Behavioral Market Design For Online Gaming Platforms

October 24, 2023

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 non-separable 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.

JEL Classification: D9, D47, C5, C13;

Keywords: Online gaming platform design, time non-separable preferences, 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 non-behavioral 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. Since the user's type seems consistent over time, the following pattern might be relevant for various interventions: loss-stoppers play more when they win, while 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 non-separable

[1] Source: www.mordorintelligence.com and www.statista.com.

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

[3] We refer to non-behavioral 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.

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, while 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, resulting 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

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

[5] An increase in session length could cause several crowding-out effects. See 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 (chess.com app). The website displays ads during the entire game on the sides of the screen. On the mobile app, it is displayed after the game. In both cases, we presume that more games (weakly) increase ad consumption.

during periods of low user activity online.

Our methodology is adaptable to other online platforms, provided two conditions are met. First, there needs to be an environment in which a person repeatedly makes a decision. Second, one needs a definition 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 these 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, while 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 non-separable 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 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 non-separable preferences and that there is substantial heterogeneity in preferences across players. By structurally estimating a model with heterogeneous time non-separable preferences, this paper contributes to a growing body of literature on structural behavioral economics (see DellaVigna (2018) for a review of studies on structural estimation of behavioral models).

The paper also contributes to the literature on the source of motivation, particularly

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

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 NBA and NCAA basketball players, Berger and Pope (2011) find an encouraging effect of being slightly behind at half time. 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.

The paper is related to the literature on reference dependence—the effect, which 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 the reference point is a player's rating at the start of a session. Second, we assume 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 to 70 times smaller compared to 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 thr Free Internet Chess Server (FICS) 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,

[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 et al. (2021), and De Sousa and Niederle (2022).

[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 to 30 minutes.

[10] On average, in our dataset, it takes a user 119 games to surpass their 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.

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, and subsequently 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, catering 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. Chess.com offers free registration, enabling anyone to engage in matches against human or computer opponents via the website or mobile app. Beyond gameplay, users can access other resources, including chess lessons and puzzles, enhancing their chess experience.

Upon registration on chess.com, a player is assigned an initial rating. During the data collection period, the default starting rating was 1200.[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 their current proficiency in chess; a higher rating signifies greater skill.[12] We recover the rating updating mechanism from the data and it closely follows the rules stated on chess.com: "When you choose to play a rated game with a specific time control (like 5 min), we try to find you an opponent who is closest to your current rating."

[11] There was a change on chess.com regarding the initial rating assignment system, and now users can choose to start from rating of 400, 800, 1200, or 1600 based on their chosen skill level.

[12] The average rating in our sample is 1218, 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 2389 and 1206, 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 1000 games and has a rating of 800.

3.2 Data collection

We conducted our data collection in two phases. Initially, we gathered usernames from the platform without imposing any restrictions on their history of play. Consequently, there are some users in our sample that had not participated in any games during the year 2017. Additionally, certain user accounts were either created but never used or used solely in years beyond the scope of our study. Utilizing Python's Selenium package and the chess.com Application Programming Interface (API), we compiled a list comprising 1,793,473 usernames.

In the subsequent step, we focused on collecting users' game histories against human players. To prevent potential issues with webpage access, we limited our analysis to a subset of users. Our approach involved randomly selecting 1000 usernames at a time and extracting their 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 dataset contains information pertaining to the user and game characteristics, including the username, the user's self-identified country of association, their 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 dataset, 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, since 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.

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

[19] 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.

Table 1: Data description

Games	50,165,970	Average Number of Sessions	630
Sessions[13]	13,237,558	Average Session Length	5.11
Users	20,997	Average Rating[14]	1,218
Rated games	99.7%	Pr(win | white pieces)	50.9
Game types:		Pr(win | 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

Note: 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 3 common game types. Finally, the bottom right quadrant presents information on the outcomes of games and highlights the small percentage (probability) of drawn games.

3.3 Definitions

A *game* g is a single game played against a human opponent. A collection of games ordered by time stamp, $(g_1, g_2, ..., g_n)$, is called a *session* if no game was played T minutes before g_1 or after g_n, and for any $i \in \{1, ..., n-1\}$, the time between g_i and g_{i+1} is less than T.[15] We call sessions that contain only one game ($n = 1$) *only-game* (O-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 and the last 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.[16]

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,

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

[16] The procedure 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 definitions in Section 2.2 *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 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 the games in that session are blitz games.

when the context is clear, instead of writing $f_W(F)$, $f_W(M)$, $f_W(L)$, $f_W(O)$, we write F, M, L, and O, to indicate the winning percentage in first, middle, last, and only-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 and define behavioral types.

3.4.1 History dependence

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

> H_0*: 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).

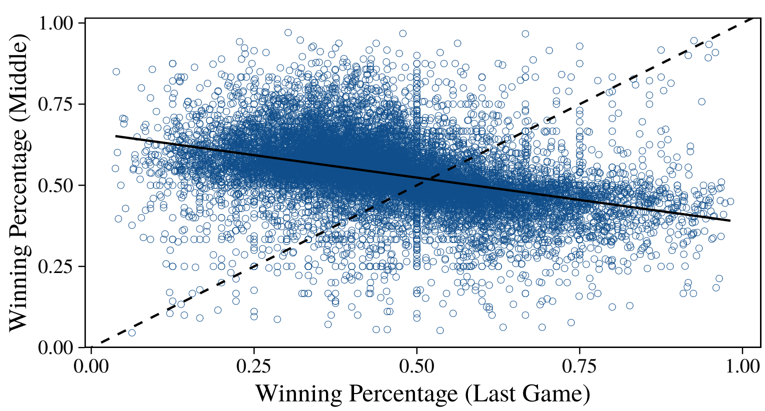

Figure 1: Winning percentage by game category

The null hypothesis H_0 posits that the correlation between the winning percentages in last-games and middle-games will be close to 1. 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 H_0.

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.

A closer examination of Figure 1 reveals an intriguing pattern: certain individuals demonstrate a considerably higher winning percentage in last-games compared to middle-games, while others exhibit a notably 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 non-behavioral 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 this user has played that lasted at least two games. For all these sessions, we examine the outcomes of only the first-games they 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 and each game type.

Now, note 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 she plays. If the decision to stop the game is random and does 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.[17] 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 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 they play in our actual data. The decision to stop or not is decided randomly with an equal chance. Simulated data shows 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 and last-games must be 100. Second, the winning percentage for loss-stoppers in both only and last-games must be 0. 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

[17] One standard deviation in winning probabilities in any game in the data is around 4.7%. We do all the analysis 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 Appendix D.2 and Appendix D.4 for the effect of changing the tolerance on behavioral decomposition and structural estimates, respectively.

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, 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- and last-games must be 0. The combination of these two observations across types and players leads to Prediction 1.

Figure 2a presents a scatter plot 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, confirming the 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.[18]

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, while 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 2b presents a scatter plot 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 Predictions 1 and 2, in Appendix B we formulate and evaluate four other predictions about the relationships between the

[18] Appendix G.1 presents winning percentages and standard deviations for each behavioral type and game category. Further, Appendix I shows that there is no correlation between types and ratings.

winning percentages in different types of games. Similar to Prediction 1 and 2, we find strong evidence in support of the 4 additional predictions, further highlighting the core findings on heterogeneous behavioral types.

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

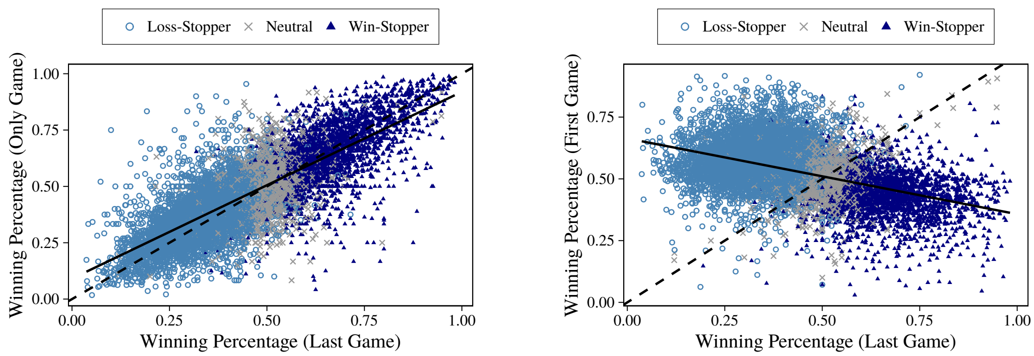

Figure 2: Winning percentage in different game types

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 2018, respectively. 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 3a 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 since the definition of types is based on a threshold level, and most movement happens near this threshold.

While 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 (due to 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 re-do the above analysis on a subsample of users who have played at least 300

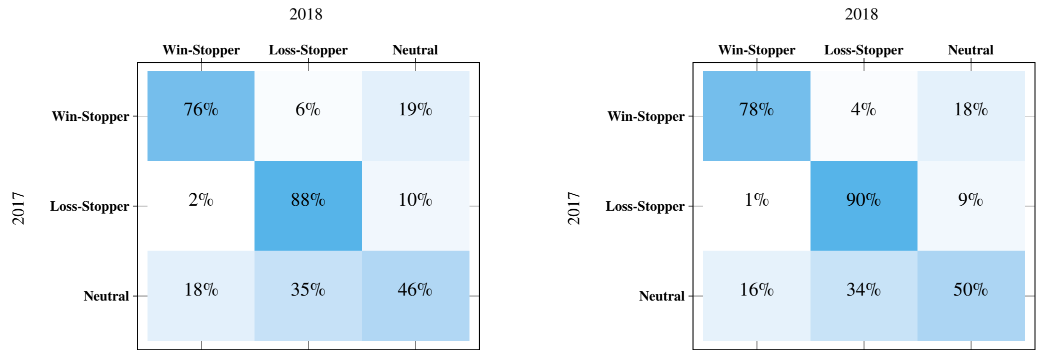

(a) Entire data set (b) Sub-sample

Figure 3: Time stability of behavioral types

games both in years 2017 and 2018. Figure 3b presents the transition matrix between types from 2017 to 2018. As expected, the more information we have on a user (more observation per user for each year) more accurate the classification is; hence, the fewer transitions we find between the categories. See Appendix D.5 where we examine the time consistency for more active users.

As highlighted in Section 3.4.2, to ensure our results are not driven by chance and 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.

3.5 Theories behind behavioral types

Before presenting our model, we explore potential explanations for the observed patterns of 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. While 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 and the hot hand fallacy.[19] 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 their chances of winning again are reduced, leading them to stop after a win. While 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 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 since 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 learning literature in 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), Erev and Roth (1998), Chen and Tang (1998), and Haruvy and Stahl (2012)). If a player plays online chess primarily for the pleasure of winning, they are more likely to continue after a win and stop after a loss. This is similar to the evidence from 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.

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

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, while we employ "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 their current rating, they 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 their 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 just concluded game, and l_θ is the magnitude of the effect of the previous game outcome.[20] 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 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 their type. In the event of a loss in the previous game, the player's

[20] Our model only considers wins and losses, omitting draws. While 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.

utility from playing another game is contingent on their 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 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:

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

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:

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

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, in the next game she is matched with a player with rating y_{-i}, and 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 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 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 to 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,

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

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) $\bar{c}(\theta_L, y, 0) < \bar{c}(\theta_L, y, 1)$;

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 to 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 their 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.[21] 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 Appendix A.

5 Estimation and counterfactual analysis

In this section, we first provide further details on 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.[22] First, we consider blitz games to ensure a more uniform time spent per game. Secondly, we focus on games where the users' pre-game rating falls within the range of 1000 to 1600. This range selection is informed by the average and standard deviation of the blitz ratings in the data.[23] The second condition is imposed to ensure that users ratings are reasonably close. This serves two purposes: firstly, we avoid matching users with significantly different ratings, and secondly, it helps mitigate

[21] We select the exponential distribution since it imposes no additional implicit assumptions beyond being a continuous distribution on $(0, \infty)$. This is due to 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 Theorem 3.3 in Conrad (2004)).

[22] Structural estimation results using the complete data can be found in Appendix D.3.

[23] The average blitz rating is 1303, with a standard deviation of 324. We round these figures to 1300 and 300, respectively, resulting in the rating range [1300-300, 1300+300].

missing values in the rating transition matrix.[24] After applying these restrictions, our data comprises 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, while the average decrease after a loss is 7.97 points. Consequently, we segment the rating space [1000, 1600] into 8-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 $\kappa = 0.07$. Appendix D.4 presents the estimation results using $\kappa = 0.05$ and $\kappa = 0.09$. Appendix D.2 presents the results of 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 Appendix A. Recall equation (1), the expected utility from playing an additional game. Let us focus on parameters: l_θ for $\theta \in \{\theta_W,\ \theta_L,\ \theta_N\}$, which represent 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 Appendix G.2 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 to 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 to after a win. For neutral types, the result of the last game has no sizable effect on their utility.

[24] Consider a user with a rating of 700. In our dataset, it is improbable that this user has ever played against an opponent with a rating of 2000. When calculating the potential new rating for a user with a rating of 700 after playing against a user with a rating of 2000, we need to consider all such games in the dataset. However, since 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.

Parameter	Mean	SD
l_{θ_W}	0.678	0.005
l_{θ_N}	-0.014	0.003
l_{θ_L}	-0.610	0.002

Table 2: Bootstrapped values for l_θ

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 (x-axis) in response to alterations in the winning percentage (y-axis), resulting from changes in the matching algorithm. The red solid line 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.

Figure 4: Winning percentage and percentage change in session length

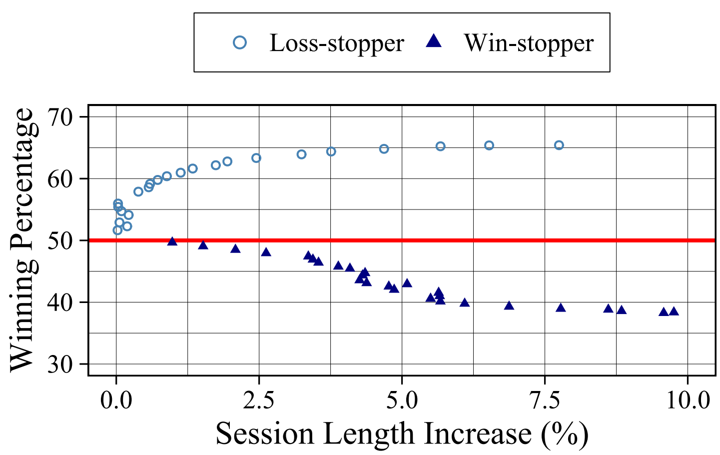

The x-axis presents the percentage change in average session length, and y-axis depicts change in the winning percentage, which in turn is a result of changing the matching algorithm. The winning percentage against a similarly rated player is around 50% (the red solid line).

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 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 an extended time on the platform. One could further argue that asymmetric matching in ratings could even reduce the length of a game since a strong player could win the game faster against a weaker player. We explore these concerns in detail in Appendix E. In particular, Appendix E.1 shows that the median correlation between minutes spent on a session and the number of games played during the session is 0.98 across users. Appendix E.2 displays the correlation between opponents' rating difference and how much time 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 close-zero correlation between rating difference and playing time is due to the nature of blitz games, which by definition, are time-constrained. Taken all the evidence together, we conclude that more games might result in more time spent on the platform.

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 they played during the same year. A linear regression (see column (1) in Table 10, 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 10, Appendix J).[25] 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 10, 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 by other factors not included in our model.

Our model assumes that players derive utility from playing another game and does not consider 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 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:

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

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

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

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 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 1 if a user won the game and 0 otherwise. The variable *Type* is assigned a value of 1 for win-stoppers and 0 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 to the baseline (which, in the estimation, is a loss-stopper who lost the game). After a loss, hazard rate is lower by 0.69 compared to 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 to a loss (baseline). The importance of those coefficients is easier to comprehend using the information provided in the third column: $\exp$(Coef). In particular, if $\exp(\text{Coef}) = 1$, it implies that a given variable has no impact on the decision to end the session, while $\exp$(Coef) being higher or lower than 1 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.

Covariate	Coefficient (Coef)	Exp(Coef)	p-value
Outcome	-0.58	0.56	< 0.005
Type	-0.69	0.50	< 0.005
Type $\times$ *Outcome*	1.10	3.01	< 0.005

Table 3: Cox Proportional Hazards Model with Type Heterogeneity

Now that we have re-affirmed 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 *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, while 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 their initial rating within a session. If a user's stopping strategy entails concluding a session once their current rating surpasses their 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 one. 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: CPH coefficients for win-stoppers and loss-stoppers

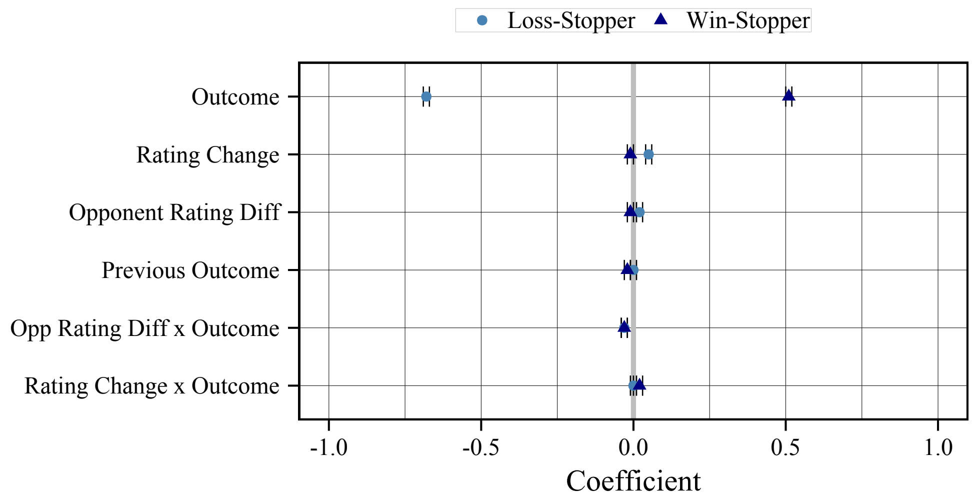

The coefficients represent the quantitative relationships between the variables and the session-stopping rate as per the Cox proportional hazards (CPH) model. These coefficients indicate whether a variable tends to increase or decrease the likelihood of ending a session. A positive coefficient suggests an increase, while a negative coefficient suggests a decrease in the session-stopping rate compared to the baseline conditions. The magnitude of these coefficients is standardized to represent the strength of the effect in standard deviations, allowing for meaningful comparisons among different variables.

Figure 5 provides a visual illustration of the results when we include all the variables together in the CPH analysis (see 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 0 and a standard deviation of 1.[26] 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 all the other variables 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.

[26] Standardization ensures that if the effect of the rating change is 15 times smaller than the effects of the last game result, it is expressed in terms of 15 standard deviations (which corresponds to approximately a 1080 rating point difference) rather than 15 rating points.

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, while the remaining 21% are considered non-behavioral (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, while 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.

While 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 pick-up as a ride-share driver). Second, the outcome and wins/losses must be well-defined (won/lost the new level).[27] 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 their 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.

[27] 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, Johannes, Armin Falk, Lorenz Goette, and David Huffman, "Reference points and effort provision," *American Economic Review*, 2011, *101* (2), 470–92.

Anderson, Ashton and Etan A Green, "Personal bests as reference points," *Proceedings of the National Academy of Sciences*, 2018, *115* (8), 1772–1776.

Berger, Jonah and Devin Pope, "Can losing lead to winning?," *Management Science*, 2011, *57* (5), 817–827.

Bertoni, Marco, Giorgio Brunello, and Lorenzo Rocco, "Selection and the age–productivity profile. Evidence from chess players," *Journal of Economic Behavior & Organization*, 2015, *110*, 45–58.

Braun, Phillip A, George M Constantinides, and Wayne E Ferson, "Time nonseparability in aggregate consumption: International evidence," *European Economic Review*, 1993, *37* (5), 897–920.

Byrne, Molly, Emily Bray, Evan Maclean, and Angie Johnston, "Evidence for win-stay-lose-shift in puppies and adult dogs.," in "CogSci" 2020.

Cai, Xiqian, Jie Gong, Yi Lu, and Songfa Zhong, "Recover overnight? work interruption and worker productivity," *Management Science*, 2018, *64* (8), 3489–3500.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor supply of New York City cabdrivers: One day at a time," *The Quarterly Journal of Economics*, 1997, *112* (2), 407–441.

Card, David and Gordon B Dahl, "Family violence and football: The effect of unexpected emotional cues on violent behavior," *The quarterly journal of economics*, 2011, *126* (1), 103–143.

Cerulli-Harms, Annette, Lorenz Goette, and Charles Sprenger, "Randomizing Endowments: An Experimental Study of Rational Expectations and Reference-Dependent Preferences," *American Economic Journal: Microeconomics*, 2019, *11* (1), 185–207.

Chen, Yan and Fang-Fang Tang, "Learning and incentive-compatible mechanisms for public goods provision: An experimental study," *Journal of Political Economy*, 1998, *106* (3), 633–662.

Conrad, Keith, "Probability distributions and maximum entropy," *Entropy*, 2004, *6* (452), 10.

Cotla, Chenna Reddy, "Learning in Repeated Public Goods Games-A Meta Analysis," *Available at SSRN 3241779*, 2015.

Crawford, Vincent P and Juanjuan Meng, "New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income," *American Economic Review*, 2011, *101* (5), 1912–32.

DellaVigna, Stefano, "Structural behavioral economics," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 613–723.

Dragone, Davide and Nicolas R Ziebarth, "Non-separable time preferences, novelty consumption and body weight: Theory and evidence from the East German transition to capitalism," *Journal of health economics*, 2017, *51*, 41–65.

Dreber, Anna, Ch Gerdes, Patrik Gränsmark, and AC Little, "Facial masculinity predicts risk and time preferences in expert chess players," *Applied Economics Letters*, 2013, *20* (16), 1477–1480.

_ , Christer Gerdes, and Patrik Gränsmark, "Beauty queens and battling knights: Risk taking and attractiveness in chess," *Journal of Economic Behavior & Organization*, 2013, *90*, 1–18.

Erev, Ido and Alvin E Roth, "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, 1998, pp. 848–881.

Farber, Henry S, "Is tomorrow another day? The labor supply of New York City cab-drivers," *Journal of political Economy*, 2005, *113* (1), 46–82.

_ , "Reference-dependent preferences and labor supply: The case of New York City taxi drivers," *American Economic Review*, 2008, *98* (3), 1069–82.

_ , "Why you can't find a taxi in the rain and other labor supply lessons from cab drivers," *The Quarterly Journal of Economics*, 2015, *130* (4), 1975–2026.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz, "Frictions in a competitive, regulated market: Evidence from taxis," *American Economic Review*, 2019, *109* (8), 2954–92.

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

Gränsmark, Patrik, "Masters of our time: Impatience and self-control in high-level chess games," *Journal of Economic Behavior & Organization*, 2012, *82* (1), 179–191.

Haenni, Simon, "Ever tried. Ever failed. No matter? On the demotivational effect of losing in repeated competitions," *Games and Economic Behavior*, 2019, *115*, 346–362.

Haruvy, Ernan and Dale O Stahl, "Between-game rule learning in dissimilar symmetric normal-form games," *Games and Economic Behavior*, 2012, *74* (1), 208–221.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Linnemer, Laurent and Michael Visser, "Self-selection in tournaments: The case of chess players," *Journal of Economic Behavior & Organization*, 2016, *126*, 213–234.

Mas, Alexandre, "Pay, reference points, and police performance," *The Quarterly Journal of Economics*, 2006, *121* (3), 783–821.

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

Morgul, Ender Faruk and Kaan Ozbay, "Revisiting labor supply of new york city taxi drivers: Empirical evidence from large-scale taxi data," in "Transportation Research Board 94th Annual Meeting," Vol. 15 2015.

Pope, Devin G and Maurice E Schweitzer, "Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes," *American Economic Review*, 2011, *101* (1), 129–57.

Posch, Martin, "Win–stay, lose–shift strategies for repeated games—memory length, aspiration levels and noise," *Journal of theoretical biology*, 1999, *198* (2), 183–195.

Post, Thierry, Martijn J Van den Assem, Guido Baltussen, and Richard H Thaler, "Deal or no deal? decision making under risk in a large-payoff game show," *American Economic Review*, 2008, *98* (1), 38–71.

Roth, Alvin E and Ido Erev, "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and economic behavior*, 1995, *8* (1), 164–212.

Sousa, José De and Muriel Niederle, "Trickle-Down Effects of Affirmative Action: A Case Study in France," Technical Report, National Bureau of Economic Research 2022.

_ , Guillaume Hollard et al., "From Micro to Macro Gender Differences: Evidence from Field Tournaments," Technical Report, Sciences Po 2021.

Tamura, Kohei and Naoki Masuda, "Win-stay lose-shift strategy in formation changes in football," *EPJ Data Science*, 2015, *4*, 1–19.

Thakral, Neil and Linh T Tô, "Daily labor supply and adaptive reference points," *American Economic Review*, 2021, *111* (8), 2417–2443.

Turnovsky, Stephen J and Goncalo Monteiro, "Consumption externalities, production externalities, and efficient capital accumulation under time non-separable preferences," *European Economic Review*, 2007, *51* (2), 479–504.

For Online Publication

A Identification proofs

The following parametric assumption is made on the distribution of the outside option, $F(c)$,

Assumption 1 *$F(c)$ is an exponential distribution with parameter λ.*

We now argue that under the Assumption 1 and by normalizing one parameter of our choice in the model, we can identify δ, l_θ and $u(\cdot)$. Let,

$$H(\theta, y) = u(y) + \delta \sum_{\substack{y',\chi' \in \\ Y \times \{0,1\}}} p(y', \chi'|y) V(\theta, y', \chi'). \tag{6}$$

Under the Assumption 1 and from equation 4, the probability of stopping and taking outside option, $h(\theta, y, \chi)$, can be written as,

$$h(\theta, y, \chi) = e^{-\lambda(H(\theta,y)+(1-\chi)l_\theta)} \tag{7}$$

Claim 2 *$\lambda H(\theta, y)$ and λl_θ are identified for all (θ, y).*

Proof. Let us look at equation (7) evaluated at $\chi = 1$. The LHS, $h(\theta, y, 1)$, can be directly calculated from the data as probability of stopping after a win. In the RHS, the second term in the power, $(1 - \chi)l_\theta$, is 0. Therefore, we can recover/identify $\lambda H(\theta, y)$ from (7).

Next, let us look at equation (7) evaluated at $\chi = 0$. The LHS can be calculated from the data as probability of stopping after a loss. In the RHS, $H(\theta, y)$ term was identified in the first part of the proof. Thus, λl_θ is identified from (7) as well. ■

Claim 3 *$\lambda V(\theta, y, \chi)$ are identified for all (θ, y, χ).*

Proof. With a little abuse of notation let us denote $\bar{c}(\theta, y, \chi)$ by $\bar{c}$. We can rewrite (2) as,

$$V(\theta, y, \chi, c) = \mathbf{1}(c > \bar{c}) * c + \mathbf{1}(c \leq \bar{c})\Big(H(\theta, y) + (1 - \chi)l_\theta\Big), \tag{8}$$

where $\mathbf{1}(\cdot)$ is an indicator function. Taking expectations of both hand sides of (8) with respect to c gives,

$$\begin{aligned}
V(\theta, y, \chi) &= E(c|c > \bar{c}) + Pr(c \leq \bar{c})\Big(H(\theta, y) + (1-\chi)l_\theta\Big) \\
&= Pr(c > \bar{c})\Big(E(c) + \bar{c}\Big) + Pr(c \leq \bar{c})\Big(H(\theta, y) + (1-\chi)l_\theta\Big) \\
&= Pr(c > \bar{c})\Big(\frac{1}{\lambda} + \bar{c}\Big) + Pr(c \leq \bar{c})\Big(H(\theta, y) + (1-\chi)l_\theta\Big) \\
&= Pr(c > \bar{c})\Big(\frac{1}{\lambda} + \bar{c} - H(\theta, y) - (1-\chi)l_\theta\Big) + H(\theta, y) + (1-\chi)l_\theta. \qquad (9)
\end{aligned}$$

Multiplying both hand sides by λ and substituting $\bar{c}(\theta, y, \chi)$ from expression (4) we get,

$$\lambda V(\theta, y, \chi) = e^{-\lambda[H(\theta,y)+(1-\chi)l_\theta]} + \lambda H(\theta, y) + (1-\chi)\lambda l_\theta. \qquad (10)$$

Claim 2 and expression (10) imply that $\lambda V(\theta, y, \chi)$ are identified for all (θ, y, χ). ■

Claim 4 *δ and $\lambda u(y)$ are identified.*

Proof. We can consider the difference $\lambda(H(\theta, y) - H(\theta', y))$ for some $\theta \neq \theta'$. This gives us,

$$\delta = \frac{\lambda(H(\theta, y) - H(\theta', y))}{\lambda(\sum_{y',\chi' \in Y \times \{0,1\}} p(y', \chi'|y)(V(\theta, y', \chi') - V(\theta', y', \chi')))}.$$

By claims 2 and 3, numerator and denominator are identified in the above equation. Finally, we can identify $\lambda u(y)$ from,

$$\lambda u(y) = \lambda H(\theta, y) - \lambda\delta \sum_{\substack{y',\chi' \in \\ Y \times \{0,1\}}} p(y', \chi'|y)V(\theta, y', \chi')$$

■

Finally, we can identify all the parameters and value functions by normalizing λ.[28] This completes the identification of the parameters of the model.

[28] In the structural estimation, we normalize $\lambda = 1$.

B Further predictions

Following a similar intuition as in Predictions 1 and 2, we derive additional predictions. We expect the correlation between winning percentages for the last-games and the middle-games to be negative; between winning percentages for the middle-games and the only-games to be negative; between winning percentages for the first-games and the only-games to be negative and finally, between winning percentages for the middle-games and the first-games to be positive. Figure 6 presents the correlation matrix with p-values in parentheses and the data support all the predictions.

Figure 6: Correlation Matrix

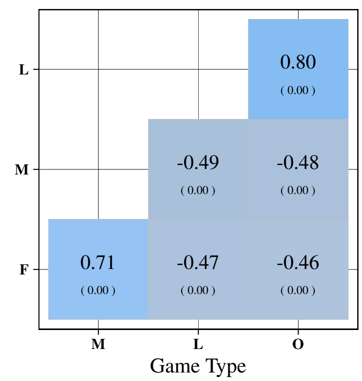

C Consistency of behavioral types

In this section, we discuss the consistency of the two behavioral classifications. In this paper we propose two distinct definitions of behavioral types (Definitions 1 and Definition 2). The two classifications are intuitively related, but do not necessarily overlap. That is, given some data, a user can be classified as a win-stopper according to Definition 1, but be identified as a loss-stopper by the model (Definition 2 based on Proposition 1). For example, consider a user whose complete playing history consists of the following set of three sessions:

$$\{WWWW, WLW, WLL\}.$$

According to Definition 1, the user is classified as a win-stopper because she won last-games more often than middle-games. For this user, the stopping probability after a loss is $\Pr(Stop|Loss) = 1/3$ and the stopping probability after a win is $\Pr(Stop|Win) = 2/7$.

Given that the probability of stopping is higher after a loss than after a win, $\Pr(Stop|Loss) > \Pr(Stop|Win)$, our model would identify the user as a loss-stopper. The fact that one definition does not necessarily imply the other strengthens any relationship we find between the two classifications, thus highlighting the consistency between our intuition and the proposed theory. Let us compare the two classifications.

For 84.6% of the users the two classifications match. This result provides strong evidence that the model captures users' behavior and that the game outcome affects the utility of the next game. Figure 7 presents a transition matrix for model types and behavioral types. We observe a large mass on the diagonal, indicating that the two classifications are fairly consistent. For example, 91% of win-stoppers identified by the model were also identified as win-stoppers under Definition 1. However, there are some mismatches; for example, some neutral types by Definition 1 are classified as behavioral types by the model and vice versa. Notably, cases in which a win-stopper (loss-stopper) under Definition 1 is identified as a loss-stopper (win-stopper) by the model are rare, occurring only about 1% of the time.

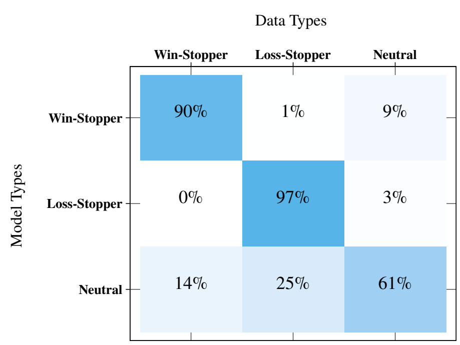

Figure 7: Model vs. Data Types

More generally, we can derive mathematical conditions when the two classifications coincide. We can modify Definition 1 to include first- and only-games and compare $\Pr(Win|Stop)$ to $\Pr(Win|Continue)$. The former corresponds to the fraction of wins in all the games after which the person stopped (it could be the last-games or only-games). The latter corresponds to the fraction of wins in all the games after which the player played at least one more game (it could be first-games or middle-games). Similarly, we can rewrite Definition 2 to compare $\Pr(Stop|Win)$ with $\Pr(Stop|Loss)$. Using conditional probability rules, one of the ways to represent the condition when two classifications coincide is, if the following holds (depending on a reader, one might rearrange terms in the conditional probability formula to make them more intuitive to understand):

$$\frac{\Pr(Continue|Win)}{\Pr(Stop|Loss)} = \frac{Pr(Continue)}{\Pr(Stop)} \tag{11}$$

Fixing the right-hand side of equation 11, if the win affects positively to continuation probability, then loss should affect positively to stop in order for the two classifications to coincide. Similarly, if the win negatively affects continuing, then the loss should also negatively affect stopping. Roughly speaking, we will have the same classifications by two definitions if both win and loss symmetrically affect the stopping decision. In other words, if the stopping decision is affected, for example, only by a win and is random after a loss, two definitions might give different classifications.

D Robustness

D.1 Changing session definition—varying break time

In the main body of the paper, while defining a session, we set the break time T to 30 minutes. To ensure that the results on the behavioral types are not sensitive to the choice of T, we classify users into types using sessions defined by break times $T \in \{5, 15, 30, 60\}$. We are interested in how the behavioral type classification changes and the transition between the different Ts. Figure 8 presents transition matrices.

In Figure 8 we see a large mass on the diagonal, which implies that the classifications mostly match. However, there are some mismatches; for example, some neutral types with $T = 5$ are classified as behavioral types with $T = 15$. What is noteworthy in the panel are the transitions between behavioral types: 0% of users are classified as a win-stopper (loss-stopper) by one classification and a loss-stopper (win-stopper) by another classification or vice versa. There are no switches in behavioral types as we vary $T \in \{5, 15, 30, 60\}$.

D.2 Robustness to tolerance thresholds

Unless we stated otherwise, throughout the paper, we set tolerance levels τ and κ to 7%. We vary τ and κ from 0 to .2 and we classify our users into types according to Definition 1 and model identification. Figure 9 presents the users' population decomposition by types with $\tau \in [0, .2]$ and $\kappa \in [0, .2]$. While there is movement in a predicted direction—higher the threshold, less behavioral types—we see that types are overall robust to changing the allowed tolerance (no abrupt, unexpected discontinuities).

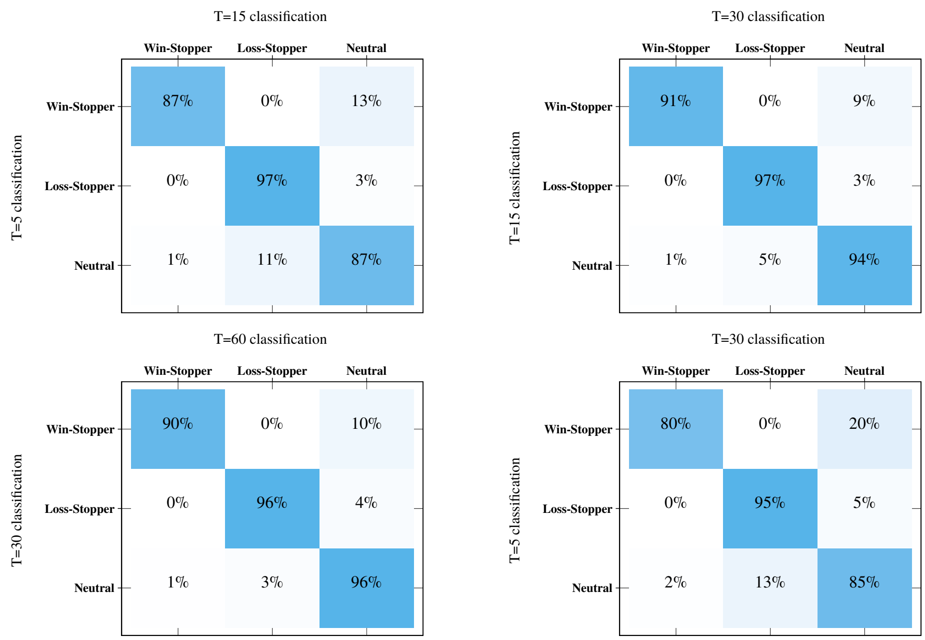

Figure 8: Transitions between 5 to 15, 15 to 30, 30 to 60, and 5 to 30

D.3 Full data structural estimation results

To show that estimation does not depend on the grid size or data range, we change both and compare the results. We divide the rating range into grids of 20 since the rating has a wide range ($[100, 2798]$). We have few observations where the rating is below 600 or above 2000. Consequently, we place all the users with a rating below 600 in the first rating grid and those above 2000 in the last rating grid (grid 71). We divide the rest of the rating range into 20 point intervals.

The main parameters that we focus on are l_θ for $\theta \in \{\theta_W, \theta_L, \theta_N\}$. The estimates are presented in Table 4. We bootstrapped 300 times to find standard deviation of the parameters. Table 4 shows that parameter estimates as well as their standard deviations are similar to the ones in Table 2 in Section 5.2. Hence, parameter estimates are stable with respect to rating range and the grid size.

Parameter	Mean	SD
l_{θ_W}	0.665	0.004
l_{θ_N}	-0.017	0.002
l_{θ_L}	-0.604	0.002

Table 4: Bootstrapped values for l_θ

(a) Behavioral Type Fractions (b) Model Type Fractions

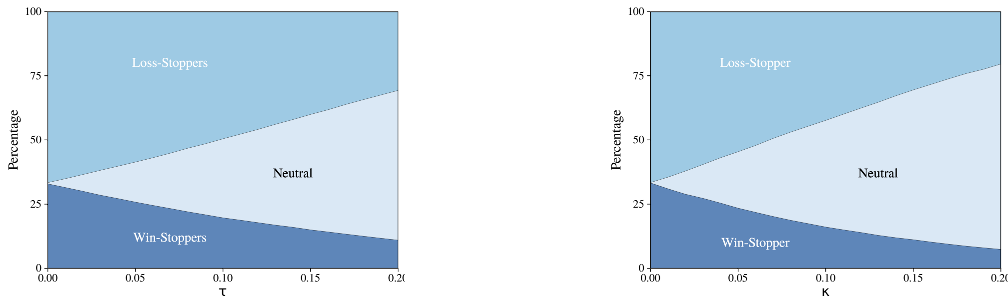

Figure 9: Type Decomposition

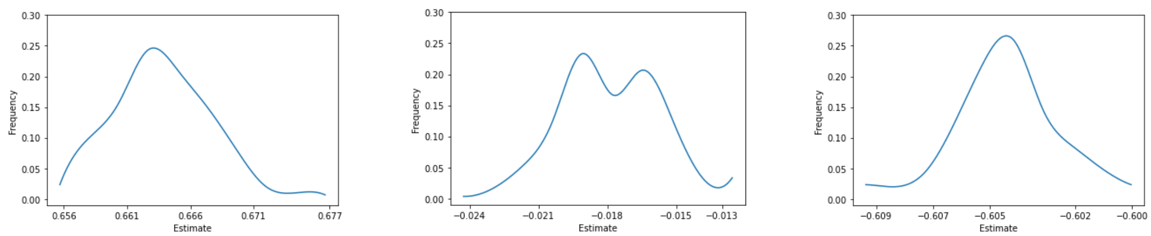

(a) l_{θ_W} (b) l_{θ_N} (c) l_{θ_L}

Figure 10: Distribution of bootstrap values for l_θ for full data.

D.4 Robustness with respect to behavioral type tolerance level

From the definition of behavioral types, it is clear that estimates of l_θ for $\theta \in \{\theta_W, \theta_L, \theta_N\}$ depend on behavioral type tolerance level. In the main text all our results are for $\kappa = 0.07$. In this section we present estimation results for two other values $\kappa = 0.05$ and $\kappa = 0.09$.

Parameter	Mean	SD
l_{θ_W}	0.622	0.004
l_{θ_N}	-0.005	0.003
l_{θ_L}	-0.580	0.002

(a) $\tau = 0.05$

Parameter	Mean	SD
l_{θ_W}	0.737	0.005
l_{θ_N}	-0.027	0.003
l_{θ_L}	-0.642	0.003

(b) $\tau = 0.09$

Table 5: Bootstrapped values for l_θ

Table 5 shows that parameter estimates changes in the expected direction. For example, when we relax non-behavioral (neutral type) constraint from $\kappa = 0.07$ to $\kappa = 0.09$, there are less behavioral types. Therefore, the users who are still behavioral types with $\kappa = 0.09$, are the ones who are "more behavioral" then the one with $\kappa = 0.07$. This implies that

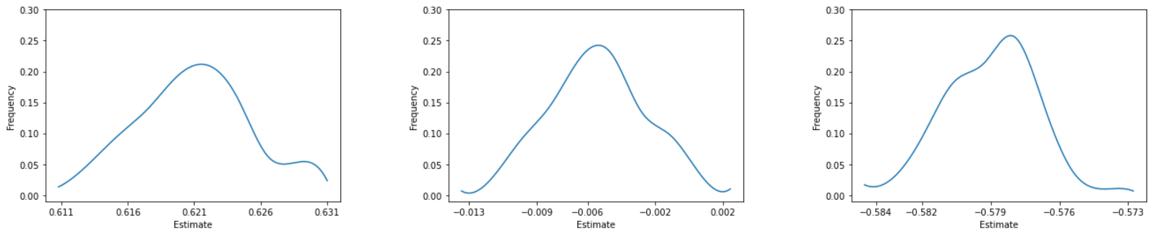

(a) l_{θ_W} (b) l_{θ_N} (c) l_{θ_L}

Figure 11: Distribution of bootstrap values for l_θ ($\tau = 0.05$).

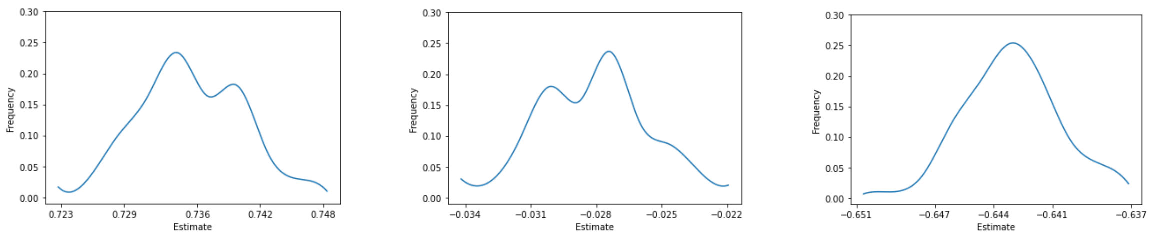

(a) l_{θ_W} (b) l_{θ_N} (c) l_{θ_L}

Figure 12: Distribution of bootstrap values for l_θ τ=0.09.

the effect from the last game result (whether negative or positive) is stronger for those behavioral types. This comparative static is met in our estimates. Win-stoppers' parameter l_{θ_W} is lower for $\kappa = 0.05$ and higher for $\kappa = 0.09$ compared to $\kappa = 0.07$. Similarly, for loss-stoppers the absolute value of l_{θ_K} is lower for $\kappa = 0.05$ and higher for $\kappa = 0.09$ compared to $\kappa = 0.07$.

D.5 Time consistency for more active users

Let us look at the level of time consistency between the years 2017 and 2018, as we look at players with at least 150, 300, 450, and 600 games in both years, 2017 and 2018. As we remove users with less than 150, 300, 450, and 600 games in both years, the sample is reduced by 7%, 11%, 12%, and 11%, respectively. The matching between the years 2017 and 2018 increases as we look at more active users. Specifically, as we look at users with at least 150, 300, 450, and 600 games in both years, the matching is 78%, 79%, 80%, and 81%, respectively. The transition between the behavioral types is in Figure 13.

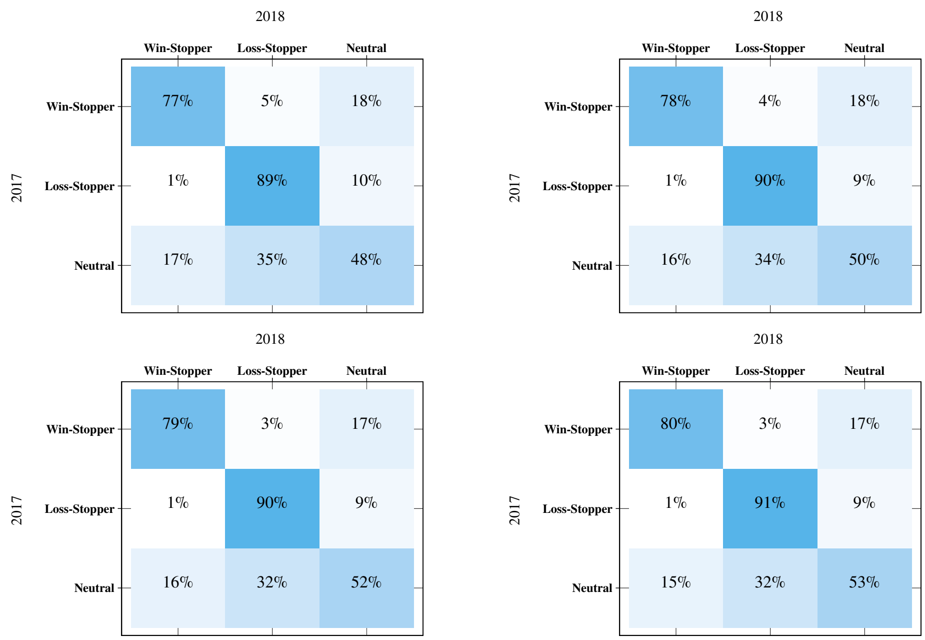

Figure 13: Sub-samples of users with at least 150, 300, 450, and 600 games each year

E Possible crowding-out effects

Our counterfactual analysis reveals that considering users' behavioral type for the matching algorithm can increase the average number of games played during a session. One might think of several crowding-out effects that an increase of a session length might have. Without a randomized controlled trial (RCT) we can not fully address such concerns; however, we provide evidence that some of these effects are not likely.

E.1 More games lead to more time on the platform

The goal of the counterfactual analysis is to increase the number of games during a session, but the market designer's goal could also be to increase the time spent on the platform. We calculated the correlation for every individual between minutes spend on the platform during a session and the number of games played in the same session. Figure 14a shows that correlation between these two variables is high. The median correlation between minutes and games during the session is 0.98 across users.

E.2 Asymmetric matching does not decrease playing time

Another issue that one might worry about is that asymmetric matching can cause fast games, in the sense that stronger users can win faster playing against weaker users. To

Figure 14: Correlations

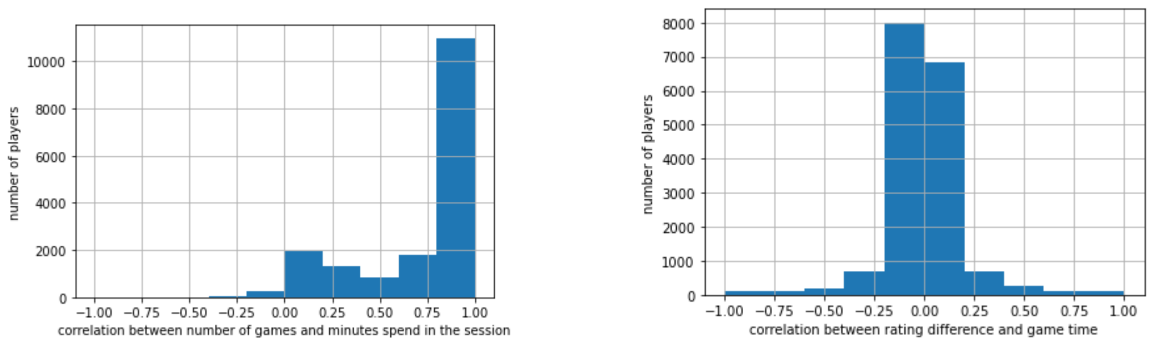

(a) Minutes spent on a session and number of games

(b) Minutes spent on a game and rating difference

show that this is not likely to be an issue, we calculate the correlation between rating differences and minutes spent on a game. Rating difference provides a measure of how much better one user is compared to another. Figure 14b shows a correlation between how much better an opponent is and how much time the game lasts is close to zero for most cases.

E.3 One long session can cause the next session to be short

One might worry that if a session time increases during the day, it can decrease the next session length (if users set out a certain amount of time to spend on the platform every day). We find that the correlation between the number of sessions played during a day, and the average length of a session is 0.0002. We also find that the session length of the previous session does not have any explanatory power on the length of the next session.

E.4 Users adjust game type based on the time they have played

The last issue that we discuss here is changing the type of the game. A person who started a session with a 5-minute blitz game can play a shorter last game (for example 3-minute game) because she has only a certain time allocated on the platform. If that is the case, we should see that people change game types during the session. We find that 96% of sessions are homogeneous in the sense of the game type. This homogeneity captures not only a change of game type in the last game but during any other time. This makes our argument even stronger that users do not choose the last game type based on the remaining time allocated for playing chess that day.

F Additional CPH model analysis

To demonstrate the stability of the Outcome variable as a main contributor to the stopping decision we conducted CPH analysis by adding all the variables sequentially. For readability, we separated Win-Stoppers and Loss-Stoppers to avoid an extra layer of interactions with behavioral types. The table 6 confirms the main findings from Figure 5. The outcome variable is stable and does not change substantially across different models neither for Win-Stoppers nor for Loss-Stoppers.

	Win-Stoppers						Loss-Stoppers					
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Outcome	0.54***	0.54***	0.56***	0.56***	0.56***	0.56***	-0.65***	-0.65***	-0.66***	-0.65***	-0.65***	-0.65***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Rating Change		-0.01***	-0.01***	-0.01***	-0.01***	-0.01***		0.01***	0.01***	0.01***	0.01***	0.01***
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Opponent Rating Diff			-0.03***	-0.03***	-0.03***	0.00***			0.00***	0.00	0.00	0.03***
			(0.00)	(0.00)	(0.00)	(0.00)			(0.00)	(0.00)	(0.00)	(0.00)
Previous Outcome				-0.02***	-0.02***	-0.02***				0.06***	0.06***	0.06***
				(0.00)	(0.00)	(0.00)				(0.00)	(0.00)	(0.00)
Rating change × Outcome					0.01***	0.01***					-0.00	-0.00*
					(0.00)	(0.00)					(0.00)	(0.00)
Opp. Rating Diff. × Outcome						-0.05***						-0.06***
						(0.00)						(0.00)
N	8,175,464						16,463,457					

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 6: Sequentially adding variables to CPH model for Win-stoppers and Loss-stoppers.

G Additional tables and figures

G.1 Winning percentage and game type

Types	O	F	L	M	All
Loss-Stopper	38.5	56.7	35.9	59.2	50.4
Neutral	50.6	50.0	50.2	50.6	50.5
Win-Stopper	64.3	43.7	65.6	43.6	50.7

(a) Winning percentage

Types	O	F	L	M	All
Loss-Stopper	10.93	8.36	8.85	7.59	3.53
Neutral	9.53	6.86	5.70	5.15	4.21
Win-Stopper	13.79	9.05	10.53	7.59	4.15

(b) Winning percentage st. deviation

Table 7

G.2 Distribution of bootstrap values for l_θ

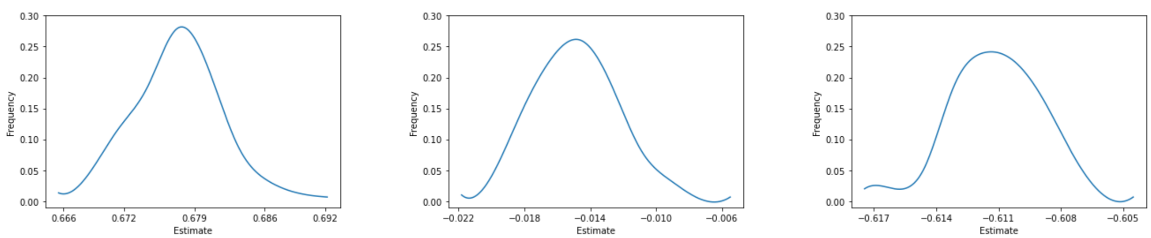

(a) l_{θ_W} (b) l_{θ_N} (c) l_{θ_L}

Figure 15: Distribution of bootstrap values for l_θ

H The effect of the opponents' rating and strength

Let us only look at games in which the difference between the own and opponent's ratings is less than 200, 100, 50, and 10. Additionally, let us include unrestricted data for comparison. Table 8 presents the results. Table 8 suggests that the fraction of behavioral types and the composition of behavioral types (fraction of win-stoppers compared to loss-stoppers) are fairly unaffected as we restrict the data. It is worth noting that we drop a significant fraction of the data when restricting the difference to less than 10 rating points. Yet, the results on behavioral types and their ratios stay consistent.

Finally, let us evaluate the effect of interaction between the opponent's strength and the stopping decision. We calculate stopping probability after a win and after a loss based on whether the user faced a stronger or weaker opponent. The results are presented in Table 9. The users are marginally more likely to stop a session after winning against a stronger

Fraction	NR	< 200	<100	<50	<10
Behavioral types	78.5	78.9	78.9	78.9	79.9
Loss-stoppers	55.1	55.4	55.2	55.3	55.6
Win-stoppers	23.4	23.5	23.6	23.6	24.3
Neutrals	21.5	21.1	21.1	21.1	20.1
Reduction in obs	NA	2%	12%	30%	77%

Table 8: Restricted data by rating difference

player. However, the differences based on the opponent's rating appear to be fairly small. Note that the overall probability of stopping is higher after a loss than after a win. This is due to the fact that there are more loss-stoppers than win-stoppers.

	Loss	Win
Stronger	0.294	0.241
Weaker	0.296	0.236

Table 9: Conditional stopping probabilities

I Rating and behavioral type

In this section, we examine whether a user's type is correlated with the type classification. In particular, whether a user with a higher rating is more likely to be one type. To study this question, we restrict our data, focusing on users with an average rating of over 1,800 (619 users). Among these users with the best rating, 81.9% are behavioral types. The ratio of win-stoppers among the behavioral types is 32.7%, similar to the full data ratio of 29.8%. A hypothesis that we may observe more of one of the behavioral types among higher rated users, for example, more of win-stoppers, has little support in our data.

J Practice improves rating

Here we ask, does practice make better? That is, we examine whether playing more games leads to higher ratings. Table 10 presents the results of regression summaries for different specifications and models. We find that playing more games results in a higher rating, implying the possible positive effect of increasing the number of games played.

	Rating			
	OLS	OLS	FE	FE
	(1)	(2)	(3)	(4)
Intercept	-0.000	-0.026**	-0.000	-0.043***
	(0.008)	(0.011)	(0.009)	(0.012)
# of Games	0.300***	0.308***	0.234***	0.251***
	(0.009)	(0.013)	(0.008)	(0.011)
Win-Stopper		0.067***		0.125***
		(0.020)		(0.022)
# of Games and Win-Stopper		0.048**		0.007
		(0.024)		(0.019)
Observations	14,788	10,472	117,070	82,770
R^2	0.090	0.091	0.055	0.058
Adjusted R^2	0.090	0.091		
Residual Std. Error	0.954(df = 14786)	0.946(df = 10468)		
F Statistic	1083.984*** (df = 1.0; 14786.0)	286.322*** (df = 3.0; 10468.0)	6806.8*** (df = 1.0; 117068.0)	1701.2*** (df = 3.0; 82766.0)
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10: Columns (1) and (3) use Ordinary Least Square and Fixed Effects models, respectively, to study if the number of games played and rating are connected. Similarly, columns (2) and (4) study if there is an asymmetry between win- and loss-Stoppers in improving their rating. Variables, # of Games and Rating (dependant variable) are standardized. Variable win-stopper takes values 1 or 0 if the observation belongs to a win or loss-stopper, respectively.

K World maps

The data includes users from 191 different countries,[29] however, there are certain countries with too few users and we exclude those countries. In particular, we remove the countries with less than 30 users and we end up with 65 countries with at least 30 users. Note, almost 50% of the users in our data indicate their country to be one of the following 5 countries: USA, India, Russia, Canada, or Norway. Figures 16, 17, 18 include three world heat maps showing number of users, average rating and the fraction of win-stoppers among behavioral types.

The fraction of behavioral types varies by country and ranges between 67.3% (Finland) and 92.3% (Vietnam). Furthermore, the ratio of win-stoppers among behavioral types varies considerably by country ([10.8, 44.1]). For example, win-stoppers make up only 10.8% of all behavioral types in Kazakhstan, while in Japan, they are 44.1% of all behavioral types. Future work could study the differences and possible underlying factors leading to heterogeneity documented above.

[29] It is important to emphasize, however, that country variable is self-reported by the users, they can choose any country they wish and they can also change it afterwards.

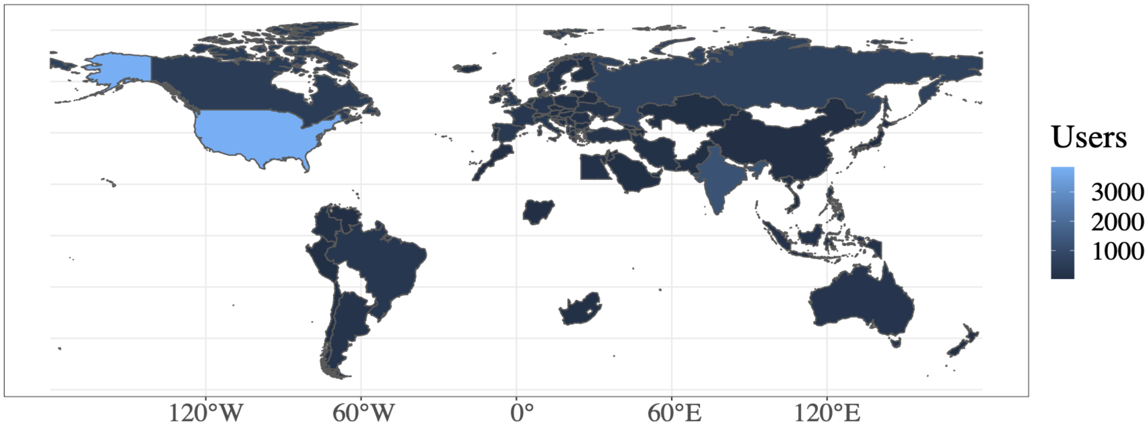

Figure 16: Number of users from a country

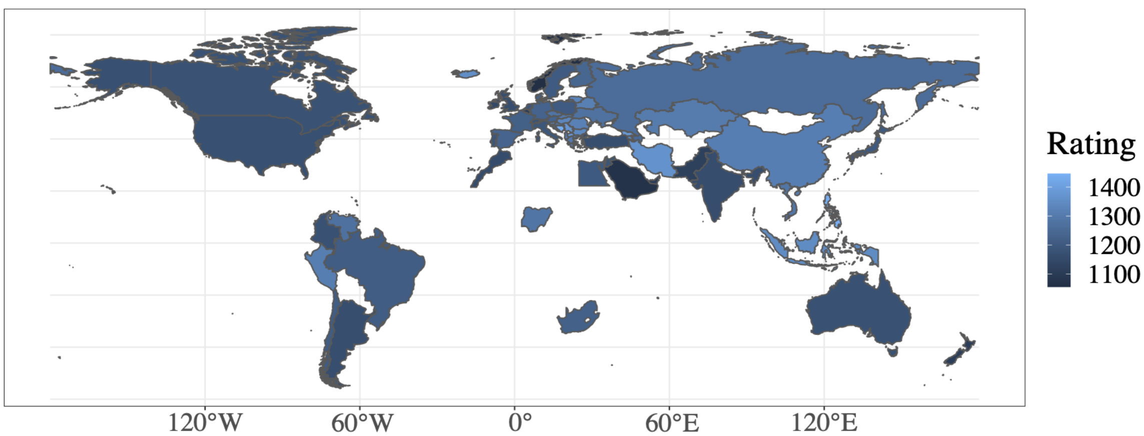

Figure 17: Average rating in a country

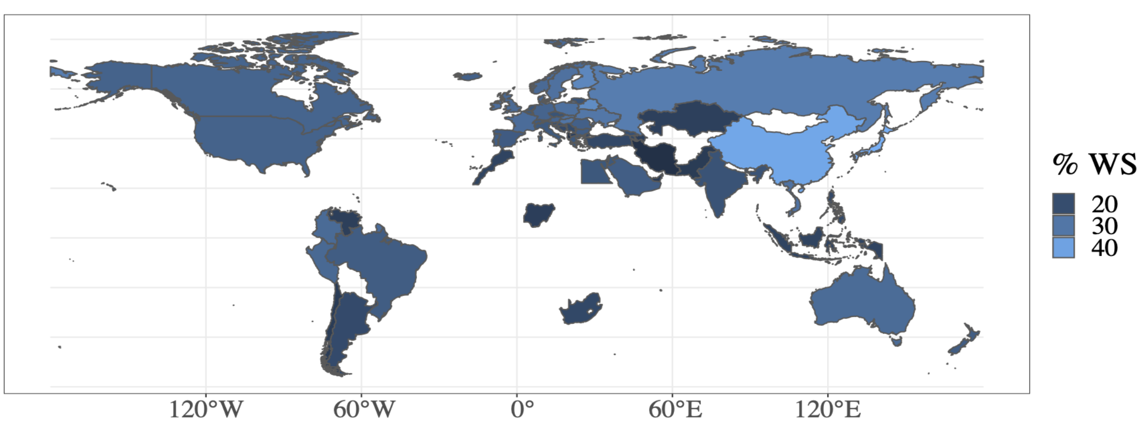

Figure 18: Fraction of win-stoppers in a country