecture that the observed high correlation between the number of games and time on the platform and the close-zero correlation between rating difference and playing time are because of the nature of blitz games, which by definition, are time constrained. Taking all of the evidence together, we conclude that more games can result in more time spent on the platform.

5.2.1. The Effects of Practice and Welfare Discussion.

Let us explore a positive externality associated with playing more games: the impact of practice on a user's rating, which serves as a proxy for skill and is a socially desirable outcome. We analyze the highest rating achieved by a user in 2017 and the total number of games that the user played during the same year. A linear regression (see column (1) in Table J.7 in Online Appendix J for details) reveals a significant positive association between the number of games played and a higher rating.

To account for player-specific variations, we employ a fixed effects (FE) panel data estimation method. We define a unit as a month in the 2017 data. For each month, we calculate the number of games played and the highest achieved rating, resulting in 12 observations per typical player. The FE estimation, with the total number of games in a month as an independent variable and the highest rating achieved as the dependent variable, reaffirms the positive relationship between practice and skill improvement (see column (3) in Table J.7 in Online Appendix J).[23] Furthermore, we find that both win-stoppers and loss-stoppers experience similar effects from practice, suggesting that neither behavioral type holds a distinct advantage in skill development (see column (4) in Table J.7 in Online Appendix J).

The above discussion raises questions about the implications of extending time on the platform on overall welfare. Increasing the time spent playing can be seen as beneficial for the platform as it leads to more user engagement and for users themselves as it enhances their skill levels. However, we acknowledge that total welfare can be influenced by other factors not included in our model. The model assumes that players derive utility from playing another game and does not consider the potential negative effects of extended play. For instance, increased time on the platform might reduce a player's productivity if that time was originally intended for work or study. In this paper, we simplify our analysis by excluding such concerns. Nevertheless, future research could explore the unintended consequences of longer play in more detail.

6. Other Factors Influencing the Stopping Decision

Up to this point, our focus has centered on the influence of the just-completed game's outcome on the decision to play a new game. In this section, we consider additional factors that could potentially affect stopping decision. To assess these factors, we employ a Cox proportional-hazards (CPH) model, which enables us to examine how various characteristics, referred to as covariates, impact the session-stopping rate, also known as the hazard rate. This approach allows us to analyze the influence of specified factors on whether a session concludes or persists.

We employ a CPH model that incorporates time-dependent covariates. The general form of this model is outlined as

$$h_j(t, x_j(t)) = h_0(t) \exp\{x_j(t)'\beta\}. \quad (5)$$

Equation (5) breaks down as follows. The left-hand side signifies the risk that game j in period t, characterized by $x_j(t)$, marks the end of the session (i.e., the session ends after this game). The right-hand side of the equation comprises two elements: baseline risk and relative risk.

The baseline risk, denoted as $h_0(t)$, represents the risk of a game being the final game in a session when all covariates are set to zero ($x_j(t) = \mathbf{0}$). The relative risk, expressed as $\exp x_j(t)'\beta$, quantifies the proportional increase or decrease in risk associated with the covariates specified in $x_j(t)$.

We aim to find whether factors beyond the outcome of the just-completed game and a user's behavioral type influence the decision to stop. Prior to evaluating additional covariates, we estimate the model with three variables: the outcome of the just-completed game, a user's behavioral type, and the interaction of the two variables. To ease the interpretation of the results, we assume that there are no neutral types and that we only have two types of users: win-stoppers and loss-stoppers. Table 3 presents the results of the CPH estimation. The variable *Outcome* takes a value of one if a user won the game and zero otherwise. The variable *Type* is assigned a value of one for win-stoppers and zero for loss-stoppers. Consequently, the baseline for the estimation is a loss-stopper who experienced a loss in the game.

Table 3 reveals that for win-stoppers, the hazard rate is higher after a win than after a loss. This is evident in the estimates. After a win, the hazard rate is lower by 0.17 ($-0.58 - 0.69 + 1.10 = -0.17$) compared with the baseline (which in the estimation, is a loss-stopper who lost the game). After a loss, the hazard rate is lower by 0.69 compared with the baseline. In simpler terms, for win-stoppers, the probability of ending a session is lower after a loss than after a win, as expected. Conversely, for loss-stoppers, we observe the opposite relationship. A win in the just-completed game reduces the hazard rate by 0.58 compared with a loss (baseline). The importance of those coefficients is easier to comprehend

Table 3. Cox Proportional-Hazards Model with Type Heterogeneity

Covariate	Coefficient	exp(Coef)	p-value
Outcome	−0.58	0.56	<0.005
Type	−0.69	0.50	<0.005
Type × *Outcome*	1.10	3.01	<0.005

Note. Coef, coefficient.

using the information provided in the third column: exp(Coef). In particular, if exp(Coef) = 1, it implies that a given variable has no impact on the decision to end the session, whereas exp(Coef) being higher or lower than one indicates increased or decreased chances of ending a session, respectively. For example, the value of exp(Coef) for the *Outcome* variable is 0.56, signifying that a loss-stopper is 44% ($(1 - \exp(\text{Coef})) \cdot 100\%$) less likely to conclude the session after a win than after a loss.

Now that we have reaffirmed our results from previous sections using CPH analysis, we add other variables that we hypothesize may affect the stopping behavior. We consider the outcome in the game before the just-completed game, the opponent's rating, the user's initial rating, and two interaction terms. In what follows, we provide a motivation for why we consider these covariates.

The outcome of the game before the just-completed game (named the *previous outcome*) could potentially impact stopping decisions. For instance, some users might be more inclined to stop after experiencing two consecutive wins as opposed to just one, whereas others may be more likely to stop playing after enduring two consecutive losses.

Another factor that might sway a user's decision to stop is the user's initial rating within a session. If a user's stopping strategy entails concluding a session once the user's current rating surpasses the initial session rating, then the rating difference between the first and last games of a session becomes significant. Including this variable and its interaction with the *Outcome* variable helps us discern if there exists reference dependence concerning the initial rating and if users react differently when the same rating change is achieved after a win or after a loss.

We have also incorporated the opponent's rating into our analysis driven by the idea of reference dependence. Winning against a stronger opponent might be more satisfying than winning against a weaker opponent. On the other hand, losing to a weaker opponent may feel like a greater setback than losing to a stronger one. In other words, winning against a weaker opponent is typically expected, so losing in such a scenario might be seen as a more significant failure. To assess whether and to what extent winning or losing against opponents of varying strengths affects stopping decisions, we introduced the "opponent rating difference," which measures the disparity in ratings between the two users. By including the "opponent rating difference" and its interaction with the *Outcome*, we can control if the outcome of the just-completed game has a different effect for a given type if that outcome is achieved against a stronger or weaker opponent.

Figure 5 provides a visual illustration of the results when we include all of the variables together in the

Figure 5. (Color online) CPH Coefficients for Win-Stoppers and Loss-Stoppers

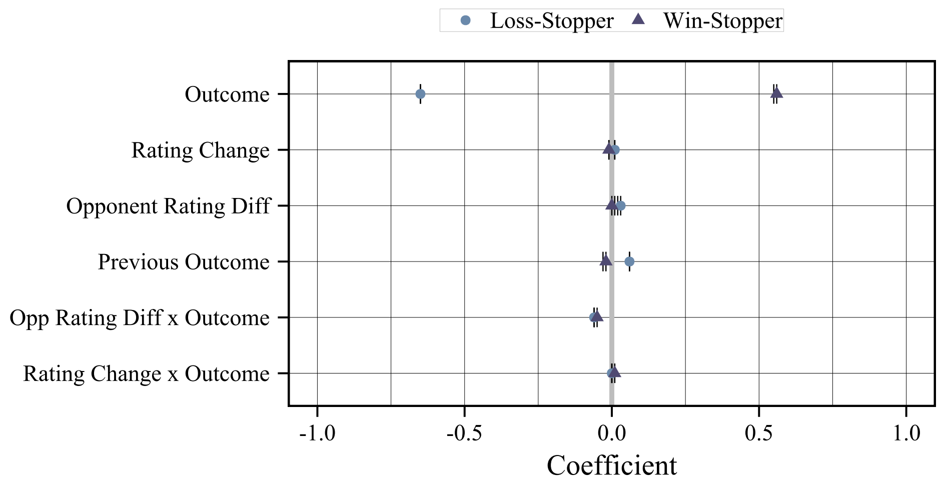

Notes. The coefficients represent the quantitative relationships between the variables and the session-stopping rate as per the Cox proportional-hazards model. These coefficients indicate whether a variable tends to increase or decrease the likelihood of ending a session. A positive coefficient suggests an increase, whereas a negative coefficient suggests a decrease in the session-stopping rate compared with the baseline conditions. The magnitudes of these coefficients are standardized to represent the strength of the effect in standard deviations, allowing for meaningful comparisons among different variables. Diff, difference; Opp, opponent.

CPH analysis (see Online Appendix F for the CPH results when we add variables one at a time). To ensure that the coefficients of these variables are easily comparable, we standardized both the rating change and opponent rating difference variables to have a mean of zero and a standard deviation of one.[24] We can easily see from Figure 5 that the magnitude of the effect of the *Outcome* is much larger for both win-stoppers and loss-stoppers than for all of the other variables that we consider. This result highlights the fact that the outcome of just-completed game carries the most substantial effect. This is not to suggest that the platform should ignore additional factors that influence a user's decision to play another game. Instead, we emphasize that in terms of magnitude, the outcome of the just-concluded game carries the most sizable effect.

7. Conclusion

This paper explores previously undocumented behavioral type heterogeneity in stopping behavior. We investigate stopping behavior on an online chess platform, shedding light on the factors influencing individuals' stopping decisions. Leveraging rich data spanning two years from chess.com, we categorize 79% of users as behavioral types, whereas the remaining 21% are considered nonbehavioral (neutral) types. Among the behavioral types, one third fall into the win-stopper category, with the remainder classified as loss-stoppers. Win-stoppers tend to halt their play after a win, whereas loss-stoppers are more inclined to stop after a loss. We then explore how platforms can utilize knowledge of user types to alter the number of games played.

Although the paper focuses on chess games on one platform, the model and the descriptive analysis can be applied to other environments as long as they satisfy two main conditions. First, there needs to be an environment where individuals repeatedly face a similar decision (for example, to play another level of the same game or to accept another passenger's request for pickup as a ride-sharing driver). Second, the outcome and wins/losses must be well defined (won/lost the new level).[25] For instance, this framework is applicable to many mobile games as they fulfill the conditions above.

It is crucial to note that for meaningful counterfactual analyses, the platform must have some control over altering *chances of whether a user succeeds or fails*. On chess.com, where game difficulty and rules are fixed, modifying winning probabilities involves influencing potential opponents' strength by adjusting the matching algorithm. In other games with variable difficulty levels or hints, the platform can alter the winning probabilities by modifying the underlying difficulty or providing hints. In the case of ride-sharing platforms, such as Uber and Lyft, the driver's type could be how the driver's stopping decision is affected by tips or the expected length of a ride. Then, the labor supply can be increased by using the information on riders' tipping behavior (or their expected length of the ride) and matching suitable riders with the driver. Future research can explore data from diverse environments to identify similar heterogeneous behavioral types. Additionally, investigating whether a person's behavioral type remains consistent across various settings could yield intriguing insights.

Acknowledgments

The authors thank Chloe Tergiman, Ran Shorrer, Kai Baron, Kalyan Chatterjee, Nino Doghonadze, Daniel Grodzicki, Peter Newberry, Nikoloz Pkhakadze, Daniela Puzzello, and Karl Schurter for their helpful comments. This paper benefited from comments from conference participants at the 2019 CESifo Area Conference on Behavioral Economics and seminar participants at Indiana University, Penn State University, the University of Pittsburgh, and the University of San Francisco. The authors also thank anonymous reviewers and the editor for their careful reading of the paper and their insightful comments and suggestions. The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the views and opinions of AlixPartners LLP.

Endnotes

1 The sources are https://www.mordorintelligence.com/ and https://www.statista.com/.

2 The top three most frequently visited chess sites, in order of popularity, are chess.com, chess24, and Lichess. Our data set includes users from 191 different countries; for a detailed country-level analysis, refer to Online Appendix K.

3 We refer to nonbehavioral types as *neutral* types. These players are about equally likely to stop playing after a win and after a loss (see Section 3.4.2 for a formal definition). The three types—win-stoppers, loss-stoppers, and neutral types—are mutually exclusive and collectively exhaustive categories.

4 See Braun et al. (1993) and Dragone and Ziebarth (2017) for evidence of time nonseparable preferences in aggregate consumption and novelty consumption, respectively. See Turnovsky and Monteiro (2007) for the effects of consumption externalities under time nonseparable preferences.

5 An increase in session length could cause several crowding-out effects. See Online Appendix E for a thorough discussion of why an increase in the number of games will likely correspond to an increase in time spent on the platform.

6 The game can be played on a computer (on the chess.com website) or a mobile app (on the chess.com app). The website displays ads during the entire game on the sides of the screen. On the mobile app, ads are displayed after the game. In both cases, we presume that more games (weakly) increase ad consumption.

7 See also Farber (2005, 2008, 2015), Abeler et al. (2011), Morgul and Ozbay (2015), Cerulli-Harms et al. (2019), Frechette et al. (2019), and Thakral and Tô (2021).

8 See the following recent papers that use chess data to study economic behavior: Gerdes and Gränsmark (2010), Gränsmark (2012), Dreber et al. (2013a, b), Bertoni et al. (2015), Linnemer and Visser (2016), De Sousa and Niederle (2022), and De Sousa and Hollard (2023).

9 On chess.com, blitz is a type of chess game in which each player has a specific amount of time (between 3 and 10 minutes) for the entire game. The blitz games analyzed in Anderson and Green (2018) lasted between 6 and 30 minutes.

10 On average, in our data set, it takes a user 119 games to surpass the previously recorded personal best rating. Further, more experienced users take longer to set new records. For example, users with at least 600 games take 275 games to reach a new personal best rating.

11 There was a change on chess.com regarding the initial rating assignment system, and now, users can choose to start from the rating of 400, 800, 1,200, or 1,600 based on their chosen skill level.

12 The average rating in our sample is 1,218, close to the initial rating. However, the majority of users in our sample are highly active players. The median and the mean number of games played by the users in both years are 2,389 and 1,206, respectively. Further, note that the rating reflects expertise conditional on experience. For example, user A, who just joined the platform and has a rating of 800, is not the same as user B, who has played 1,000 games and has a rating of 800.

13 For the main analysis, we set $T = 30$ minutes; we then vary T to check the robustness of the results, and we find no substantial differences. See Online Appendix D.1 for more details.

14 The procedure that we used to label the games within a session involves two steps. The first step of handling the data is removing daily games. The second step is to define sessions and label games according to the definitions in Section 3.3 *before* cleaning the data any further. This way, we avoid a game being classified as the last game of a session simply because the user changed the type of chess game that they are playing. We find that 96% of all sessions are homogeneous in terms of game type and game length. That is, players do not change the game type and the game time lengths within a session. Furthermore, among sessions that contain at least one blitz game, 97.98% are homogeneous. In other words, if a session has one blitz game, in 97.98% of times, all of the games in that session are blitz games.

15 One standard deviation in winning probabilities in any game in the data is around 4.7%. We do all of the analyses in the paper for tolerance levels of 5%, 7%, and 9%, which are around 1, 1.5, and 2 standard deviations, respectively. For the main part of the paper, we present the results with a 7% tolerance level. See Online Appendix D for the effects of changing the tolerance on behavioral decomposition and structural estimates.

16 Online Appendix G presents winning percentages and standard deviations for each behavioral type and game category. Further, Online Appendix I shows that there is no correlation between types and ratings.

17 Miller and Sanjurjo (2018) found evidence suggesting that the hot-hand fallacy might not be a fallacy in the context of basketball free throws.

18 Our model only considers wins and losses, omitting draws. Although draws are more common in classical chess, they are less frequent in fast chess. In our data, only 3.2% of the games ended in a draw. We treated draws as wins if they occurred against stronger opponents (players with higher ratings) and as losses if they occurred against weaker opponents. As a robustness check, we also estimated our model by excluding games that ended in a draw, and our estimation results remained unaffected by this change.

19 We select the exponential distribution because it imposes no additional implicit assumptions beyond being a continuous distribution on $(0, \infty)$. This is because of the fact that exponential distribution represents the distribution with maximum entropy among the class of continuous distributions on $(0, \infty)$ with a given mean (see Conrad 2004, theorem 3.3).

20 Structural estimation results using the complete data can be found in Online Appendix D.3.

21 The average blitz rating is 1,303, with a standard deviation of 324. We round these figures to 1,300 and 300, respectively, resulting in the rating range $[1{,}300 - 300, 1{,}300 + 300]$.

22 Consider a user with a rating of 700. In our data set, it is improbable that this user has ever played against an opponent with a rating of 2,000. When calculating the potential new rating for a user with a rating of 700 after playing against a user with a rating of 2,000, we need to consider all such games in the data set. However, because there may not be any single game with such a vast rating difference, we encounter missing values. To minimize estimation errors, we

limit our analysis to games for which we have a sufficient quantity of data.

[23] It is worth highlighting that most of our players are not top professional chess players. For professionals, practice might only help a little because they could have reached the limits of their abilities.

[24] Standardization ensures that if the effect of the rating change is 15 times smaller than the effect of the last game result, it is expressed in terms of 15 standard deviations (which corresponds to approximately a 1,080-rating-point difference) rather than 15 rating points.

[25] In the case of drivers on ride-sharing apps, it is not straightforward to define what would constitute a win or a loss. Receiving a tip or not is one of the possibilities. Other options could be the expected length of the ride or the drop-off location.

References

Abeler J, Falk A, Goette L, Huffman D (2011) Reference points and effort provision. *Amer. Econom. Rev.* 101(2):470–492.

Anderson A, Green EA (2018) Personal bests as reference points. *Proc. Natl. Acad. Sci. USA* 115(8):1772–1776.

Berger J, Pope D (2011) Can losing lead to winning? *Management Sci.* 57(5):817–827.

Bertoni M, Brunello G, Rocco L (2015) Selection and the age–productivity profile. Evidence from chess players. *J. Econom. Behav. Organ.* 110:45–58.

Braun PA, Constantinides GM, Ferson WE (1993) Time nonseparability in aggregate consumption: International evidence. *Eur. Econom. Rev.* 37(5):897–920.

Byrne M, Bray E, Maclean E, Johnston A (2020) Evidence for win-stay-lose-shift in puppies and adult dogs. Working paper, University of Arizona, Tucson.

Cai X, Gong J, Lu Y, Zhong S (2018) Recover overnight? Work interruption and worker productivity. *Management Sci.* 64(8):3489–3500.

Camerer C, Babcock L, Loewenstein G, Thaler R (1997) Labor supply of New York City cabdrivers: One day at a time. *Quart. J. Econom.* 112(2):407–441.

Card D, Dahl GB (2011) Family violence and football: The effect of unexpected emotional cues on violent behavior. *Quart. J. Econom.* 126(1):103–143.

Cerulli-Harms A, Goette L, Sprenger C (2019) Randomizing endowments: An experimental study of rational expectations and reference-dependent preferences. *Amer. Econom. J. Microeconomics* 11(1):185–207.

Chen Y, Tang F-F (1998) Learning and incentive-compatible mechanisms for public goods provision: An experimental study. *J. Political Econom.* 106(3):633–662.

Conrad K (2004) Probability distributions and maximum entropy. *Entropy* 6(452):10.

Cotla CR (2015) Learning in repeated public goods games—A meta analysis. Preprint, submitted September 12, https://dx.doi.org/10.2139/ssrn.3241779.

Crawford VP, Meng J (2011) New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *Amer. Econom. Rev.* 101(5):1912–1932.

DellaVigna S (2018) Structural behavioral economics. Bernheim BD, DellaVigna S, Laibson D, eds. *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 1 (Elsevier, Amsterdam), 613–723.

De Sousa J, Hollard G (2023) From micro to macro gender differences: Evidence from field tournaments. *Management Sci.* 69(6):3358–3399.

De Sousa J, Niederle M (2022) Trickle-down effects of affirmative action: A case study in France. NBER working paper, National Bureau of Economic Research, Cambridge, MA.

Dragone D, Ziebarth NR (2017) Non-separable time preferences, novelty consumption and body weight: Theory and evidence from the East German transition to capitalism. *J. Health Econom.* 51:41–65.

Dreber A, Gerdes C, Gränsmark P (2013a) Beauty queens and battling knights: Risk taking and attractiveness in chess. *J. Econom. Behav. Organ.* 90:1–18.

Dreber A, Gerdes C, Gränsmark P, Little AC (2013b) Facial masculinity predicts risk and time preferences in expert chess players. *Appl. Econom. Lett.* 20(16):1477–1480.

Erev I, Roth AE (1998) Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Amer. Econom. Rev.* 88(4):848–881.

Farber HS (2005) Is tomorrow another day? The labor supply of New York City cabdrivers. *J. Political Econom.* 113(1):46–82.

Farber HS (2008) Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *Amer. Econom. Rev.* 98(3):1069–1082.

Farber HS (2015) Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *Quart. J. Econom.* 130(4):1975–2026.

Frechette GR, Lizzeri A, Salz T (2019) Frictions in a competitive, regulated market: Evidence from taxis. *Amer. Econom. Rev.* 109(8):2954–2992.

Gerdes C, Gränsmark P (2010) Strategic behavior across gender: A comparison of female and male expert chess players. *Labour Econom.* 17(5):766–775.

Gränsmark P (2012) Masters of our time: Impatience and self-control in high-level chess games. *J. Econom. Behav. Organ.* 82(1):179–191.

Haenni S (2019) Ever tried. Ever failed. No matter? On the demotivational effect of losing in repeated competitions. *Games Econom. Behav.* 115:346–362.

Haruvy E, Stahl DO (2012) Between-game rule learning in dissimilar symmetric normal-form games. *Games Econom. Behav.* 74(1):208–221.

Hotz VJ, Miller RA (1993) Conditional choice probabilities and the estimation of dynamic models. *Rev. Econom. Stud.* 60(3):497–529.

Linnemer L, Visser M (2016) Self-selection in tournaments: The case of chess players. *J. Econom. Behav. Organ.* 126:213–234.

Mas A (2006) Pay, reference points, and police performance. *Quart. J. Econom.* 121(3):783–821.

Miller JB, Sanjurjo A (2018) Surprised by the hot hand fallacy? A truth in the law of small numbers. *Econometrica* 86(6):2019–2047.

Morgul EF, Ozbay K (2015) Revisiting labor supply of New York City taxi drivers: Empirical evidence from large-scale taxi data. *Transportation Res. Board 94th Annual Meeting*, vol. 15 (Transportation Research Board, Washington, DC), 1–19.

Pope DG, Schweitzer ME (2011) Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *Amer. Econom. Rev.* 101(1):129–157.

Posch M (1999) Win–stay, lose–shift strategies for repeated games—Memory length, aspiration levels and noise. *J. Theoret. Biol.* 198(2):183–195.

Post T, Van den Assem MJ, Baltussen G, Thaler RH (2008) Deal or no deal? Decision making under risk in a large-payoff game show. *Amer. Econom. Rev.* 98(1):38–71.

Roth AE, Erev I (1995) Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games Econom. Behav.* 8(1):164–212.

Tamura K, Masuda N (2015) Win-stay lose-shift strategy in formation changes in football. *EPJ Data Sci.* 4:1–19.

Thakral N, Tô LT (2021) Daily labor supply and adaptive reference points. *Amer. Econom. Rev.* 111(8):2417–2443.

Turnovsky SJ, Monteiro G (2007) Consumption externalities, production externalities, and efficient capital accumulation under time non-separable preferences. *Eur. Econom. Rev.* 51(2):479–504.