# First Impressions and Technology Adoption: Separating Primacy Effects from Selection

Miguel Ángel Talamas Marcos *

February 24, 2026

**Abstract**

First impressions are widely believed to exert persistent effects on behavior — a primacy effect. However, causally identifying the primacy effect is challenging due to selection: individuals with higher returns are more likely to adopt and to have a positive first experience. Leveraging historical playing data from hundreds of thousands of online chess players, this paper disentangles the causal primacy effect from selection. The paper studies the adoption of "berserk"—a feature in online chess tournaments that increases the payoff of winning a match, but also the difficulty. Using a random forest, I identify players for whom winning the first time they experiment with berserk is quasi-random. While descriptive evidence suggests large and permanent primacy effects, quasi-experimental estimates reveal a primacy effect that fades within fifty games. Consistently, leveraging data from the same player on faster chess matches allows for the construction of a measure of returns to adoption, and controlling for returns to adoption in the full sample of players also renders a short-lived primacy effect. These findings demonstrate that conflating selection with the outcome of the first experimentation can substantially overstate the persistence of first impressions.

*JEL* D83, O33, O10, O30

*Talamas Marcos: Inter-American Development Bank, miguelta@iadb.org.

I thank Henrique Brasiliense, Matias Busso, Julian Cristia, Joshua Dean, Andrea de la Garza, Daniel Martin, David Salama, Yuval Salant, Carlos Scartascini, Jörg Spenkuch, Carolina Tojal, Diego Vera, and Chris Woodruff for helpful conversations and suggestions. Excellent research assistance by Maria Emilia Bullano is gratefully acknowledged.

Why do initial experiences with new technologies strongly predict long-term adoption patterns? A leading explanation is primacy effects: positive or negative first impressions create path dependencies that shape subsequent usage (Asch, 1946; Rabin and Schrag, 1999). However, a fundamental identification challenge limits causal interpretation. Individuals who experience positive outcomes during initial experimentation may adopt at higher rates either because the positive first impression causally increased their propensity to adopt—a true primacy effect—or because they possess higher returns to the technology, which generated both the favorable first experience and continued use. Distinguishing between these two mechanisms—primacy effects and selection—is the central challenge this paper addresses.

This distinction matters for understanding technology diffusion and designing policies to promote adoption, particularly for technologies whose returns depend on unobserved individual characteristics. When returns are driven by such unobserved idiosyncratic factors, individuals cannot extrapolate from others' experiences, making personal experimentation the primary source of information (Munshi, 2004; Foster and Rosenzweig, 2010). Establishing whether initial experiences causally drive adoption or merely proxy for underlying returns is therefore essential for understanding what determines long-run technology use.

Disentangling primacy effects from selection involves overcoming four challenges. First, it requires exogenous variation in the outcomes of early experimentation. Second, the analysis must focus on a technology whose returns are initially unknown to the experimenter and are influenced by unobservable idiosyncratic factors, thereby limiting social learning. Third, ideally, individual returns to adoption would be known regardless of whether an individual adopts. Fourth, the setting should provide sufficient statistical power by allowing observation of many individuals before and after the first impression, with adoption tracked over extended periods.

This paper addresses these four challenges by employing a novel quasi-experimental design in a highly rich context. The empirical analysis draws on millions of online chess

tournaments played by hundreds of thousands of players on Lichess.org, the second-largest chess platform. The technology in question is "berserk", a voluntary feature that halves the player's thinking time in exchange for earning an extra tournament point if they win the game. Players choose to go berserk after being paired with an opponent and observing their Elo rating, a precise measure of ability, but before the match starts.

While online chess tournaments may seem a highly specific setting, they provide a uniquely advantageous context for overcoming the identification challenges at hand, making the setting a strength rather than a limitation (Palacios-Huerta, 2025; List, 2020). The setting offers exceptionally high statistical power, with hundreds of thousands of players and many games per individual, both before and after their first experimentation. Moreover, the individual returns to the "berserk" option—the technology under study—are likely initially unknown to players, as the feature is exclusive to Lichess.org tournaments and unavailable on Chess.com, the other major online chess platform. Additionally, online chess has several time-control variations (how much time players have to think), allowing the use of players' skill in faster games as a measure of the returns to adopting berserk. Berserk thus represents an ideal setting for studying primacy effects.

The online chess setting and the literature on experimentation and technology adoption face the same challenge of disentangling the true primacy effect from selection. This paper addresses this challenge by focusing on players for whom winning or losing the first berserk is as good as random. To do so, it uses a random forest to predict the probability that each player will win their first berserk based on observable player characteristics and history. As expected, players with a predicted high (low) probability of winning are indeed more (less) likely to win. However, within the neighborhood of a predicted probability of winning of 0.5 (0.485-0.515), the empirical strategy section shows that the actual win rates are not statistically different from 50%, i.e., these players were as likely to win as they were to lose their first berserk.

The identifying assumption is that within this narrow threshold, the remaining shocks that determine the outcome of the match (win vs lose) are not systematically related to returns to adoption. While being around the 0.5 prediction of the random forest guarantees that players are similar in a weighted combination of their observables, it does not necessarily guarantee that returns to adoption are orthogonal to a positive first impression. To ameliorate this concern, I show that while the constructed measure of returns to adoption based on the player's skill in faster modalities is a significant driver of the outcome of the first berserk in general, it has no relationship with the outcome of the first berserk in the window close to the 0.5 prediction of the random forest. This identification strategy has some parallels to the assumptions used in close elections, where the winner is not random in general but is essentially believed to be random in close contests (Lee et al.; Lee, 2008).

Descriptive evidence indicates that those who had a positive first impression (winners) in their first berserk are 5.6 percentage points (22%) more likely to use the feature in the future, even after 100 games. Moreover, the better the first impression was, the more likely they are to use the feature in the future: those who won with more than 5% of time left on the clock have the highest adoption, followed by those who won but with less than 5% of time remaining, then those who lost due to a checkmate, and finally, those who had the worst experience in their first berserk — losing by running out of time. However, this pattern that could be interpreted as consistent with a primacy effect, could also be driven by selection: those with higher returns (who are better at playing fast chess) are more likely to win their first berserk and are also more likely to go berserk in the future.

The causal estimates of the primacy effect based on the sample of players for whom winning or losing was as good as random yield a different conclusion. There is indeed a primacy effect: players with a positive first impression are 5 percentage points (25%) more likely to go berserk in their first 30 games. However, this effect begins to fade within the first 10 games, and by the 50th game after the first berserk, it is entirely

diminished. This evidence indicates that the primacy effect observed descriptively persists due to selection, but it is not a true persistent primacy effect.

Leveraging the fact that players compete under different time controls (the time available to think per game) allows for the construction of a measure of returns to adoption based on how well a player performs in faster games. In particular, I study the adoption of berserk in Blitz tournaments, where players have between 3 and 8 minutes, and I use their elo rating in Bullet tournaments (between 30 seconds and 3 minutes) as a measure of returns to adoption. The underlying assumption is that players who are better at faster modalities will have higher returns on a technology that halves their time to move. Estimating the primacy conditional on the players' returns to adoption yields the same conclusion: primacy effects exist but are short-lived.

Comparing trends among those who won with those who lost their first berserk shows that the primacy effect is likely driven by overoptimism. While the adoption of those who lost the first berserk remains constant across time, the adoption for those who won is high early on but starts diminishing until it reaches the level of those who lost. The primacy effect is probably temporary, because these overconfident players gained no real advantage from their extra berserking and eventually recognized it. In the first 50 games after their first berserk, they go berserk, on average, against opponents they hold less of an advantage over, leading to lower win probability and fewer points in these games. Consequently, it is rational for them to adjust their berserking behavior, converging with that of those who lost the first berserk.

This paper contributes to two strands of literature. The first contribution of the paper is to the theoretical and empirical literature on path dependencies in belief formation, confirmation bias, and primacy (e.g., Rabin and Schrag, 1999; Macchiavello et al., 2020; Beaman et al., 2009; Pouget et al., 2016; Bohren et al., 2019; Compte, 2021; Hogarth and Einhorn, 1992; Asch, 1946; Murdock Jr, 1962). The key contribution to this literature is to empirically disentangle the causal effect of primacy from selection.

The second one concerns learning and technology adoption (e.g., Conley and Udry, 2010; Duflo et al., 2011; Foster and Rosenzweig, 2010; Feder et al., 1985), particularly the work discussing low take-ups of efficient technologies and firms' non-adoption of profitable technologies (Suri, 2011; Verhoogen, 2023; Bloom et al., 2013; Atkin et al., 2017). The contribution to this literature is to highlight and provide evidence for the role (or lack thereof) of path dependencies, particularly the first experimentation outcome, in the long-term adoption of productive technologies.

The remainder of the paper is structured as follows: Background and Data, Empirical Strategy, Estimation, and Conclusion.

# 1 Background and Data

The data used is the open database of licess.org (Lichess, 2025), a free, open-source chess server. There are more than five million games on the platform every day, and all the rated games, over six billion, are available for download under the Creative Commons CC0 license. For each of these games, I can observe who played, the time control (how much time they had available to make their moves), the color they played, the match outcome, their Elo rating (a measure of their skill for that *time control*), among other variables.

Lichess classifies its games into *time controls* based on the total expected time players have available to make moves in the match: Ultrabullet ($<$ 29 seconds), Bullet (30-179 seconds), Blitz (180-479 seconds), Rapid (480-1499 seconds), and Classical ($\geq$ 1500 seconds).[1] The elo rating is a precise numerical system for calculating players' relative skill levels in zero-sum games, such as chess. Lichess uses the Glicko-2 system, where players initial elo rating is 1500. After each game, the player's elo is updated as a function of the outcome of the match, the player's rating deviation, the opponent's elo,

[1]Lichess computes the expected time to make moves as initial time on the clock plus the increment (how much additional time the player gets after each move) times 40 (an average benchmark for the length of a match). See Lichess (2024) for additional details.

and the opponent's rating deviation. Glickman (2022) provides details and examples of how the Glicko-2 ratings are computed.[2]

The berserk technology is only available in tournaments. In the arena tournaments studied, players return to the arena after each match, and they are matched with available players, ensuring they play as many games as possible within the allotted time. The tournament winner is the player with the most points. Players get 2 points for a win, 1 for a tie, and 0 for a loss. If a player wins two games in a row, they are on a streak, meaning the next games will earn double points until they fail to win. If the player goes berserk and wins, they earn an additional point (only one, regardless of whether the player is on a streak). Hence, berserk lowers the winning probability (half the time to make moves) and increases the payoff for winning (by 1 point).

The data spans from 2013 to 2020 and includes 1,274,062 tournaments and 268,704,933 games. I dropped tournaments with variable match duration (938 tournaments, 42,211 games) and games with missing results (33,132). I use the clock variable to identify berserking. Specifically, when the player's starting time left is half the time she should have in that game, she went berserk. Therefore, I dropped games before the clock variable became available in April 2017 (26,827,443 games) and those with missing clock values after April 2017 (5,492,183). I focus on Blitz tournaments, the most popular variant (54.08% of tournaments), where players have an expected time to make moves of between three and eight minutes in total. Among these, I keep only official tournaments (106,200), and within these, only those where the berserk option is available (77,893).

The sample includes players who berserked for the first time in an official Blitz tournament and who did not play any other tournament prior to March 2017 (152,609 players) because, for those who played prior to when the clock variable is available, it is unknown whether their first berserk in the data is their actual first berserk or their first observable berserk. I also exclude players whose opponent berserked during their first

[2]Lichess additionally takes into account color when calculating rating changes, with more points awarded for results with the black pieces.

berserk (33,882 players), players who made fewer than seven moves in their first berserk (1,984 players), players who tied during their first berserk (4,455 players), and players who always berserked after their first berserk (15 players).[3] Under these criteria, the remaining sample has 112,273 players.

I also restrict the sample to players who had an elo score higher than that of their opponent the first time they went berserk, because for those with a lower elo score than their opponent, the probability of winning would have been less than 50% even without going berserk. As an attempt to prevent the inclusion of accounts that are not the primary account of the individual, I restrict the sample to players who did not use berserk right away.[4] For the game-level analysis, I study individuals who played at least 10 games before their first berserk, yielding a sample of 30,977 players. Similarly, at the tournament level, the estimation includes players who participated in at least one tournament before their first berserk, for a total of (40,029) players.

Table A.1 in the Appendix presents summary statistics at the game and tournament levels for these players through their 150th game and the 10th tournament following their first berserk. The probability of going berserk is 18%, the average elo rating is 1,782, the average opponent elo rating is 1,774, and the average elo rating at the first berserk is 1,774. Across the tournaments studied, the average number of players is 900, and each player plays approximately 6 games per tournament.

## Measuring returns to adoption.

The measures of returns to adoption are based on the average elo rating for the player in its observed games in Bullet tournaments, a faster modality. Players in Bullet games

[3] I exclude players whose opponent berserked during their first berserk, because it is not possible to identify which player berserked first. I also drop players who had fewer than seven moves in their first berserk game because they need to make at least seven moves to get the extra berserk point.

[4] Technically, players could have multiple accounts. In at least one account, they are careful not to degrade their elo score. Yet, they may have another account where they go berserk in every match, including the very first one.

have, on average, 30 seconds to 3 minutes to make moves, whereas in the studied Blitz games, they have 3 to 8 minutes and 1.5 to 4 minutes when they go berserk. The underlying idea of this measure is that players who are better at faster modalities will also benefit more from a technology that requires them to make moves with only half the available time. The main measure of returns to adoption includes games prior to and after the first berserk to maximize the number of players with a measure of returns to adoption (still, 40% of players have a missing measure).[5] The empirical section shows that the missing behave as missing at random relative to the outcome of interest: the estimated effects on adoption for the full sample and for those with a measure of returns to adoption are very similar.

A potential concern with the primary measure of returns is that it includes games after the first berserk, which may introduce endogeneity. Therefore, the empirical section shows that the results are robust to using a measure constructed solely from games that occurred prior to the first berserk. While the sample size declines significantly, the patterns are practically identical.

## Descriptive Evidence

Based on raw data, players with a positive first experimentation outcome are more likely to use the technology in the future. The top of Figure 1 displays the probability of going berserk for the next one hundred and fifty games after the first berserk, showing a clearly higher likelihood of using the technology if the first impression is a win ($\approx$3 percentage points). The bottom of Figure 1 suggests that the intensity of the first impression is also important. Players who achieve the best possible first outcome—winning with more than 5% of the clock remaining—are the most likely to adopt the technology in the future. Their adoption is followed by the group with the second-best possible first outcome: winning with less than 5% of the clock remaining. The next group comprises those who lost by checkmate. Finally, the group with the lowest adoption rate is also

[5] Using the elo score of ultrabullet, the fastest modality, the missing share increases beyond 75%.

the one that had the worst first impression: those who lost due to running out of time. These patterns are consistent and very similar when estimated at the tournament level instead of at the game level (see Figure A.1).

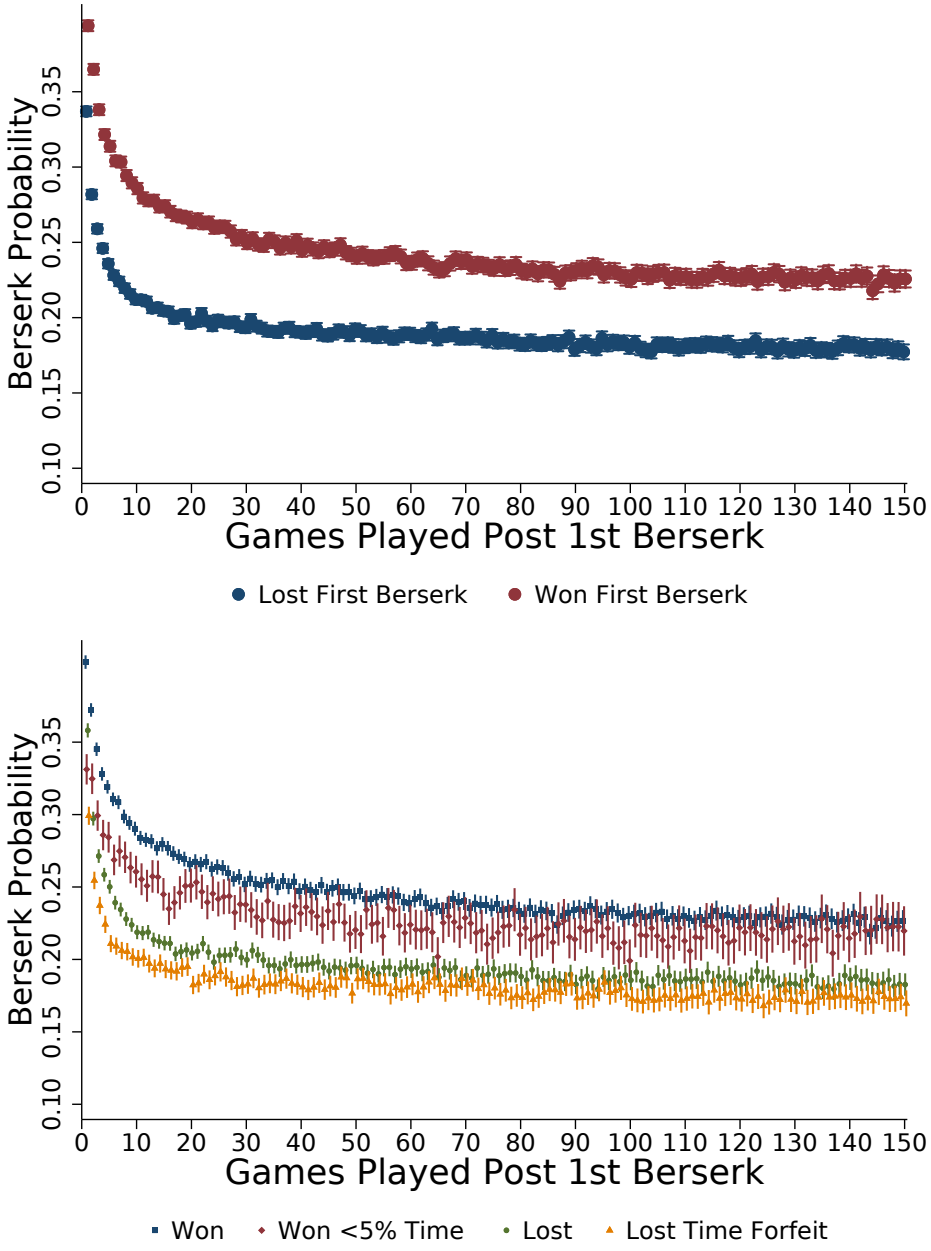



Figure 1: Descriptive Evidence: First Impression and Technology Adoption

Note: The top figure shows the probability of using berserk in the 150 games following its first use, distinguishing between players who won or lost the first berserk. The bottom figure further disaggregates by the intensity of the first impression signal: won (with more than 5% of time left), won with <5% of time remaining, lost (by checkmate), or lost by time forfeit. Standard errors are clustered at the player level.

Although these patterns are clear, other factors may be driving the outcome of the first experiment and subsequent adoption. In particular, players who are better at

playing faster are more likely to win on the first attempt with the technology and will also have a higher return to using it in the future. Figure 2 shows that players with higher returns to adoption adopt more: an above-median elo score in Bullet (faster matches) are approximately 14.7 percentage points (57%) more likely to use berserk in the following 150 games.

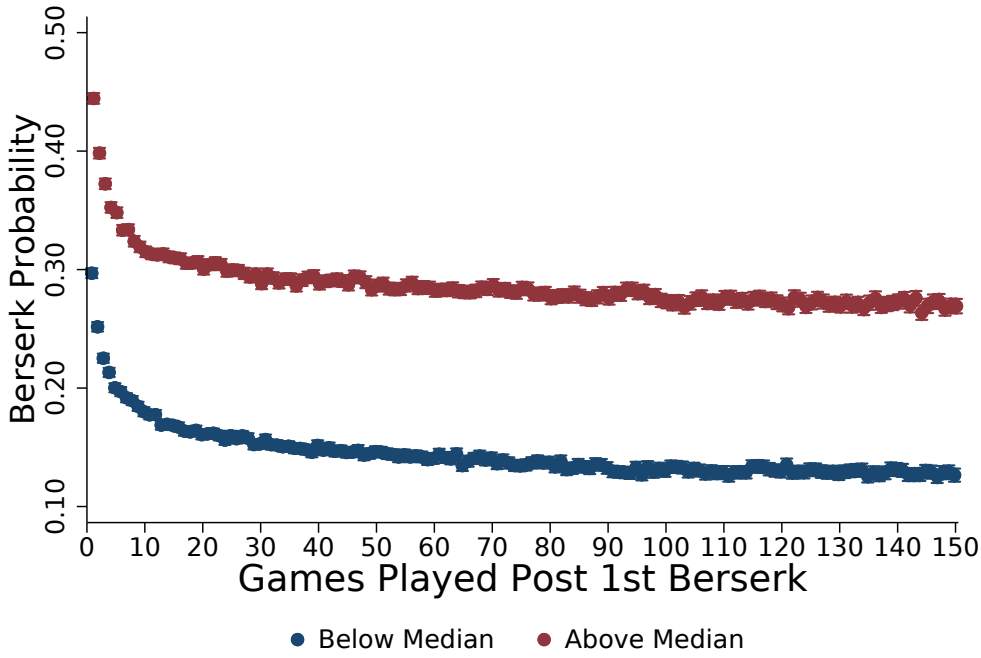



Figure 2: Descriptive Evidence: Adoption Returns and Technology Adoption

Note: The figure shows the probability of using berserk in the 150 games following its first use, distinguishing between players above and below the median of the returns-to-adoption measure, defined as the player's elo rating in bullet games. Standard errors are clustered at the player level.

Based on the descriptive evidence, both primacy and selection could account for the later higher take-up, and these are positively correlated. Players who won the first berserk have, on average, a 0.3 standard deviation higher return to adoption (Elo in Bullet games).

## Primacy vs Recency

An alternative explanation for the observed descriptive patterns is recency: more recent impressions carry greater weight in belief formation. If recency were the prevailing mechanism, the pattern observed relative to the first impression would be similar to that observed relative to any other impression. Descriptively, I find no evidence of

recency effects.

Figure A.5 compares the probability of going berserk for the next 150 games after winning or losing the 10th, 20th, 30th, and 40th berserk. The first takeaway is that berserking is serially correlated. The probability of going berserk is higher immediately after a berserk. This makes intuitive sense because if a player's ranking in a tournament in game $t$ was such that the player decided to go berserk, it is likely that the incentive is also high in game $t + 1$. For example, if the player is in 5th place and wants to reach a podium position, the player would likely play aggressively (go berserk) in the remaining games to make it to the podium.

Second, players take into account their streak for their decision of going berserk. The benefit of winning on berserk is independent of being on a streak: one extra point. Yet the cost differs when on a streak. If the player is on a streak and loses, the streak ends, and double points are no longer awarded. Therefore, the player should be more willing to risk losing (going berserk) after losing the last game. This is indeed what I observe: players who lost the 10th, 20th, 30th, and 40th berserk are more likely to go berserk immediately after losing. The streak consideration is not dominant when observing the first berserk (Figure 1). This is probably because either the primacy, the selection effect, or their sum dominates the streak effect.

Finally, the outcomes of the 10th, 20th, 30th, and 40th berserks have no predictive power for whether a player will go berserk in the future. In particular, one might have expected that those with higher returns would be more likely to win these berserks and adopt more. However, players may already be choosing to go berserk optimally, and, in particular, those with higher returns could be taking greater risks by going berserk against better opponents, reducing their winning chances. This is indeed the case. Players with an above median return to adoption go berserk against players to whom, on average, they hold a lesser advantage, decreasing their winning chances when going berserk. Their average Elo difference (own Elo minus opponent Elo) in their berserk games is 20 points (12%) smaller than for players with below-median returns.

# 2 Empirical Strategy

The primary concern for identification is that lower (unobservable) types may have lower returns to adoption and poorer experimental outcomes. This leads to omitted-variable bias in an Ordinary Least Squares (OLS) estimate of the effect of the first experimentation outcome on adoption, because the player's returns to adoption (an omitted variable) affect both adoption and experimentation outcomes. This section introduces two strategies for addressing this problem: i) identify players for whom the first experimentation outcome is as good as random, i.e., independent from returns to adoption, and ii) explicitly control for returns to adoption.

To identify players for whom the outcome of the first berserk is as good as random, I estimate a random forest model with a binary outcome variable equal to one if the player wins the game. The explanatory variables include the player's elo, the opponent's elo, their elo difference, an indicator for playing with the white pieces, the number of games previously played in Blitz, Bullet, Classical, Rapid, Standard, and Ultrabullet formats, and the number of times the player previously observed opponents berserking in each category. The model's out-of-bag error is 0.343.

Figure 3 displays the relationship between the predicted win probability from the random forest and the actual probability of winning. As expected, players predicted to win are more likely to win (top-right), and those predicted to lose are more likely to lose (bottom-left). Theoretically, when the predicted win probability is 0.5, it means that, based on the observables, the player is as likely to win as to lose. Empirically, I observe that in the vicinity of the 0.5 predicted win probability, players are indeed as likely to win as they are to lose their first berserk.

The key identifying assumption is that, within this narrow neighborhood, the residual shocks that determine the match outcome are not systematically related to the returns to adoption. Although being near the 0.5 prediction of the random forest ensures that players are similar in a weighted combination of observables, it does not ensure that returns to adoption are orthogonal to a favorable first impression. To ad-

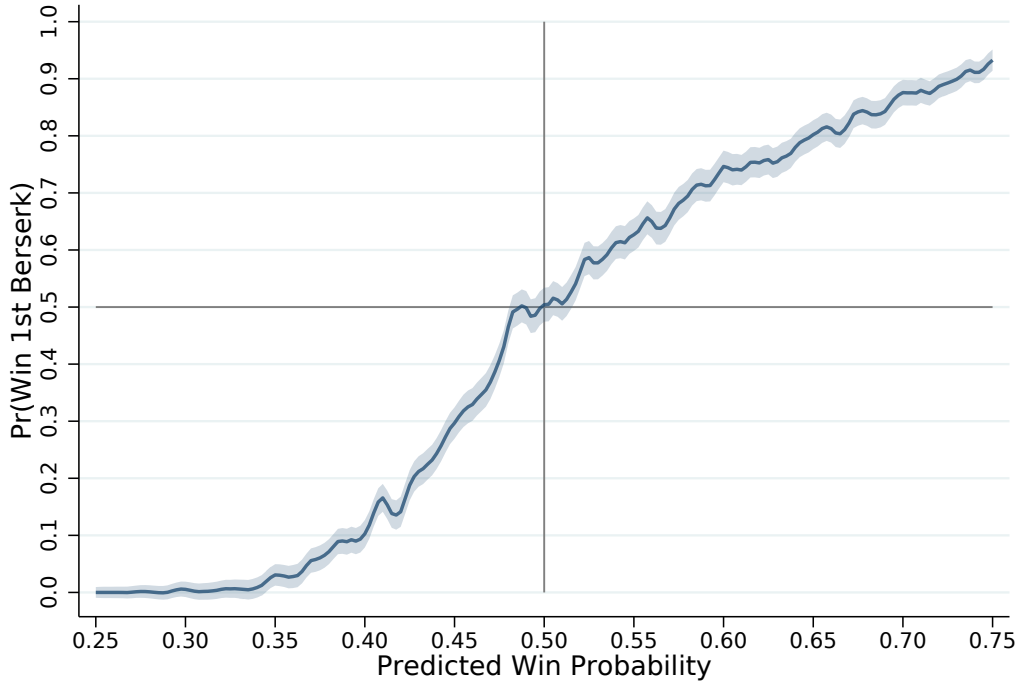



Figure 3: Winning Probability and Predicted Winning Probability

Note: The figure plots a nonparametric estimate of the probability of winning in the first berserk as a function of the predicted win probability (estimated using a Random Forest). The solid line shows a local polynomial regression of degree 3 evaluated on a grid, using a bandwidth of 0.005. Shaded areas indicate 95% confidence intervals.

dress this concern, Table A.2 document that, while the constructed measure of returns to adoption—based on a player's skill in faster time controls—is generally an important determinant of the outcome of the first berserk, it is unrelated to winning the first berserk for players in the window around the 0.5 predicted win probability. In general, a one-standard-deviation higher returns to adoption is associated with 10.7 pp higher probability of winning the first berserk (Column 4). However, in the neighborhood of the 0.5 predicted win probability, this relationship does not exist (Column 8).

This identification approach is motivated by the one used in the close-elections literature, where the winner is not random in the full sample but is effectively random in very tight races (Lee et al.; Lee, 2008). A key difference is that there is no running variable, and moreover, winning is a discrete variable. The variable used instead to identify the target sample is the predicted winning probability. In particular, the main specification focuses on players with a predicted winning probability between 0.485 and 0.515. Interestingly, these players' actual winning probability is not statistically

different from 0.5.

A potential concern with this type of identification strategy is that players could strategically sort into a neighborhood around 0.5 in win probability. This is unlikely in this setting, because if players sorted, they would prefer to sort toward a higher winning probability rather than toward being close to a toss-up. Nonetheless, the density function of the predicted win probability shows no evidence of bunching into any category (see Figure A.2).

The first analysis sample comprises players with winning probabilities between 0.485 and 0.515 who played at least 10 games or 2 tournaments before the first berserk, yielding 1,323 players for the game-level analysis and 1,686 for the tournament-level analysis. Table A.3 in the Appendix includes the summary statistics at the game and tournament levels.

The second empirical strategy to isolate the effect of primacy from that of selection based on returns to adoption is to explicitly control for returns to adoption. In these specifications, I drop the restriction in the sample of players having a predicted win probability between 0.485 and 0.515. Instead, I use the player's elo in Bullet games, which are faster, to control for unobservable returns to adoption. The control is multiplied by the time fixed effects. Therefore, the interpretation of the estimates is conditional on the players' returns to adoption.

I estimate the effect of winning the first berserk on future adoption at the game and tournament level using the following difference-in-difference equation:

$$Y_{i,t} = \phi_i + \sum_{s=-1}^{S} \delta_s G_{s,t} + \sum_{s=-1}^{S} \beta_s \left( \mathbb{1}[\text{Won 1st Berserk}]_i \times G_{s,t} \right) + \Gamma X_{i,t} + \epsilon_{i,t} \qquad (1)$$

Where $Y_{i,t}$ denotes whether the player $i$ went berserk at game $t$ relative to the first berserk (or the share of berserks in the tournament $t$ relative to the first berserk in the tournament-level specification). The player fixed effects are $\phi_i$ and the time relative to the first berserk fixed effects are $\delta_s$ (groups of games in the game-level specification and

tournaments in the tournament level specification), $G_{s,t}$ indicates whether the game $t$ belongs to the $s$-th 5-game period relative to the first berserk (in the tournament specification, it is a dummy indicating the tournament time relative to the first berserk).

The omitted period is $s = -2$ (the 6th to 10th game prior to the first berserk or the second tournament prior). Estimates of adoption prior to the first Berserk are included for completeness, but they are zero by design. The variable $\mathbb{1}[\text{Won 1st Berserk}]_i$ indicates whether player $i$ won her first game using berserk. The model also includes a set of control variables ($X_{i,t}$) that, at the game level, account the opponent's Elo rating, and the difference in Elo between the player and the opponent, as well as the player's winning streak.[6] At the tournament level, the controls include tournament size (number of players), number of moves per tournament, average player rating, total number of berserk games, and total number of games played in the tournament.[7] These controls could influence the potential cost and benefit of going berserk. Therefore, controlling for these variables allows the estimates to be interpreted as the effect of technology adoption under similar circumstances.

# 3 Estimation

Consistent with the descriptive evidence, the estimates of equation 1 indicate that having a positive first impression leads to higher adoption in the short term (Figure 4). However, unlike the descriptive evidence, the effect is not long-lasting. It fades during the first 50 games and the first 5 tournaments. Players who won the first berserk are 4 pp (21%) more likely to go berserk in the next 50 games, and their share of berserk games in tournaments increases by 3 pp (25%) in the next 5 tournaments. Yet the effect begins to fade after the first 10 games, and by the 50th game or the 5th tournament,

[6]The player winning streak fixed effect has three categories: the player did not win the last game, the player won the last game, or the player won at least the last two games. The Elo and difference in Elo controls comprise 10 dummy variables for each control, indicating the decile in the distribution.

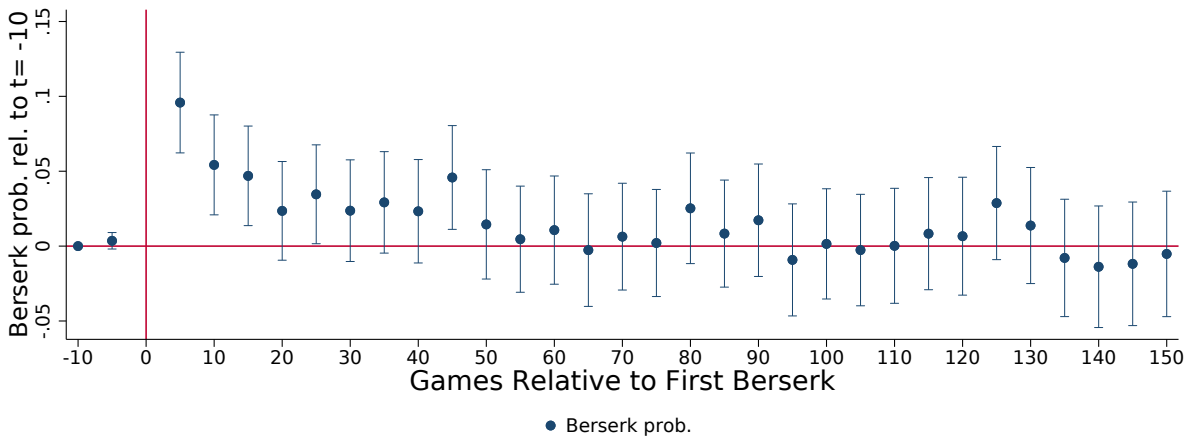
[7]Each tournament control variable is included as 10 dummy variables indicating the decile in the distribution.

it is no longer present. Moreover, the first impression does not affect continued use of the platform, i.e., the probability of reaching the next 150 tournament games after the first berserk (Figure A.3).

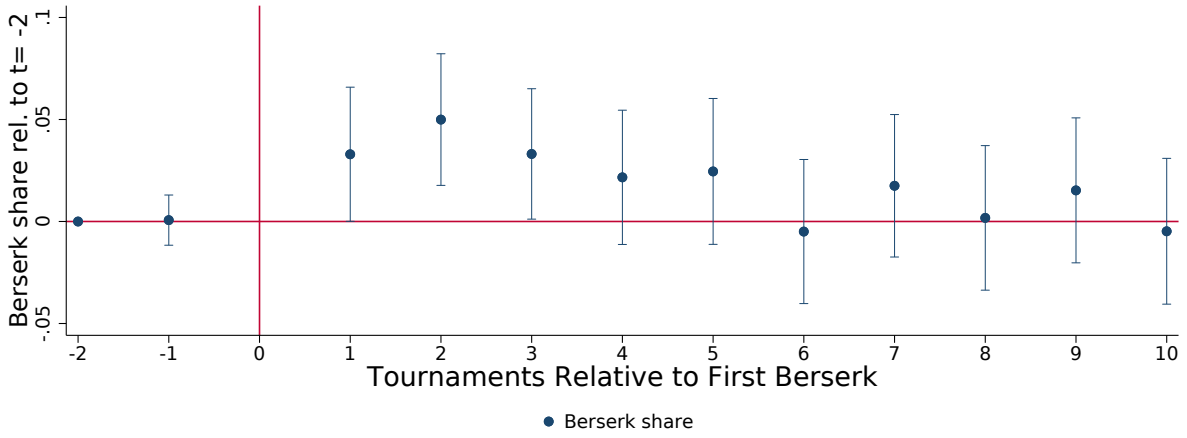
The second alternative to address the selection problem and estimate the causal effect of primacy on adoption is to control explicitly for returns to adoption. I use the player's elo in faster games (Bullet) as the measure of returns to adoption, since players who are better at playing faster games are more likely to benefit from a technology that cuts their available time in half. Figure 5 presents three sets of results: i) the estimates of equation 1 for the full sample ii) the same estimates for the sample with a measure of returns to adoption, and iii) the same estimates controlling for returns to adoption.

Similar to the descriptive evidence, Panel A of Figure 5 shows that players who had a positive first impression are more likely to continue to use the technology in the future. Yet, this effect may be biased by selection: players who had a positive first impression may also be those who are better at playing chess faster (have higher returns to adoption). Panel B shows a similar pattern when restricting the sample to players with a measure of returns to adoption, confirming that the attrition in this measure is likely unrelated to the outcome of interest. Finally, Panel C shows that, after controlling for returns to adoption, the primacy effect is no longer permanent. Similar to the effects estimated from the sample in which the first outcome was as good as random, an early-primacy effect fades over time.

Figure A.4 displays a consistent pattern at the tournament level. Players who had a positive first impression in their first berserk use berserk in a larger share of games in subsequent tournaments (Panel A). Attrition in the measure of returns to adoption appears unrelated to the outcomes of interest, as the sample with a measure of returns to adoption exhibits the same pattern (Panel B). However, after controlling for returns to adoption, this primacy effect fades within the first 5 tournaments (Panel C).




Panel (a): Game Level

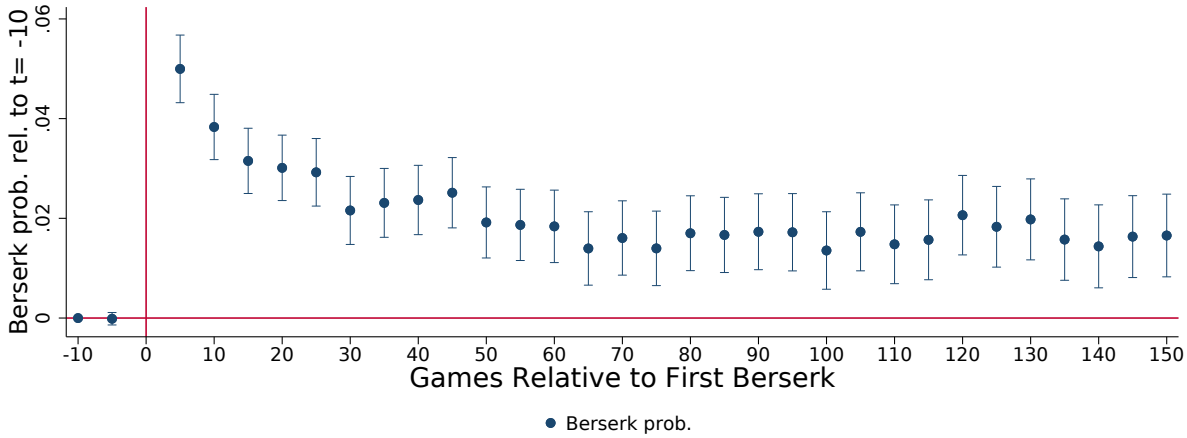



Panel (b): Tournament Level

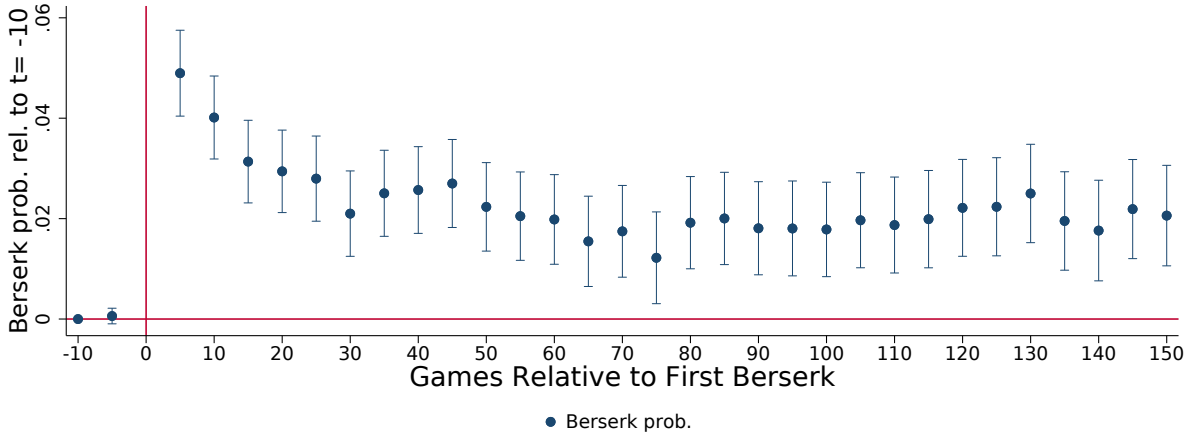
Figure 4: Effect of Winning the First Berserk on Adoption

Note: The figure displays the effect of winning the first berserk on the probability of using berserk, estimated using equation 1. Panel A presents game-level estimates and covers the 10 games prior to and the 150 games following the first berserk. The specification includes player fixed effects, 5-game-interval fixed effects, decile fixed effects for the opponent's elo rating and the Elo difference, and winning-streak fixed effects. The sample is restricted to players who played at least 10 games prior to their first berserk and whose predicted win probability lies between 0.485 and 0.515. Panel B presents tournament-level estimates and covers 2 tournaments prior and 10 tournaments following the first berserk. The specification includes player fixed effects, 1-tournament-interval fixed effect, and deciles fixed effects for tournament size (number of players), total number of moves, average player rating, total number of berserk games, and total number of games played in the tournament. The sample is restricted to players who participated in at least one tournament prior to their first berserk and whose predicted win probability lies between 0.485 and 0.515. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.

Panel (a): All




Panel (b): Returns to Adoption Available




Panel (c): Controlling for Returns to Adoption

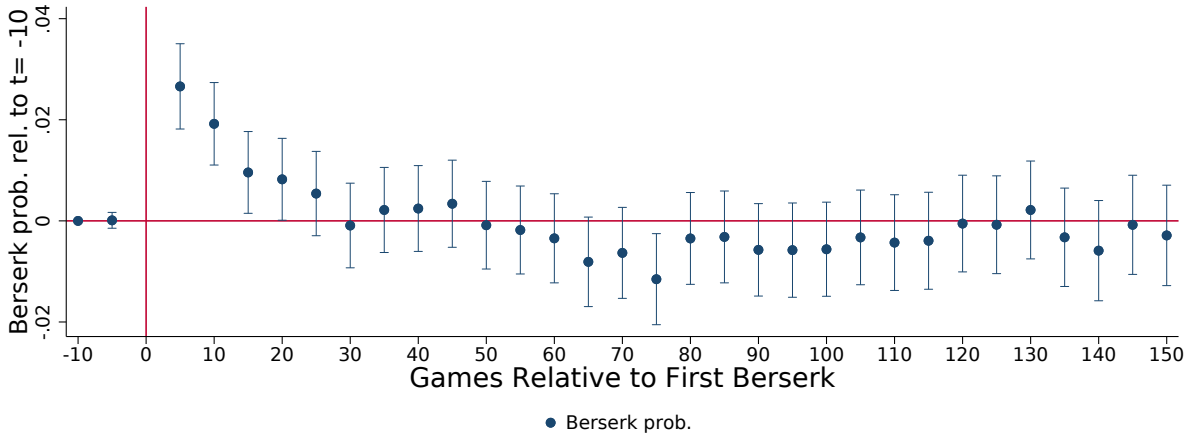



Figure 5: First Berserk Outcome, Returns to Adoption, and Adoption

Note: The figure displays the effect of winning the first berserk on the probability of using berserk, estimated using equation 1. It presents game-level estimates and covers the 10 games prior to and the 150 games following the first berserk. The sample is restricted to players who played at least 10 games prior to their first berserk. In Panel (a) the controls include player fixed effects, 5-game-interval fixed effects, decile fixed effects for the opponent's elo rating and the Elo difference, and winning-streak fixed effects. Panel (b) restricts the sample to players for whom the returns-to-adoption measure is observed. Panel (c) additionally controls for returns to adoption. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.

## 3.1 Robustness

A standard concern in the literature on treatment status assignment being as good as random in the local neighborhood of a cutoff is how close to the cutoff is considered a local neighborhood. To alleviate this potential concern, I present results for the samples for which the winning probability ranges between .48 and .52, .49 and .51 and between .495 and .505 (Figure A.7). The patterns of the estimates using these alternative cutoffs are practically indistinguishable from those of the main specification.

One potential issue with the primary returns-to-adoption measure is that it also includes games played after the first berserk, which could introduce endogeneity. Figure A.8 shows that using a measure based solely on games that occurred prior to the first berserk yields a pattern that is practically identical: a short-lived primacy effect.

## 3.2 Discussion

The primacy effect could manifest as overoptimism or scarring. When players have a positive first impression, confirmatory bias and similar mechanisms can lead to overoptimism—believing they are better at or have higher returns to the technology than they really do. On the other hand, when players have a negative first impression, the same mechanism can lead to scarring, causing them to believe they are worse at or receive lower returns from the technology than they actually are.

Looking at the patterns for winners and losers of the first berserk separately suggests that the short-term primacy effect is driven by those who overadopt after a positive first impression (see Figure A.6). While those who had a negative first impression have an almost constant probability of going berserk in the future, those who have a positive first impression are more likely to go berserk soon after the positive experience. Then the effect begins to fade and converges toward the adoption of those who had a negative first impression.

Absent effects like those described by Rabin and Schrag (1999), where first impressions matter due to confirmatory bias, it is not surprising that the primacy effect fades

with time. If players were Bayesian, the initial primacy effect would diminish over time as they continue to receive signals about the costs and benefits of this technology. Since players who are equally likely to win or lose their first berserk, by design, face the same costs and returns to adoption regardless of their first outcome, they should converge on the same level of adoption.

Therefore, the primacy effect is likely short-lived because these overconfident players were not receiving a benefit from their additional berserking, and realized it over time. The berserking outcomes in the 50 games after the first berserk of players in the sample, near the 50% chance of winning the first berserk, are consistent with this pattern. Controlling for the difference in elo with the opponent in the match, when berserking, players who won their first Berserk do not perform worse than those who lost it. They are not less likely to win or earn fewer points (Columns 1 and 2 of Table A.4).

However, Column 3 of Table A.4 shows that they take higher risks: they go berserk, on average, against opponents they hold 56 elo points less advantage (46%). Therefore, in their overall berserking (without conditioning on the difference in elo with the opponent), they are 3.5 pp (8%) less likely to win, and they make .09 fewer points per game (6%). Consequently, it is rational for them to adjust their berserking behavior and eventually align with those who lost the first berserk.

# 4 Conclusion

This paper offers causal evidence on how early experimentation outcomes influence technology adoption by separating causal primacy effects from selection bias. Using a near-threshold-as-random design in a setting with exceptionally detailed data, the analysis demonstrates that positive first impressions causally boost short-term adoption, but these effects fade quickly as individuals gain more experience. The observed gaps in adoption in the raw data are primarily driven by differences in returns rather than by persistent belief distortions stemming from initial outcomes. These findings imply that while early experiences are influential, they do not necessarily cause long-term path

dependence in adoption when agents continue to receive informative feedback. More broadly, the results reconcile common beliefs about the significance of first impressions with a learning process in which experimentation is repeated and feedback is frequent: primacy exists, but its impact is temporary.

These findings have implications for policy. The results warn against relying too heavily on one-time interventions designed to create positive initial experiences, such as trials or demonstrations, to promote long-term adoption. Instead, policies that enhance the fit between technologies and users, lower experimentation costs, or support ongoing learning may prove more effective in pushing adoption towards its optimal level. Although the empirical setup is intentionally simplified, the underlying mechanisms—such as unique returns, limited social learning, and repeated experimentation—are key factors in many real-world adoption decisions.

# References


S. E. Asch. Forming impressions of personality. *The journal of abnormal and social psychology*, 41(3):258, 1946.

D. Atkin, A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen. Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. *The Quarterly Journal of Economics*, 132(3):1101–1164, 2017.

L. Beaman, R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova. Powerful women: does exposure reduce bias? *The Quarterly journal of economics*, 124(4):1497–1540, 2009.

N. Bloom, B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts. Does management matter? evidence from india. *The Quarterly journal of economics*, 128(1):1–51, 2013.

J. A. Bohren, A. Imas, and M. Rosenberg. The dynamics of discrimination: Theory and evidence. *American Economic Review*, 109(10):3395–3436, 2019.

O. Compte. On belief formation and the persistence of superstitions. *Paris School of Economics Working Paper*, 2021.

T. G. Conley and C. R. Udry. Learning about a new technology: Pineapple in ghana. *American Economic Review*, 100(1):35–69, 2010.

E. Duflo, M. Kremer, and J. Robinson. Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American economic review*, 101(6):2350–2390, 2011.

G. Feder, R. E. Just, and D. Zilberman. Adoption of agricultural innovations in developing countries: A survey. *Economic development and cultural change*, 33(2): 255–298, 1985.

A. D. Foster and M. R. Rosenzweig. Microeconomics of technology adoption. *Annual Review of Economics*, 2(1):395–424, 2010.

M. E. Glickman. Example of the Glicko-2 System. Technical report, Boston University, Mar. 2022. URL https://www.glicko.net/glicko/glicko2.pdf. Revision dated March 22, 2022.

R. M. Hogarth and H. J. Einhorn. Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24(1):1–55, 1992.

D. S. Lee. Randomized experiments from non-random selection in us house elections. *Journal of Econometrics*, 142(2):675–697, 2008.

D. S. Lee, E. Moretti, and M. J. Butler. Do voters affect or elect policies? evidence from the us house. *The Quarterly Journal of Economics*, 119(3):807–859.

Lichess. Lichess time controls definition. https://lichess.org/faq#time-controls, 2024.

Lichess. Lichess database. https://database.lichess.org/, 2025.

J. List. Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial. Technical report, 2020.

R. Macchiavello, A. Menzel, A. Rabbani, and C. Woodruff. Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector. Technical report, National Bureau of Economic Research, 2020.

K. Munshi. Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1):185–213, 2004.

B. B. Murdock Jr. The serial position effect of free recall. *Journal of experimental psychology*, 64(5):482, 1962.

I. Palacios-Huerta. The beautiful dataset. *Journal of Economic Literature*, 63(4): 1363–1423, December 2025. doi: 10.1257/jel.20241616. URL https://www.aeaweb.org/articles?id=10.1257/jel.20241616.

S. Pouget, J. Sauvagnat, and S. Villeneuve. A mind is a terrible thing to change: Confirmatory bias in financial markets. *The Review of Financial Studies*, 30(6): 2066–2109, 12 2016. ISSN 0893-9454. doi: 10.1093/rfs/hhw100.

M. Rabin and J. L. Schrag. First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82, 1999.

T. Suri. Selection and comparative advantage in technology adoption. *Econometrica*, 79(1):159–209, 2011.

E. Verhoogen. Firm-level upgrading in developing countries. *Journal of Economic Literature*, 61(4):1410–1464, 2023.

# Online Appendix

## Figures

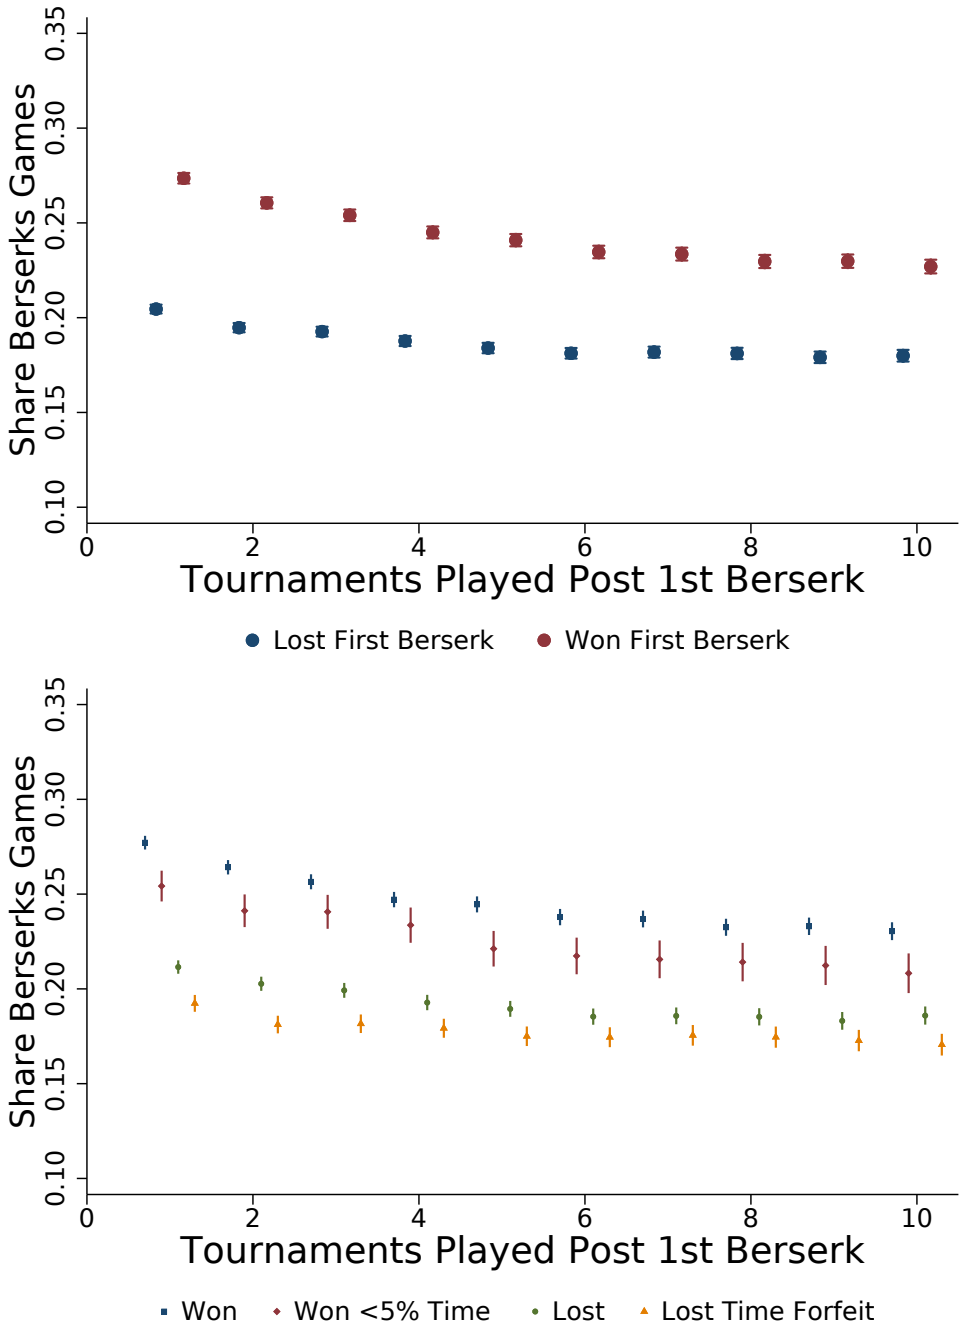



Figure A.1: Descriptive: First Impression and Adoption - Tournament level

Note: The top figure shows the share of games using berserk in the 10 tournaments following its first use, distinguishing between players who won or lost the first berserk. The bottom figure further disaggregates by the intensity of the first impression signal: won (with more than 5% of time left), won with <5% of time remaining, lost (by checkmate), or lost by time forfeit. Standard errors are clustered at the player level.

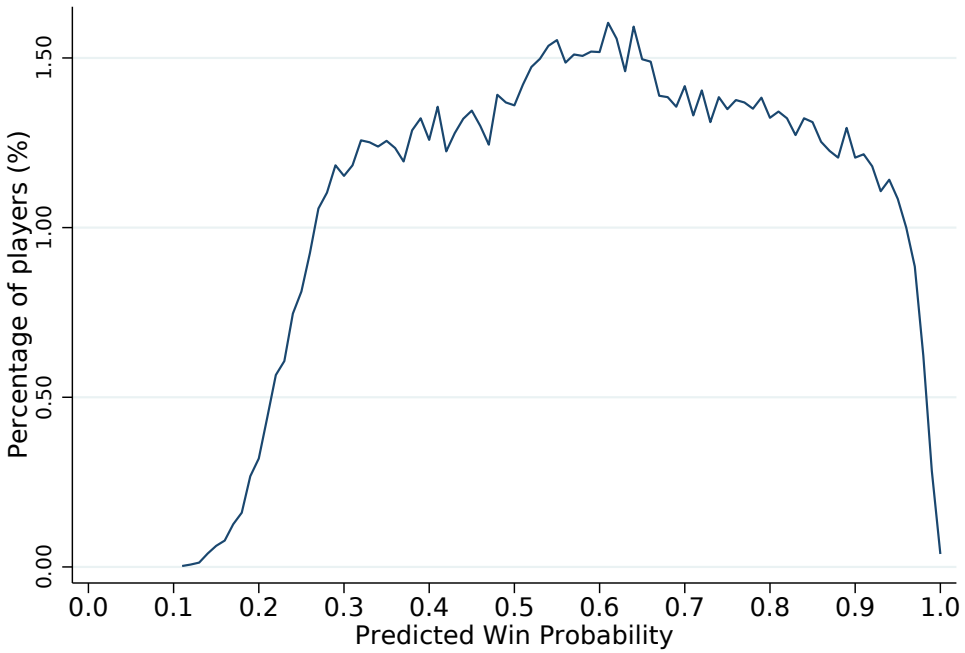



Figure A.2: Distribution of Predicted Win Probability

Note: The figure displays the distribution of predicted win probabilities generated by the Random Forest model. Predictions correspond to player-level win probabilities and are grouped into two-decimal bins. The y-axis reports the percentage of players in each predicted-probability bin.

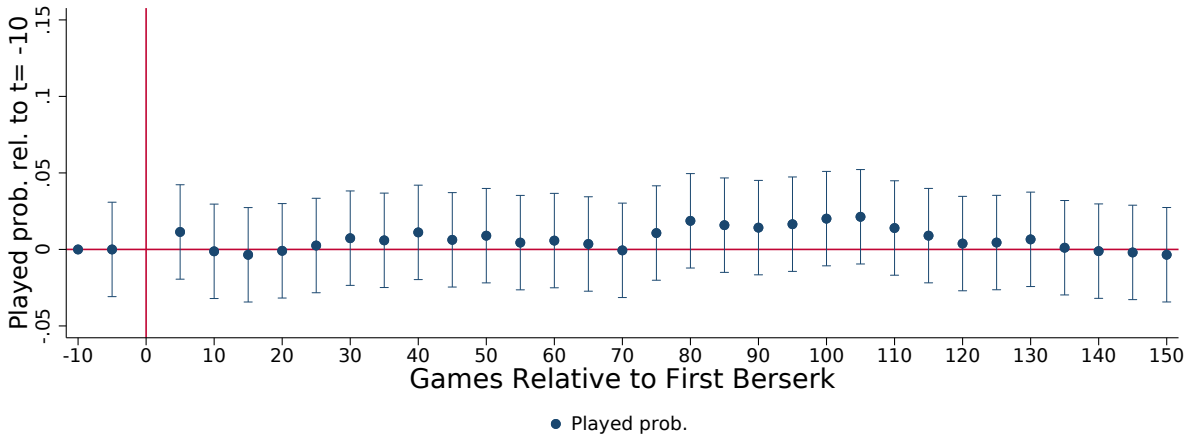
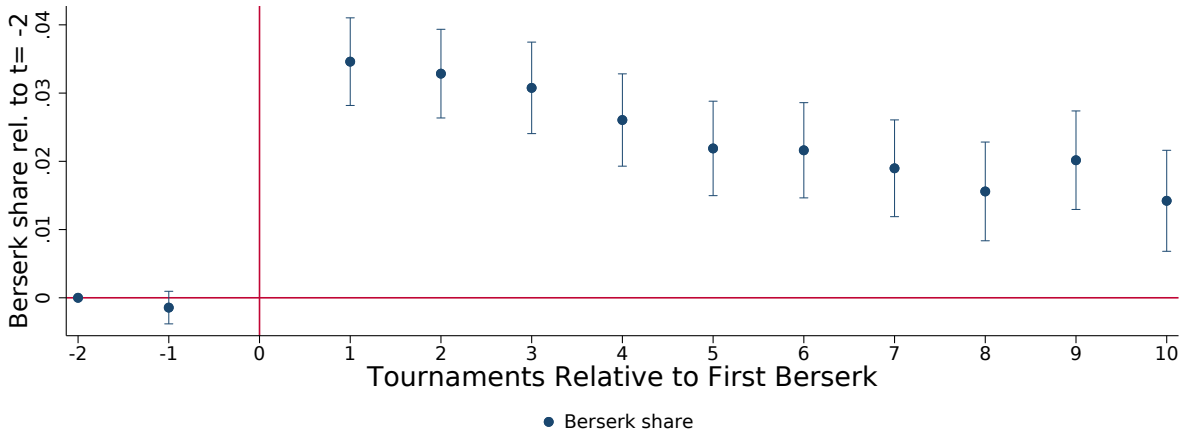
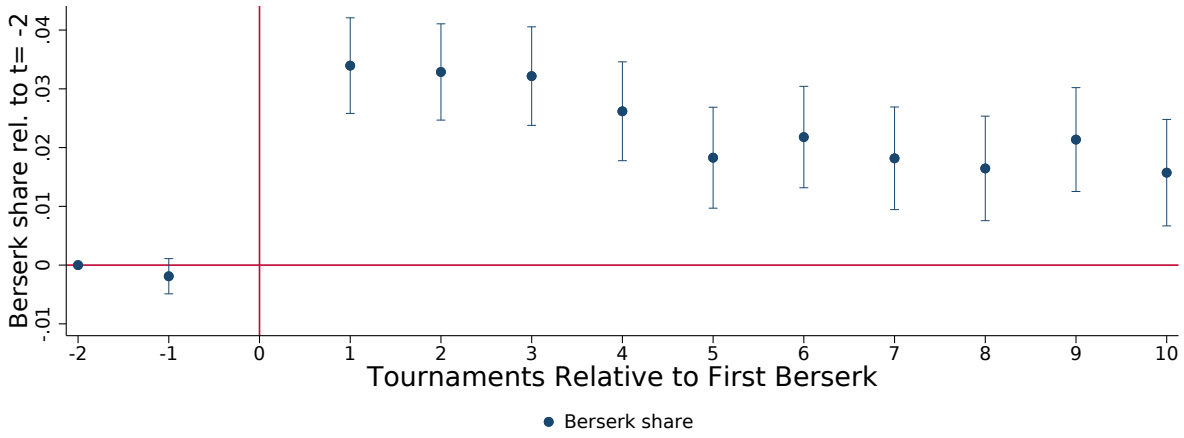



Figure A.3: Extensive Margin: Probability of Playing

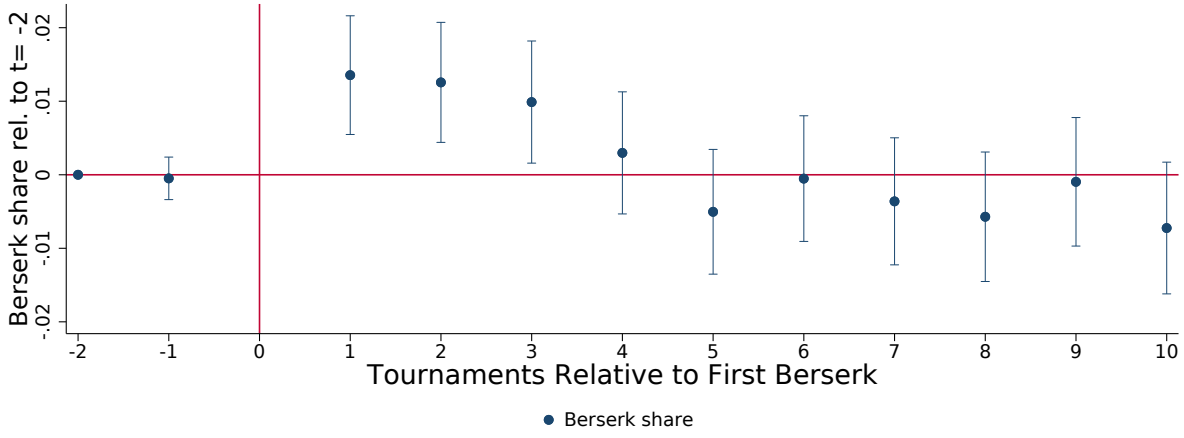
Note: The figure displays the effect of winning the first berserk on the probability of reaching a certain number of games after the first berserk, estimated using equation 1. The analysis is conducted at the game-level over a balanced window spanning 10 games prior to and 150 games following the first berserk. The sample is restricted to players who played at least 10 games prior to their first berserk and whose predicted win probability lies between 0.485 and 0.515. The outcome variable is a dummy variable that takes the value of 1 if the player played that game and 0 if the player did not, i.e., stopped playing tournaments before reaching that game. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.




Panel (a): All




Panel (b): Returns to Adoption Available




Panel (c): Controlling by Returns to Adoption

Figure A.4: First Berserk Outcome, Returns to Adoption, and Adoption - Tournament Level

Note: The figure displays the effect of winning the first berserk on the probability of using berserk, estimated using equation 1. It presents tournament-level estimates and covers the 2 tournaments prior to and the 10 tournaments following the first berserk. The sample is restricted to players who played at least 1 tournament prior to their first berserk. In Panel (a) the controls include player fixed effects, 1-tournament-interval fixed effect and deciles fixed effects for tournament size (number of players), total number of moves, average player rating, total number of berserk games, and total number of games played in the tournament. Panel (b) restricts the sample to players for whom the returns-to-adoption measure is observed. Panel (c) additionally controls for returns to adoption. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.

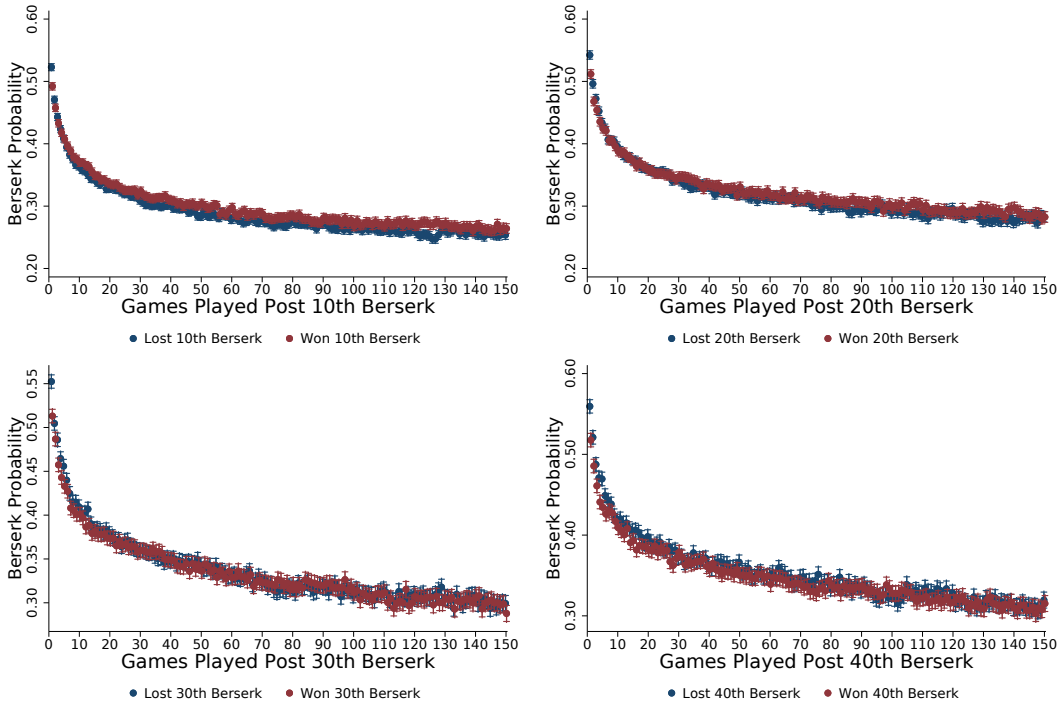



Figure A.5: Descriptive Evidence: Primacy vs Recency

Note: The figure shows the probability of using berserk in the 150 games following the 10th, 20th, 30th, and 40th berserk, distinguishing between players who won or lost the corresponding berserk. Standard errors are clustered at the player level.

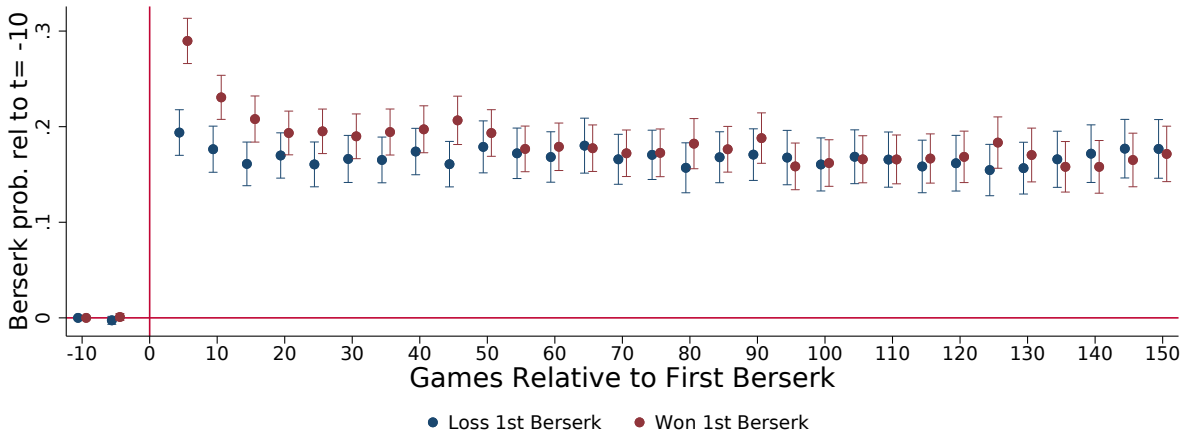
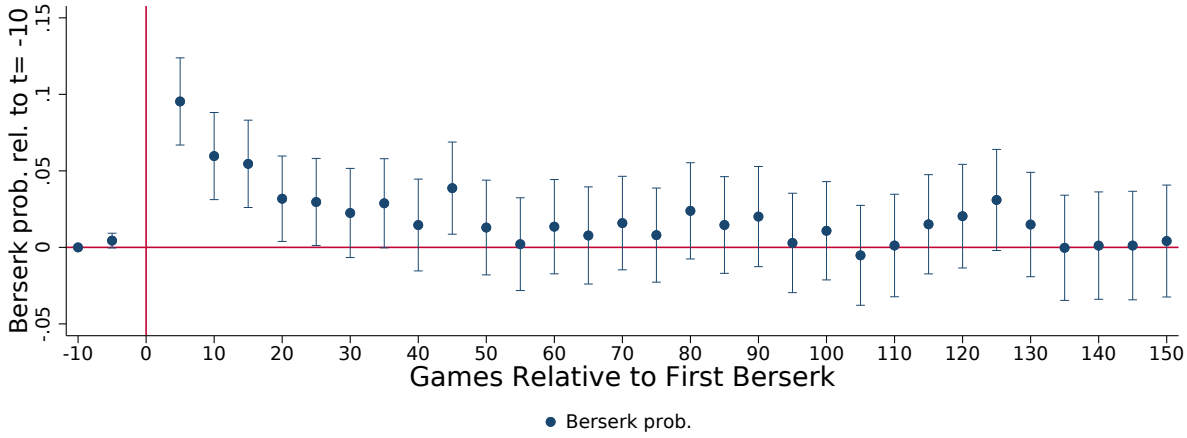
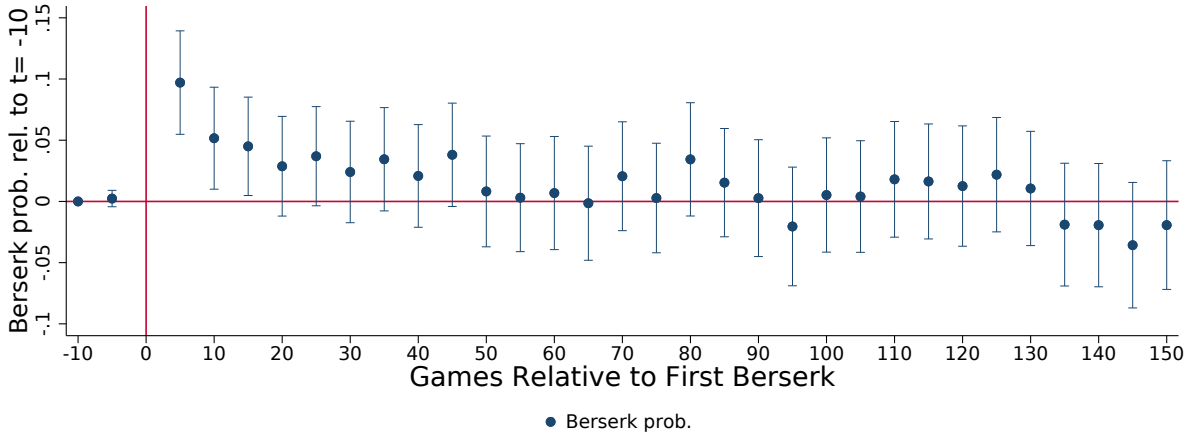



Figure A.6: First Berserk Outcome and Adoption

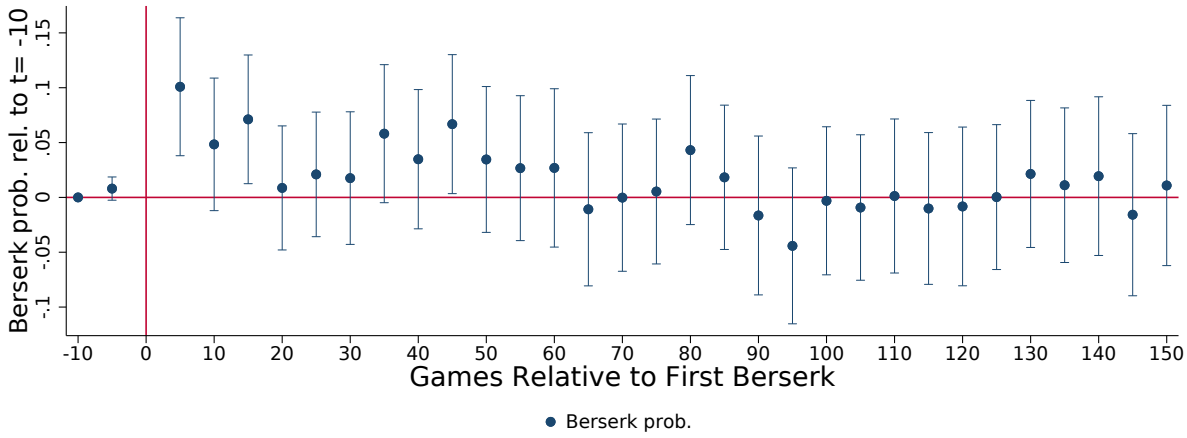
Note: The figure displays the first differences in berserk usage for players who won their first berserk and for those who lost it, estimated using equation 1. It presents game-level estimates and covers the 10 games prior to and the 150 games following the first berserk. The specification includes player fixed effects, 5-game-interval fixed effects, decile fixed effects for the opponent's elo rating and the Elo difference, and winning-streak fixed effects. The sample is restricted to players who played at least 10 games prior to their first berserk and whose predicted win probability lies between 0.485 and 0.515. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.




Panel (a): Estimated Winning Probability ∈ (0.48, 0.52)




Panel (b): Estimated Winning Probability ∈ (0.49, 0.51)

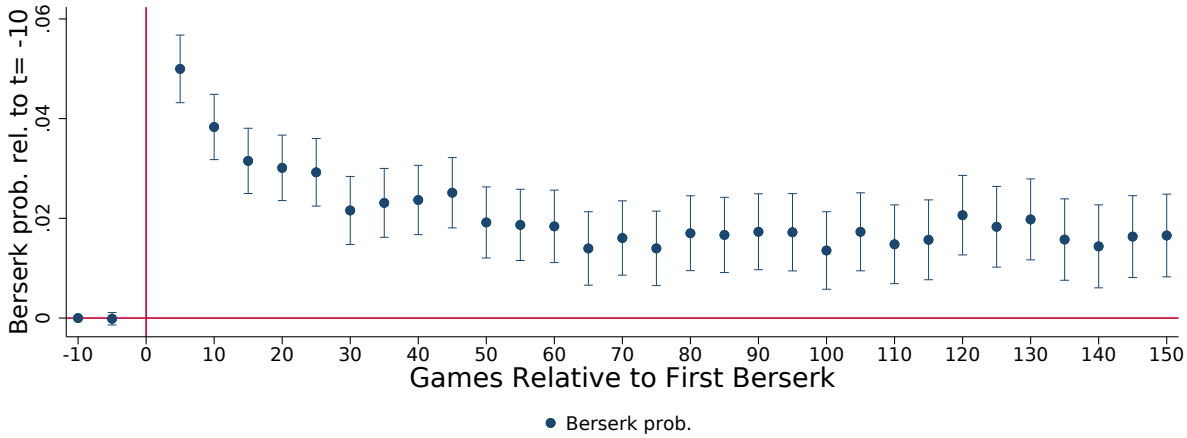



Panel (c): Estimated Winning Probability ∈ (0.495, 0.505)

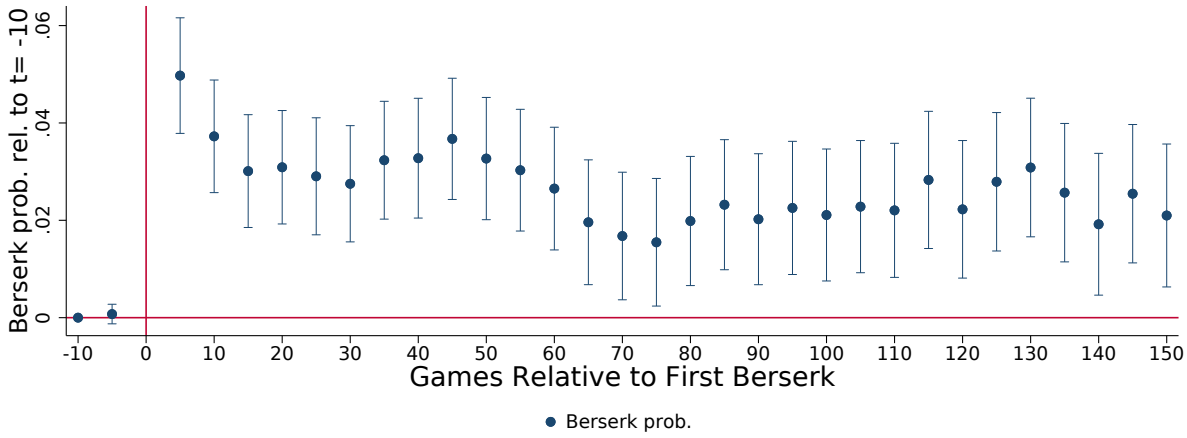
Figure A.7: Robustness Cutoff in Random Forest Estimated Winning Probability: Effect of Winning the First Berserk on Adoption

Note: The figure displays the effect of winning the first berserk on the probability of using berserk, estimated using equation 1. Each panel presents game-level estimates and covers the 10 games prior to and the 150 games following the first berserk and restricts the sample to players who played at least 10 games prior to their first berserk. The specification includes player fixed effects, 5-game-interval fixed effects, decile fixed effects for the opponent's elo rating and the Elo difference, and winning-streak fixed effects. In Panel A the sample is also restricted to whose predicted win probability lies between 0.48 and 0.52. Panel B contemplate players whose predicted win probability lies between 0.49 and 0.51. Panel C considers players whose predicted win probability lies between 0.495 and 0.505. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.

Panel (a): All




Panel (b): Returns to Adoption Available (Pre–First-Berserk games only)




Panel (c): Controlling for Returns to Adoption (Pre–First-Berserk games only)

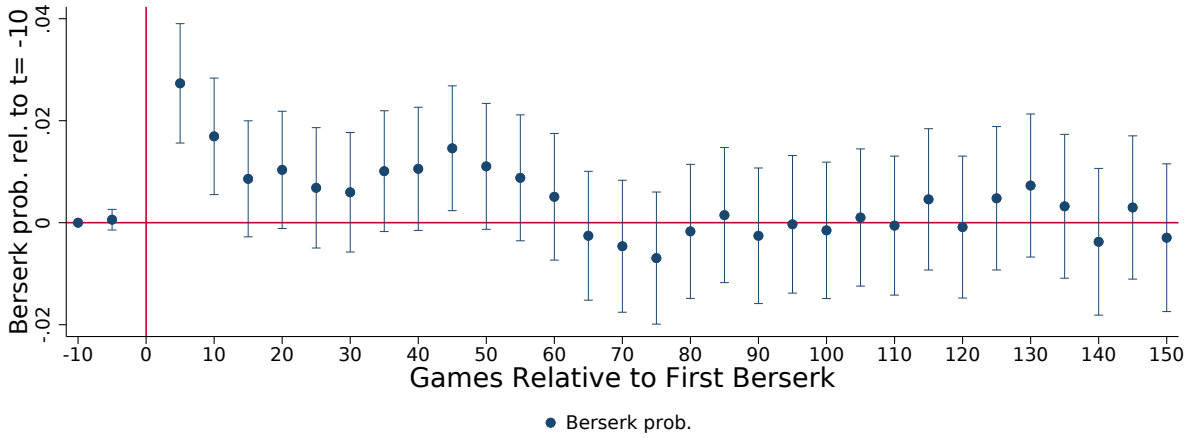



Figure A.8: Robustness: First Berserk Outcome, Returns to Adoption (pre–first-Berserk games only), and Adoption

Note: The figure displays the effect of winning the first berserk on the probability of using berserk, estimated using equation 1. It presents game-level estimates and covers the 10 games prior to and the 150 games following the first berserk. The sample is restricted to players who played at least 10 games prior to their first berserk. In Panel (a) the controls include player fixed effects, 5-game-interval fixed effects, decile fixed effects for the opponent's elo rating and the Elo difference, and winning-streak fixed effects. Panel (b) restricts the sample to players for whom the returns-to-adoption measure—constructed exclusively using games played prior to the first berserk—is observed. Panel (c) additionally controls for this pre–first-berserk measure of returns to adoption. All coefficients are displayed with their 95% confidence interval. Standard errors are clustered at the player level.

# Tables

Table A.1: Summary Statistics

| | Mean (1) | SD (2) | P25 (3) | P75 (4) |
|---|---|---|---|---|
| **PANEL A: Games** | | | | |
| Berserk Rate after First Berserk | 0.183 | 0.221 | 0.020 | 0.267 |
| Winning Streak | 0.198 | 0.124 | 0.111 | 0.263 |
| Elo Rating | 1,782 | 256.0 | 1,603 | 1,959 |
| Elo Opponent | 1,774 | 148.4 | 1,696 | 1,870 |
| Elo Rating at First Berserk | 1,777 | 265.0 | 1,591 | 1,962 |
| Games after First Berserk | 92.3 | 59.8 | 28.0 | 150.0 |
| **PANEL B: Tournaments** | | | | |
| % Berserk Games in the Tournament | 0.126 | 0.169 | 0.000 | 0.189 |
| Top 5% | 0.075 | 0.154 | 0.000 | 0.083 |
| Ranking | 0.542 | 0.190 | 0.408 | 0.679 |
| Games Played per Tournament | 6.0 | 3.3 | 3.8 | 7.7 |
| Tournaments after First Berserk | 7.1 | 3.9 | 3.0 | 10.0 |
| Players per Tournament | 903.8 | 939.9 | 364.1 | 1,107 |
| Moves per Tournament | 239,745 | 524,203 | 62,164 | 229,212 |
| Elo Player | 1,725 | 105.4 | 1,688 | 1,793 |
| Berserks per Tournament | 1,011 | 2,470 | 238.5 | 924.0 |
| Games per Tourmament | 3,417 | 7,465 | 883.5 | 3,269 |

Note: The table presents summary statistics. Panel A reports statistics at the game level, including berserk rate, winning streak rate, Elo ratings, and games played after first berserk. The sample includes 30,977 players who played at least 10 games before their first berserk, with a positive Elo rating difference in their first berserk game. Players are observed over a window spanning 10 games prior to and 150 games following the first berserk. Panel B displays summary statistics for tournament-level analysis, including the share of berserks played in the tournament by the player, probability of being at the top 5%, raking, games played by tournament and tournaments played after the first berserk. The table also reports statistics on tournament characteristics (e.g., number of players, total moves, average Elo, berserk games, and total games), computed at the player–tournament level. The sample includes 40,029 players who participated in at least one tournament before their first berserk, with a positive Elo rating difference in their first berserk game. Players are observed over a window spanning 2 tournaments before and 10 tournaments after the first berserk. The columns display the mean, standard deviation (SD), 25th percentile (P25), and 75th percentile (P75).

Table A.2: Returns to Adoption and First Berserk

| Dependent Variable: | Win First Berserk | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Sample: All First Berserk | | | | Sample: Estimated Winning Probability $\in$ (0.485, 0.515) | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Ret. To Adoption (std) | 0.1118*** | 0.1071*** | 0.1073*** | 0.1067*** | 0.0414 | -0.0013 | -0.0017 | -0.0002 |
| | (0.003) | (0.003) | (0.003) | (0.003) | (0.026) | (0.015) | (0.015) | (0.014) |
| Observations | 68,052 | 68,052 | 68,052 | 68,052 | 828 | 828 | 828 | 828 |
| Mean Dep. Var | 0.45 | 0.45 | 0.45 | 0.45 | 0.54 | 0.54 | 0.54 | 0.54 |
| Elo (std) | Y | Y | Y | Y | Y | Y | Y | Y |
| Difference in Elo (std) | | Y | Y | Y | | Y | Y | Y |
| Streak | | | Y | Y | | | Y | Y |
| Color | | | | Y | | | | Y |

Note: The table presents OLS regression estimates predicting the probability of winning a player's first berserk game. The dependent variable across all specifications is a binary indicator for winning the first berserk match. The independent variable is the returns to adoption variable, which, along with other continuous Elo measures, has been standardized within the sample of all first berserk players. Columns 1-4 includes all players experiencing their first berserk game. Columns 5-8 restrict the sample to players who played at least 10 games prior to their first berserk and whose predicted win probability lies between 0.485 and 0.515. Control variables are introduced sequentially and include the player's standardized elo and elo difference between opponents, the player's streak, and color. Standard errors in parentheses are clustered at the player level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.3: Summary Statistics: Estimated Winning Probability ∈ (0.485, 0.515)

| | Mean | SD | P25 | P75 |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| **PANEL A: Games** | | | | |
| Berserk Rate after First Berserk | 0.187 | 0.223 | 0.020 | 0.267 |
| Winning Streak | 0.186 | 0.118 | 0.106 | 0.238 |
| Elo Rating | 1,746 | 236.8 | 1,579 | 1,906 |
| Elo Opponent | 1,759 | 145.2 | 1,680 | 1,854 |
| Elo Rating at First Berserk | 1,737 | 248.4 | 1,565 | 1,901 |
| Games after First Berserk | 94.0 | 60.2 | 28.0 | 150.0 |
| **PANEL B: Tournaments** | | | | |
| % Berserk Games in the Tournament | 0.129 | 0.178 | 0.000 | 0.186 |
| Top 5% | 0.062 | 0.142 | 0.000 | 0.083 |
| Ranking | 0.530 | 0.186 | 0.402 | 0.662 |
| Games Played per Tournament | 5.9 | 2.9 | 3.8 | 7.6 |
| Tournaments after First Berserk | 7.1 | 3.9 | 3.0 | 10.0 |
| Players per Tournament | 847.1 | 908.6 | 332.2 | 1,041 |
| Moves per Tournament | 223,510 | 505,833 | 57,453 | 211,608 |
| Elo Player | 1,718 | 107.1 | 1,680 | 1,791 |
| Berserks per Tournament | 946.1 | 2,438 | 216.0 | 851.3 |
| Games per Tourmament | 3,187 | 7,213 | 814.0 | 3,025 |

Note: The table presents summary statistics. Panel A reports statistics at the game level, including berserk rate, winning streak rate, Elo ratings, and games played after first berserk. The sample includes 1,323 players who played at least 10 games before their first berserk, with a positive Elo rating difference in their first berserk game and with winning probabilities between 0.485 and 0.515. Players are observed over a window spanning 10 games prior to and 150 games following the first berserk. Panel B displays summary statistics for tournament-level analysis, including the share of berserks played in the tournament by the player, probability of being at the top 5%, raking, games played by tournament and tournaments played after the first berserk. The table also reports statistics on tournament characteristics (e.g., number of players, total moves, average Elo, berserk games, and total games), computed at the player–tournament level. The sample includes 1,686 players who participated in at least one tournament before their first berserk, with a positive Elo rating difference in their first berserk game, and winning probabilities between 0.485 and 0.515. Players are observed over a window spanning 2 tournaments prior to and 10 tournaments following the first berserk. The columns display the mean, standard deviation (SD), 25th percentile (P25), and 75th percentile (P75).

Table A.4: Performance in Berserks Games

| Dependent Variable: | Win | Points | Difference in Elo | Win | Points |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| 1[Win First Berserk] | 0.015 | 0.065 | -56.454*** | -0.035*** | -0.093** |
| | (0.012) | (0.040) | (9.280) | (0.012) | (0.039) |
| Observations | 9,629 | 9,629 | 9,912 | 9,912 | 9,912 |
| Mean Dep. Var | 0.46 | 1.49 | 120.81 | 0.46 | 1.49 |
| Elo | Y | Y | Y | Y | Y |
| Streak | Y | Y | Y | Y | Y |
| Difference in Elo | Y | Y | | | |

Note: The table reports game-level OLS estimates of the effect of winning the first berserk on subsequent performance in berserk games. The sample is restricted to berserk games played between games 1 and 50 after the first berserk, for players whose predicted win probability in the first berserk lies between 0.485 and 0.515 and who played at least 10 games prior to their first berserk. Columns (1)–(2) and (4)–(5) report estimates for win (an indicator equal to one if the player wins the game) and gained points, while column (3) reports the difference in Elo between the player and the opponent. All specifications include player's Elo and streak fixed effects. Columns (1)–(2) additionally include difference in Elo fixed effects. Standard errors in parentheses are clustered at the player level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.