

NBER WORKING PAPER SERIES

BELIEFS THAT ENTERTAIN

Ashvin Gandhi
Paola Giuliano
Eric Guan
Quinn Keefer
Chase McDonald
Michaela Pagel
Joshua Tasoff

Working Paper 32295
<http://www.nber.org/papers/w32295>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

We thank Isabelle Brocas, Juan Carrillo, Andrew Caplin, Emir Kamenica, Alex Frankel, Andrey Simonov, Troy Skinner, Jonathan Frye, Tom Cadwell, Jeffrey Lin, and participants from seminars at the University of Exeter, Claremont Graduate University, the ASSA 2023 Annual Meeting, the USC LABEL lab meeting, St. Lawrence University, and Columbia University for helpful comments. We thank Huizi Yu for her excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, nor Riot Games.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w32295>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Ashvin Gandhi, Paola Giuliano, Eric Guan, Quinn Keefer, Chase McDonald, Michaela Pagel, and Joshua Tasoff. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Beliefs that Entertain

Ashvin Gandhi, Paola Giuliano, Eric Guan, Quinn Keefer, Chase McDonald, Michaela Pagel,
and Joshua Tasoff

NBER Working Paper No. 32295

April 2024

JEL No. D8,D9

ABSTRACT

Economic research on entertainment is scant despite its large share of time use. We test economic theories of belief-based utility in the context of video-game engagement. Using data on 2.8 million matches from League of Legends, we find evidence supporting reference-dependent preferences, loss aversion, preferences for surprise and suspense, preferences for clumped surprise, and flow theory from psychology. We then leverage our estimated model and an evolutionary algorithm to find the information-revealing process that maximizes player engagement. We find that the optimal version of the game has increased game play equivalent to 43% of the winner-loser gap.

Ashvin Gandhi
University of California at Los Angeles
Anderson School of Management
110 Westwood Plaza
Los Angeles, CA 90025
and NBER
ashvin.gandhi@anderson.ucla.edu

Paola Giuliano
Anderson School of Management
UCLA
110 Westwood Plaza
C517 Entrepreneurs Hall
Los Angeles, CA 90095-1481
and IZA
and also NBER
paola.giuliano@anderson.ucla.edu

Eric Guan
Riot Games
e.guan.yw@gmail.com

Quinn Keefer
California State University, San Marcos
qkeefer@csusm.edu

Chase McDonald
Riot Games and
Carnegie Mellon University
chasemcdonald@gmail.com

Michaela Pagel
Washington University
Olin Business School
John Simon Hall 242
1 Brookings Drive
St. Louis, MO 63130
and NBER
mpagel@wustl.edu

Joshua Tasoff
Claremont Graduate University
Department of Economics 160
East Tenth Street Claremont,
CA 91711
joshua.tasoff@cgu.edu

1 Introduction

People spend a lot of time on leisure. The American Time Use Survey reveals that adults in the United States spend roughly one-fifth of their waking hours consuming entertainment (Aguiar et al., 2012). This is about 25% more time than they spend working. Nearly 9% of all leisure time is spent on “playing games and computer use for leisure.” Imputing the average wage of an American, this represents a shadow value of \$540 billion per year. The U.S. video game industry earned revenues of \$180 billion in 2020, which is more than the revenues of the global movie and North American sports industries combined.¹ Yet the economic and scientific literature on video games is scarce. In this study, we analyze preferences for entertainment in the context of video games.

We focus on how players’ experiences throughout the game affect continued play. We posit that an important component of the player experience is how their beliefs about victory evolve over the course of the match. Advances in economic theory model preferences over beliefs. We thus start this paper by presenting the leading theoretical approaches to belief-based utility and their predictions for continued play.

First, players may have a preference for anticipating victory, and therefore prefer to expect to win (Loewenstein, 1987; Caplin and Leahy, 2001). Second, players may have prospect-theory preferences and compare the match outcome to a reference point (Kahneman and Tversky, 1979). Third, players may value suspense – the expectation that an important event will occur – or surprise – experiencing the unexpected (Ely, Frankel and Kamenica, 2015). And fourth, players may have reference-dependent utility with a reference point determined by their expectations (Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006). Our theoretical analysis then informs our empirical model for testing the key predictions of each of these theories.

We also examine the importance of flow theory, one of the most prominent theories in psychology for understanding motivation and enjoyment in leisure activities. Flow is described as a pleasurable highly-engaged state in which a person loses sense of oneself (Csikszentmihalyi and Larson, 2014; Nakamura and Csikszentmihalyi, 2009). This state of flow occurs when the challenge level of an activity is matched to the skill level of the actor.

We explore the effect of how beliefs evolve over the course of a game on subsequent game-play using data from *League of Legends*, one of the most popular PC titles played globally with 151 million monthly users.² In *League of Legends*, two 5-player teams fight for approximately 30 minutes to destroy the opposing team’s base. We proxy players’ beliefs about winning by estimating the objective probability of winning at time t given the game state. Using a data set of 2.8 million matches composed of approximately 84 million minutes, we construct a minute-by-minute prediction of which team will be victorious. The estimated belief-paths represent the rational expectations of a well-calibrated player. We then observe how the belief-paths affect 28 million players’ subsequent decisions to play another match in due course.

¹See <https://www.marketwatch.com/story/videogames-are-a-bigger-industry-than-sports-and-movies-combine-d-thanks-to-the-pandemic-11608654990>.

²This number is from <https://activeplayer.io/league-of-legends/> sourced on Aug 4, 2023.

Our data set allows us to test for the presence of different kinds of preferences for beliefs. First, we operationalize anticipatory utility as a preference for a high average probability of victory. The more likely victory is, the more anticipatory utility is generated. Second, we operationalize prospect theory by setting the reference point to the average beliefs of victory throughout the game. Third, we estimate the effects of suspense and surprise throughout the game. Our predictive model allows us to create a minute-by-minute measure of surprise and suspense, and we can test whether matches that are more surprising or more suspenseful lead to greater engagement, as measured by the probability of playing another match. Fourth, we test whether people have a preference for clumped rather than smooth information revelation as predicted by expectations-based reference-dependent utility. Fifth, to test for flow theory, we proxy the challenge level of the match as the probability of victory at minute zero which varies depending on the player compositions. After all, flow theory predicts that players are most engaged when they are challenged to the right degree.

Overall, we find evidence that attributes of the belief-path have a significant effect on continued game-play.³ We first establish that winning a match leads to a higher probability of playing another match in the next 60 minutes. We use this win-loss gap as a benchmark for comparing effect sizes of the different belief-paths.

We then document our main findings. First, we find that the average probability of winning has a strong and significant negative effect on continued game-play, conditional on victory or defeat. A standard deviation increase in the average probability leads to a decrease in continued play equal to $1.5\times$ the win-loss gap for losers and $1\times$ the win-loss gap for winners. This significantly larger effect for losers is consistent with loss aversion. Second, we find that average surprise is associated with less continued play, but average suspense is associated with more continued play, especially for losers. Because surprise and suspense are correlated, these effects tend to counteract each other somewhat. Third, we find evidence for the flow-theory prediction, that there is an interior optimal challenge level at minute zero. We find that this interior optimum exists conditional on winning, but it does not exist conditional on losing. Winners play the most when the initial chance of victory is slightly higher than even, at 0.57. Finally, we find that clumped surprise leads to more game-play for losers, but less for winners.

These results naturally lead to a broader question: how should one design the revelation of information in an entertainment product? One can think of a match as if the winner is pre-destined but unrevealed, and the match is a gradual revelation of the winner’s identity. As a result, *League of Legends*, the game, can be abstracted as a data-generating process that produces belief-paths,⁴ and we can estimate the optimal distribution of these belief-paths.

We use the predictions of our theoretically motivated regression model as the objective function. Using an evolutionary algorithm, we then search for locally optimal versions of *League of Legends* in game-space. The evolutionary structural analysis implies that the information revelation could be changed to produce meaningfully more game-play. The effect is large enough to make players

³Throughout the paper, we will refer to the vector of victory probabilities over time and the vector of suspense over time collectively as the belief-path.

⁴Ely, Frankel and Kamenica (2015) refer to this as the belief martingale.

feel like 43% more of a winner in expectation. The optimized game exhibits notable differences from information revelation in *League of Legends*. It has greater information revelation in the first five minutes but decreased information revelation for later minutes. We also use gradient-boosted estimation to construct an objective function for evolutionary optimization. This optimization of the information structure produced larger gains on the order of 103% more of a winner in expectation. We conclude that this high-dimensional atheoretic estimation method discovered other dimensions of the game capable of improving subsequent game-play.

We view this paper as making two key contributions. First, we extend the literature on how beliefs affect the value of entertainment products. Second, we contribute by developing a method for counterfactual game-optimization analysis. We show how structural modeling with evolutionary optimization can improve the data-generating process to produce more engaging beliefs. As far as we are aware, this is a novel structural approach in economics and has implications for game-design, entertainment design more generally, as well as theories of beliefs-based utility.

After reviewing related literature, we present our conceptual framework in Section 2. Section 3 explains our data, and Section 4 presents our match-level regression analysis. In Section 5 we show our counterfactual game-design analysis. The discussion is in Section 6 and Section 7 concludes.

1.1 Related Literature

The academic literature on video game-play begins with [Malone \(1981\)](#), who hypothesized that challenge, fantasy, and curiosity create engagement. He surveyed children on their experiences playing several simple computer games. This style of analysis, surveying players on their motivations and experiences and connecting it to outcomes like self-reported engagement and motivation, has been extended by [Ryan, Rigby and Przybylski \(2006\)](#), [Przybylski, Rigby and Ryan \(2010\)](#), [Hoffman and Nadelson \(2010\)](#), and [Tang, Kucek and Toepfer \(2022\)](#). [Ryan, Rigby and Przybylski \(2006\)](#), for instance, show that a self-reported sense of autonomy and competence are correlated with continued game play.⁵ [Huang, Jasin and Manchanda \(2019\)](#) estimate a hidden-Markov model on the frequency and volume of game-play, using a large panel data set of video-game players in a first-person shooter game. Their primary variable is the hidden psychological state of “engagement”, which, when high, is meant to induce greater frequency and volume. They find that the length of time of neither moving up nor down in rank within the meta-tournament of the game leads to lower psychological engagement. They then use their structurally estimated model to conduct counterfactual analysis on the player-matching algorithm and find that a “learn and exploit” algorithm leads to greater predicted game-play than a nearest-rank matching algorithm.

Since [Ely, Frankel and Kamenica \(2015\)](#), several published papers in the domain of sports have estimated suspense and surprise empirically. First, [Olson and Stone \(2014\)](#) solve for the

⁵Relatedly, [Nguyen \(2020\)](#) presents a philosophical theory of games as an art of agency. Core to their theory, is the notion that players temporarily adopt goals that they otherwise do not care about (e.g., getting a ball through a hoop) so that they can experience a sense of playful striving. In their theory, the pleasure of playing games comes from this striving, and comes from the capacity to experience alternative clearly defined “agencies” across different games with different goals and different rules.

tournament format that maximizes suspense for college football playoffs. Five subsequent papers estimate the effects of in-game suspense and surprise in professional sports games on spectator viewership. [Bizzozero, Flepp and Franck \(2016\)](#), using 80 men’s single tennis matches, find that suspense and surprise increases viewership. The effects of surprise are roughly 2 to 5 times larger than suspense, and the effects of both are larger in the latter half of the game. [Buraimo et al. \(2020\)](#) find that, in 540 British football matches, suspense and surprise increase viewership. In contrast to [Bizzozero, Flepp and Franck \(2016\)](#), they find that the effects of suspense are much larger than surprise and they do not see meaningful differences between the first and second halves. They also define the concept of “shock,” a measure of deviation from one’s prior. The “shock” at time t , calculated as the Euclidean distance between the beliefs at the beginning of the match and beliefs at time t , has an effect, though small. [Kaplan \(2020\)](#) examines the effect of player skill, suspense, and surprise in 477 NBA basketball games. He also finds positive significant effects from suspense and surprise on viewership. [Simonov, Ursu and Zheng \(2020\)](#) estimate the effects of suspense and surprise in e-sports. Using a sample of 104 professional games in a Counterstrike GO tournament and observing viewership on Twitch.TV channels, they find positive effects of suspense but little effect of surprise on viewership. Finally, [Liu, Shum and Uetake \(2020\)](#) turn to professional Japanese baseball. The unique advantage of their data is that they have eye-gaze metrics on their viewers, so they know whether the viewers are attentive or not. Using 877 games they estimate a statistical model and apply it to 41 games with eye-gaze data. They find that viewers attend more from suspense than from surprise. However, commercials are more highly attended during surprising moments than during suspenseful moments.

There are several important differences between our paper and the recent empirical literature on suspense and surprise in professional sports and e-sports. Rather than using third-party viewership as an outcome, we use players’ continued engagement as our outcome. This has three implications. First, unlike the viewership studies, we know who is rooting for whom – players are rooting for themselves. This allows us to examine directional (up vs. down) effects of beliefs on the outcome. Specifically, we look at the effect of trailing or leading on subsequent game-play. Second, the outcome is not a contemporaneous response to what is happening at the current moment. Instead, the outcome at the end of the game is based on the entire belief-path that transpired. Players make an ex-post evaluation and then decide to engage or walk away. In this sense, continued game-play is an appropriate metric for comparing matches to each other, as players are evaluating matches as a whole. Ex-post evaluative outcomes are arguably a better metric for game design as ultimately players and viewers decide to engage with subsequent matches based on their experiences with past matches. In contrast, viewership is an appropriate metric for within-match comparisons, as viewers tune-in or -out based on exciting moments in the current match.⁶ Third, some players win and

⁶While contemporaneous outcomes are useful metrics for game design, there are serious causes for concern. For example, suppose a player breaks a rule and the referee permits the behavior. This is a shocking event that may retain much viewership. An econometrician may then recommend that players break more rules and that referees ignore more violations. But a game without rules is probably boring to watch, and so the econometrician’s advice may lower total viewership. In contrast, if the econometrician used an ex-post outcome such as subsequent viewership of another match, this outcome would likely capture growing dissatisfaction with the game.

some players lose. The ex post evaluation of the probability paths may be colored by the final outcome. We can therefore differentiate the effect of beliefs on losers and winners in our analyses.

Another difference from the literature, with the exception of [Simonov, Ursu and Zheng \(2020\)](#), is that we use video games instead of traditional sports. Video games are understudied in economics relative to their market size. However, due to their digital nature, the opportunities for high-quality data are better. Our 2.8 million matches represent roughly one thousand times more matches than all of the previous studies on suspense and surprise combined, and about an order of magnitude more matches than the sum of all Major League Baseball, National Football League, and National Basketball Association matches through their respective histories.

The literature on reference-dependent preferences in games is thinner but present. [Coates, Humphreys and Zhou \(2014\)](#) use a model of reference-dependent preferences to motivate the uncertainty of outcome hypothesis, which states that greater outcome uncertainty leads to greater attendance at sports matches. They find evidence for reference dependence using Major League Baseball data. [Pope and Schweitzer \(2011\)](#) show that professional golfers are loss averse. [Bartling, Brandes and Schunk \(2015\)](#) show that professional soccer players breach the rules of the game, measured by the referee’s assignment of cards, significantly more often if their teams are behind the expected match outcome.

There are a couple of related papers that connect the psychological concept of flow to games. [Acland \(2020\)](#) uses panel data in an online word game to classify players by their relative skill and the relative challenge that they seek. He classifies players as flow-players if they have high-relative skill and seek high-relative challenge. He finds that flow-players are the most likely to use the in-game commitment device to limit play. [Abuhamdeh and Csikszentmihalyi \(2012\)](#) studied flow in online chess games. Because flow theory states that flow occurs when high skill is met with high challenge, they hypothesized an inverted-U shape pattern of enjoyment as a function of one’s relative skill over one’s opponent. They found this inverted-U pattern: player enjoyment was maximized when facing an opponent approximately 250 Elo points higher.⁷

2 Conceptual Framework

In this section, we present a theoretical motivation for our empirical analysis, largely following the existing literature. We define a player’s utility over features of the path that their beliefs about the probability of winning follow throughout a match and derive implications for players’ preferences over matches. In our model, players don’t choose their beliefs. Rather players’ beliefs are updated rationally throughout a match as the game generates more data indicating which team is likely to win.

Our key measure of a player’s engagement is whether the player chooses to continue playing another match after the current one finishes. Importantly, we assume positive adjacent complementarity ([Becker and Murphy, 1988](#)), which means that when players complete a higher-utility

⁷Elo is a rating metric used in chess.

match it increases demand for another match (i.e., it increases engagement). We briefly discuss five models: anticipatory utility, prospect theory, suspense and surprise, expectations-based reference dependence, and flow theory. After presenting a model for utility over matches, we explore the game-design implications.

2.1 Utility over Matches

2.1.1 Anticipatory Utility

We first consider a model of anticipatory utility following [Caplin and Leahy \(2001\)](#) which allows for a general class of anticipatory emotions. This model incorporates anticipatory emotions such as fear and anxiety and models how lotteries influence these emotions.

When we think of any feelings of anticipatory utility or anxiety, we posit that both winners and losers prefer the higher probability of winning throughout the game. Formally, we define $w \in \{0, 1\}$ be the outcome of a single *match*, where 1 corresponds to a win and 0 corresponds to a loss. Let $p = \{p_t\}_{t=0}^T$, where $p_t := \mathbb{E}[w|h_t]$ denotes the probability of winning given the history h_t of game states up to t . Then p is a probability path that represents rational beliefs about the probability of victory at each minute. Note that $p_T = w$, since the outcome is realized in the last minute of the match.

We define the following function that characterizes the anticipatory utility received by a player from a match that followed the probability path p .

$$U^{AU}(p) := \sum_{t=0}^T v(p_t), \quad (1)$$

where $\frac{dv(p_t)}{dp_t} > 0$.

Proposition 1 (Anticipatory Utility or Anxiety). *Under U^{AU} , players would prefer to think they are winning, i.e.,*

$$\frac{dU^{AU}}{d\bar{p}} > 0.$$

Proof. The sum of increasing functions is increasing in the average of the argument, i.e.,

$$\frac{dU^{AU}}{d\bar{p}} = \frac{d \sum_{t=0}^T v(p_t)}{d \sum_{t=0}^T \frac{1}{T} p_t} = \frac{d \sum_{t=0}^T v(p_t) dp_t}{d \sum_{t=0}^T \frac{1}{T} p_t dp_t} = \underbrace{\frac{d \sum_{t=0}^T v(p_t)}{dp_t}}_{>0} \underbrace{\left(\frac{d \sum_{t=0}^T \frac{1}{T} p_t}{dp_t} \right)^{-1}}_{>0} > 0.$$

□

2.1.2 Prospect Theory

Using the same setup, we now assume that the player derives utility at the end of the match based on how the realized outcome (w) compares to a reference point formed throughout the match. Specifically, she compares the utility $z > 0$ from winning and 0 from losing to a reference point $\bar{p} := \frac{1}{T} \sum_t p_t$, the average probability of winning she perceived throughout the match. In other words, the player compares the realized outcome of the match to the average of what they thought was going to be the outcome throughout the match. The prospect theory (PT) utility from the realization of the match p is:

$$U^{PT}(p) := v(zw - \bar{p}z) = v((w - \bar{p})z), \quad (2)$$

where $v(x)$ equals ηx (with $\eta > 0$) for $x \geq 0$ and $\eta\lambda x$ (with $\lambda > 1$) for $x < 0$.

Proposition 2 (Comparative statics with respect to the reference point). *Under U^{PT} and $\lambda > 1$, players would prefer to think they were losing, especially when they are losing; i.e.,*

$$\left. \frac{dU^{PT}}{d\bar{p}} \right|_{w=0} < \left. \frac{dU^{PT}}{d\bar{p}} \right|_{w=1} < 0$$

Proof. The result follows from the fact that $1 - \bar{p} \geq 0 \geq -\bar{p}$. Formally:

$$\frac{dU^{PT}}{d\bar{p}} = -zv'((w - \bar{p})z) = \begin{cases} -\eta\lambda z & \text{if } w = 0 \\ -\eta z & \text{if } w = 1, \end{cases} \quad (3)$$

where $-\eta\lambda z < -\eta z$ because $\eta, z > 0$ and $\lambda > 1$. □

2.1.3 Suspense and Surprise

Players derive utility from suspense and surprise (SS) throughout the game, defined as:

$$\begin{aligned} surprise_t &:= |p_t - p_{t-1}| \\ suspense_t &:= \sqrt{\mathbb{E}[(p_{t+1} - p_t)^2]} \end{aligned}$$

We assume that players collect minute-by-minute utility from suspense and surprise as follows:

$$u_t^{SS} = \beta_1^{surp} surprise_t + \beta_1^{susp} suspense_t.$$

Then, the relevant utility of the match is the average of u_t^{SS} over the game:

$$U^{SS}(p) := \overline{\beta_1^{surp} surprise_t} + \overline{\beta_1^{susp} suspense_t}.$$

Definition 1 (Preference for surprise). Players have a preference for surprise when $\beta_1^{surp} > 0$.

Definition 2 (Preference for suspense). Players have a preference for suspense when $\beta_1^{susp} > 0$.

2.1.4 Expectations-Based Reference Dependence

We now assume that players derive utility from their beliefs about winning p_t throughout the game relative to an expectations-based reference point as in [Kőszegi and Rabin \(2006, 2007, 2009\)](#). More specifically, the agent derives gain-loss utility in each period t from being positively or negatively surprised relative to his or her reference point. The reference point updates throughout the game. Overall utility equals:

$$U^{KR}(p) = \sum_{t=0}^T \gamma^t \int \int n(u(c) - u(r)) dF_{p_{t+1}}(c) dF_{p_t}(r)$$

where $n(x)$ equals ηx for $x > 0$ and $\eta\lambda$ for $x \leq 0$.

Proposition 3 (Preference for clumped versus smooth). *If $\gamma^t \approx \gamma^{t+1} \approx 1$ and t is small, then $U(\tilde{p}) > U(p)$ if the belief-path \tilde{p} is more clumped than p , i.e., for at least one t , $t + 1$, and $t + 2$ when $p_{t+2} = \tilde{p}_{t+2}$, $p_t = \tilde{p}_t$, $p_{t+1} \neq \tilde{p}_{t+1}$, and $\tilde{p}_{t+1} = \tilde{p}_t$, then $U(\tilde{p}) > U(p)$.*

The proof is in [Appendix A](#). In turn, note that σ_t in the proof represents the volatility of p_t , i.e., $\sigma_t = \frac{1}{N} \sum (p_t - p_{t-1})^2$, which is proportional to the squared surprise, i.e., *surprise* _{t} ² of the game. The expectations-based reference-dependent agent thus prefers games with higher *surprise* _{t} ².

2.1.5 Flow Theory

The final theory considering beliefs at specific moments is the theory of flow ([Csikszentmihalyi and Larson, 2014](#)). When a person is engaged in an activity that challenges but does not overwhelm their skill, with proximate goals and immediate feedback, the person may enter a flow state. A flow state is characterized by intense concentration on the present moment, a loss of self-awareness, a distortion of the sense of time, and intrinsic reward. Video games are highly sophisticated fine-tuned media that reliably deliver flow states to players ([Cowley et al., 2008](#)).

The primary relevance of flow to our analysis of beliefs is that flow states are induced only when the challenge matches a person’s skill level. The probability of winning at the start of the game (p_0) reflects only the challenge level as there has not yet been a record of performance. Variation in the initial conditions of a match creates variation in the ex ante chances of winning centered at 50%, but with non-negligible variance that is plausibly random.⁸⁹

Definition 2.1 (Flow). *Enjoyment has a unique interior maximum as a function of p_0 .*

⁸There are two sources of variation prior to the start of the match. The first is player skill level, which is not observable to the players. The matching algorithm randomly assembles teams so that the skill is balanced on average, but there will be variation as all team members will not be at the same skill level. Second, after teams have been assembled but before the clock starts, players choose a champion, each with different abilities, and a role (this is observable to the players). The designers balance each game to make champions equally powerful but the balance is never perfect, and champions may have strategic advantages or disadvantages depending on the champion selection of the other players.

⁹Our interpretation of the effect of p_0 in the context of flow is that this represents the ex ante challenge of the match. Our use of the p_t in the context of the other theories is as a proxy for player beliefs.

2.2 Empirical Analogy

Taking all five utility functions into account, the utility of belief-path p is:

$$U(p) := U^{AU}(p) + U^{PT}(p) + U^{SS}(p) + U^{KR}(p) + U^F(p).$$

We can then test the key predictions of all utility functions in the following specification:

$$\begin{aligned} ContinuedPlay = & (\beta_1^{AU} + \beta_1^{PT})\bar{p} + \beta_1^{surp}\overline{surprise}_t + (\beta_2^{surp} + \beta_2^{KR})\overline{surprise}_t^2 \\ & + \beta_1^{susp}\overline{suspense}_t + \beta_1^F p_0 + \beta_2^F p_0^2 \end{aligned} \quad (4)$$

We estimate this specification separately for winners and losers to allow the coefficients to differ conditional on the outcome.

We are unable to identify β_1^{AU} and β_1^{PT} individually, but we are able to identify the sum of the two. Because anticipatory utility predicts a positive coefficient estimate of \bar{p} and prospect theory predicts a negative one, the sign on the coefficient indicates the comparative relevance of the two theories. A positive coefficient implies that anticipatory utility dominates, and a negative coefficient implies that prospect theory dominates. If $\beta_1^{AU} = 0$, then the regression specification identifies η , the relative importance of gain-loss utility. Moreover, by comparing the coefficients on \bar{p} for winners and losers the specification also identifies λ , the degree of loss aversion.

A preference for surprise and suspense is tested by estimating the coefficients on the match-averages of these variables, as well as their squared counterparts. The squared term tests whether there is a preference for clumped surprise. That said, we are unable to separately identify β_2^{surp} and β_2^{KR} , though again we can identify the sum of the two. The latter is assumed to be weakly positive. Thus if the sum of the two are negative, we can infer $\beta_2^{surp} < 0$ and $\beta_2^{surp} < -\beta_2^{KR}$, i.e., if the coefficient on the average of squared surprise is negative, we can infer that whatever expectations-based reference-dependent preferences exist for clumpy information, they are dominated by a preference for smooth surprise.

Preferences for flow are captured by the coefficients on p_0 and p_0^2 . Flow theory predicts that $\beta_1^F > 0$ and $\beta_2^F < 0$, which would generate single-peaked preferences for “challenge.”

The belief-path is high-dimensional, so there are certainly other specifications one can run based on many possible features of the belief-path. These are the ones we chose, motivated by theory.¹⁰

¹⁰We suspect that there is an emotional difference between positive surprises, that is shifts in beliefs towards the desired outcome (win), and negative surprises, shifts in belief to the undesired outcome (loss). Moment to moment, these shifts may be experienced very differently. However, as we aggregate over an entire match, the total positive surprise minus the total negative surprise must always equal $-p_0$, on a loss, or $1 - p_0$, on a win. Thus at the match level, conditional on win status, average positive surprise and average negative surprise are linear combinations of p_0 and average surprise. Thus we cannot add them to our specifications.

2.3 Theory of Game Design

There are two levels of empirical analysis we consider in this paper. The first is the match-level analysis in which a player experiences a belief-path and then makes a continuation choice. The second level is what we call the game-level analysis, which is relevant for the game designers. Belief-paths are not born in isolation. They come from a distribution of belief-paths generated by the game. In other words, we can think of a game, G , as a probability measure over the set of possible belief-paths. We refer to a single draw of a belief path from the game as a match. The aim of this section is to provide some intuition and theorems for the game design problem. We draw on Lemma 1 from [Ely, Frankel and Kamenica \(2015\)](#) which allows us to focus on G without having to consider the information policy that induces G .

The analogues to all previous propositions and definitions of anticipatory utility, suspense and surprise, reference-dependence, and flow extend right away. For prospect theory, the game design problem provides a prediction for early versus late resolution of uncertainty.

2.3.1 Prospect Theory

We first consider the game design problem assuming only the prospect theory utility from the realized outcome. The prospect theory utility that a player expects from game G is:

$$\tilde{U}^{PT}(G) := \int U^P(p) dG(p) = \int v((w - \bar{p}; \eta, \gamma)w) dG(p).$$

Under this utility function, the player forms expectation \bar{p} throughout the match and receives prospect theory utility based on the outcome relative to their expectation. Proposition 2 tells us that, given the shape of v , players would always prefer to have thought they were losing throughout the match—i.e. to have a lower \bar{p} .

From a game design perspective, an important question is what this implies about whether a game should reveal information about who is likely to win earlier or later in the match. To explore this, we consider a perturbation of a game that reveals some information slightly earlier and ask whether players would prefer the perturbed game.

For any game G , define G^t to be the same game as G , except the information from t and $t + 1$ are both revealed at t . In other words, the game G^t is almost the same game as G , except it reveals the information slightly earlier, from $t + 1$ one minute earlier at t .¹¹ Formally, for any belief-path p generated by G , the game G^t generates belief-path p' so that $p'_\tau = p_\tau$ for all $\tau \neq t$ and $p'_t = p_{t+1}$.

Comparing preferences for G^t relative to G provides insights into whether players prefer games with early or late information revelation. Players who prefer G^t to G prefer earlier revelation, whereas players that prefer G to G^t prefer later revelation. Proposition 4 shows that losers prefer games that reveal information earlier, and winners prefer games that reveal information later.

¹¹Note that this makes minute $t + 1$ a “dead minute” with no new information in matches generated by G^t .

Proposition 4 (Preference for early versus late resolution). *For any game G and perturbation G^t , losers prefer G^t —i.e. that information be revealed earlier—and winners prefer G —i.e. that information be revealed later.*

The proof is in Appendix A. Intuitively, the proof follows from the fact that earlier revelations will tend to move \bar{p} towards the realized outcome. This will be preferred by the eventual loser, since $w - \bar{p}$ will tend to be less negative if the game generates matches with smaller \bar{p} . For the eventual winner, on the other hand, moving towards the realized outcome means that $w - \bar{p}$ will be less positive.

2.3.2 Suspense and Surprise

Ely, Frankel and Kamenica (2015) provide a thorough treatment of the game design problem in Section IV and V of their paper. Their Lemma 2 shows that the suspense-maximizing information policy—i.e., the game that maximizes suspense—fully reveals the outcome by the end of the match. Moreover, it shows that fixing game length and p_0 , any game that fully reveals the outcome has a fixed budget of suspense. One consequence of this is that if preferences for suspense are linear in total suspense, then all games with the same distribution of p_0 and match length provide the same expected utility from suspense. However, if players have concave preferences in minute-by-minute suspense, a suspense-optimal game will smooth the suspense over the match. In fact, the authors show how to construct a suspense-optimal information policy (i.e. game) that holds suspense fixed throughout the match. Ely, Frankel and Kamenica (2015) likewise derive the properties of surprise-optimal information policies for a more limited set of cases. These policies are less straightforward than suspense-optimal policies and the results mostly demonstrate that little is guaranteed for surprise-optimal policies. For example, they may reveal the winner before the final period or not at all, uncertainty may increase or decrease over time, and realized surprise is stochastic and varies over time. We refer the reader to Ely, Frankel and Kamenica (2015) for the full treatment of suspense- and surprise-optimal policies.

3 Data and Empirical Strategy

3.1 League of Legends

League of Legends, developed by Riot Games, Inc., has been one of the most popular PC games since its release in 2009. In July of 2023, the game had 151 million users who played at least once per month.¹² The player base is approximately 82% male, with 77% of players between 16 and 34 years old.¹³

In *League of Legends*, players select a “champion” from a menu of over one hundred options, such as knights, wizards, ninjas, pirates, monsters, and robots before a match begins. Each champion

¹²Estimates come from <https://activeplayer.io/league-of-legends/>.

¹³Estimates from <https://www.statista.com/statistics/1018232/league-of-legends-fans-by-gender/> and <https://www.statista.com/statistics/1018224/league-of-legends-fans-by-age-usa/>.

has a unique suite of abilities and lends itself to a particular play-style. Players manipulate their champion in real-time and form teams of five with the ultimate objective of destroying the opposing team’s base. The teams’ bases are located at opposite corners of a diamond-shaped map, and there are three “lanes” that connect the bases to each other. The lanes are guarded by static defenses called “towers” that attack enemy champions advancing into the defending team’s territory.

Players compete for resources to strengthen their champions’ offensive and defensive capabilities to both overcome the towers and opposing champions. The two most relevant resources are gold and experience points (XP). During a match, players spend gold to purchase items to enhance their champions’ abilities and customize their strategies to the current match. When a player earns enough XP they “level up” which enhances his or her champion’s overall abilities and grants the player a skill point to enhance one aspect of their champion’s abilities strategically. Champions that are killed get a time-out before they reappear at their home base. Matches typically last around 30 minutes.¹⁴ As the match progresses, champions become more powerful and time-outs for dying become longer making errors more costly. As a result, matches rarely exceed 45 minutes.

When players enter the queue for a match, the matchmaker algorithm randomly assigns players to teams with the objective to make each team’s total skill level approximately equal. Each player has a hidden rating, generated from a proprietary algorithm using the player’s history of performance. Players do not know their ratings but they are present in our data set.

3.2 Data

We obtained a proprietary data set from Riot Games that contains information on 2,882,101 games from the North American server from March 1, 2018 to August 1, 2018. The data are observational and there were no experimental interventions conducted by the researchers. In addition to game-specific data such as date and time, we also observe information for each player such as their chosen champion and rating. We can observe detailed minute-by-minute game-state variables at the player-level. For example, we observe each player’s minute-by-minute accumulation of gold and XP and their position on the map. As a result, the data contain 28,009,918 observations at the player-game level, approximately 97 million observations at the minute-game level, and approximately 960 million observations at the player-minute-game level. The data constitute a repeated cross-section, as players may appear multiple times in the data. However, we are unable to construct a panel dataset, as we do not have player account identifiers. Player records do not contain any personally-identifying information.

3.3 Estimating the Probability of Winning

A necessary ingredient for our analysis is each team’s probability of winning at various points throughout the game. In this subsection, we describe the process by which we estimate these win probabilities at every minute during each game. The estimated minute-by-minute probability

¹⁴In our sample we include only matches between 20 to 45 minutes in length, which leaves us with an average game length of 30.65 minutes.

of winning is then used to construct our variables of interest, such as the average probability of winning, average suspense, and average surprise. The specific variables used to estimate the probability of winning are summarized in Table B1 in Appendix B. At each minute, we have a total of 133 variables describing the game state. Players have access to numerous game-state variables including the best predictors of victory and so our statistical model should approximate the rational expectations of an experienced player.

We denote by s_{mt} the state of match m at minute t . There are a large number of variables in s_{mt} , including the gold, experience, kills, and deaths accumulated by each player by time t , as well as each player’s position at t and his or her chosen champion and role. At minute t during the match, beliefs about the probability of winning are formed based on the history of game states through t , which we denote by $h_{mt} := (s_{m0}, s_{m1}, \dots, s_{mt})$. Formally, the probability of team j winning given the history observed at time t is:

$$p_{mjt} := \mathbb{E}[y_{mj} | h_{mt}], \tag{5}$$

where y_{mj} is an indicator for team j winning match m .

The probability of winning at t depends on the entire history (h_{mt}) rather than just the current state (s_{mt}) of the game. The full history captures the strategies and skills of each team, leading to a better model of p_{mjt} . Additionally, this flexibility allows the probability of winning to account for how the game evolves. For example, it may be that “momentum” in a team’s performance is informative about its probability of winning (Green and Zwiebel, 2018). We use machine learning—specifically Microsoft’s LightGBM implementation of gradient-boosted decision trees—to empirically approximate p_{mjt} . See Appendix B for details.

3.4 Variable Construction

The estimated probability of winning, \hat{p}_t , forms the basis of our analysis. We construct additional key variables using \hat{p}_t , as outlined in Section 2. Surprise is simply the absolute change in the probability of winning over a minute. Suspense is the expected absolute change in the probability of winning. It is not a function of \hat{p}_t and \hat{p}_{t+1} but a function of \hat{p}_t and $E[\hat{p}_{t+1} | h_{mt}]$. Suspense, therefore, needs to be estimated. We estimate $\sqrt{E[(\hat{p}_{t+1} - \hat{p}_t)^2]}$ in a second stage using a machine-learning approach analogous to that used to estimate \hat{p}_t above. Details of the estimation are in Appendix B.2.

3.5 Empirical Strategy

Our primary goal is to estimate the effect of various attributes of the probability path on engagement. We measure engagement at player level by the decision to continue play for another match. Specifically, we use a binary variable indicating whether or not the individual played another match

within 60 minutes of completing the current one. We estimate regressions of the form:

$$y_{mji} = \beta_0 + \beta_1 q_{mj} + \mathbf{x}'_{mji} \lambda + \epsilon_{mji} \tag{6}$$

where y_{mji} is the measure of continued play for player i on team j after match m . Here, q_{mj} represents the characteristics of the probability path \vec{p}_{mj} that theory posits may affect engagement. Formally, $q_{mj} = q(\vec{p}_{mj})$, where the function $q(\cdot)$ computes a vector of theory-relevant characteristics from probability path \vec{p}_{mj} . Our coefficient of interest are the vector β_1 , the marginal effects of the measures on continued game-play. There are features of the game that may have a direct effect on game-play; \mathbf{x}'_{mji} represents a vector of control variables. Unfortunately, as mentioned above, we cannot track players over time. However, we observe several measures of previous game-play, so we can control for the number of recent games played. The controls in our continuation of play estimations include linear controls for games played in the previous 2 hours, games played in the previous 2 hours with the same champion, log minutes from the previous game, log minutes from the previous game using the same champion, quadratic controls for player skill rating, and fixed effects for date, day of the week, hour of the day, team, role, champion, and game length. Because attributes of the probability path q_{mj} may have different effects for losers and winners, we estimate our regression models separately for the two groups.

3.6 Summary Statistics

Figure 1 shows the belief-paths in a small sample of matches (wins only). The bar graphs below the belief-paths show the minute-to-minute values of surprise and suspense. We visualize the variation and distribution of the belief-paths in Figure 2 (wins only). The figure shows the belief-paths of 1,000 randomly selected matches from the winner’s perspective. We note that there is a lot that can happen in these matches. Victory may be rapid and direct or there could be a long circuitous path.

Figure 3 shows a scatter plot of surprise and suspense, as a comparison to the games featured in Ely, Frankel and Kamenica (2015): tennis, soccer, blackjack, the Clinton-Obama primary, a surprise-optimum process, and a suspense-optimal process. Panel (a) shows surprise and suspense at the minute level and panel (b) shows average surprise and average suspense at the match level. The correlations between the two are 0.27 and 0.87, respectively.

Table 1 shows summary statistics of our variables of interest. We show the statistics from the perspective of winners. The full sample is double, 2.8 million matches. We can see that, conditional on winning, the average probability of winning is 0.70 over the course of the game, whereas the initial probability of winning is, by construction, close to 0.5. The average match is 30.65 minutes long. The average frequency of playing another match within the next 60 minutes (continued play) is 59.64% for winners and 58.76% for losers.

In Table 2, we perform our first estimation exercise. We regress continued play on win, weekend, and prime-time (defined as the time between 1:00pm and 1:00am CST) to establish benchmark

Figure 1: Sample games showing \hat{p} -path, surprise and suspense, only winners.

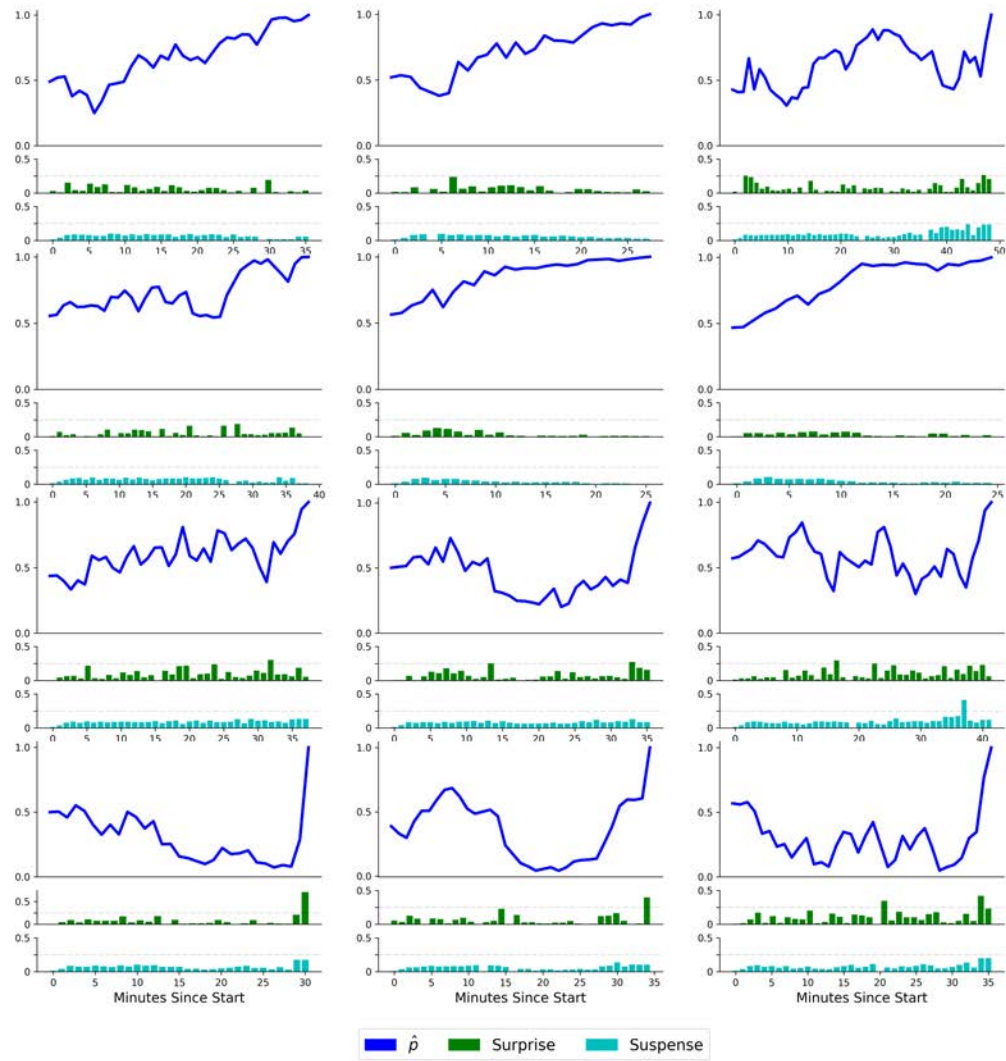


Figure 2: The path of \hat{p} , surprise, and suspense for 1,000 different matches, only winners.

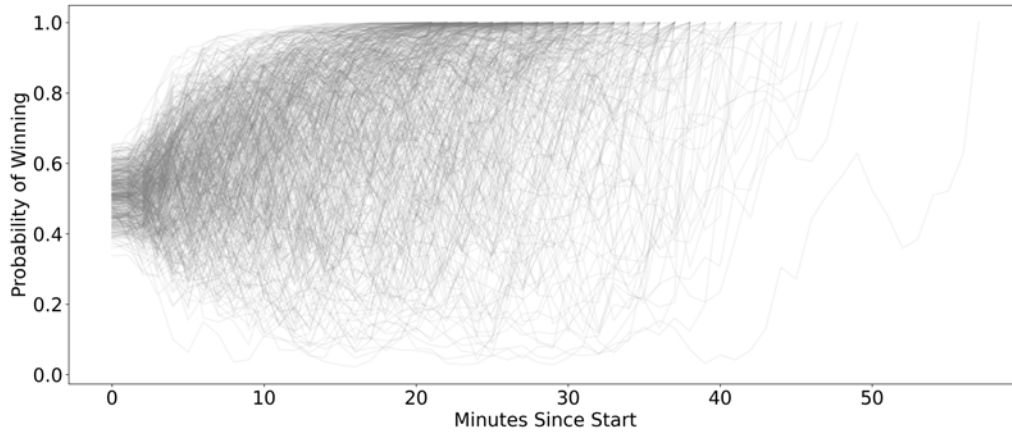
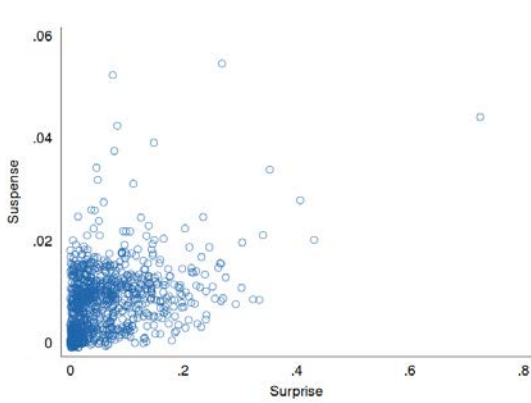
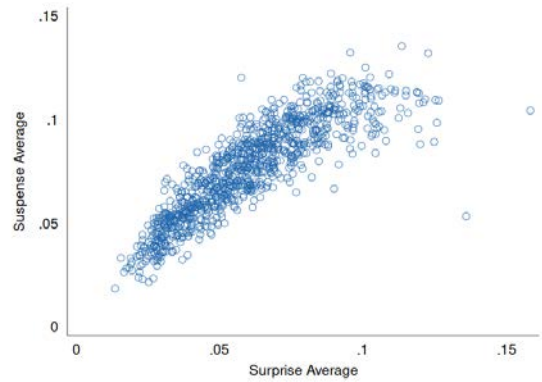


Figure 3: The association between surprise and suspense.



(a) Random selection of 1,000 minutes.



(b) Random selection of 1,000 matches, with surprise and suspense averaged at the match level.

Notes: The correlation between surprise and suspense, minute by minute, is 0.27, as visualized in panel (a). The correlation between surprise and average suspense, averaged at the match level, is 0.87, as visualized in panel (b).

Table 1: Summary Statistics

	Winners		Losers	
	Mean	SD	Mean	SD
\hat{p} Average	0.70	0.14	0.30	0.14
\hat{p}_0	0.51	0.06	0.49	0.06
Surprise Average	0.06	0.02	0.06	0.02
Surprise ² Average	0.01	0.01	0.01	0.01
Suspense Average	0.08	0.02	0.08	0.02
Game Minutes	30.65	6.03	30.65	6.03
Continued Play (60 min.)	59.64	49.06	58.76	49.23
N	1.4e+07		1.4e+07	

effects on continued game-play with the aforementioned control variables. In all of our estimations continued play is multiplied by 100 and thus coefficients represent percentage points. Column (1) shows that winning increases the probability of playing another game by 0.88 percentage points. Weekends have an even stronger effect, increasing the probability by 2.12 percentage points. Playing during primetime also plays an important role, increasing the probability by 1.68 percentage points. Though these effects appear modest, they belie that much variation is at the individual level. Some people play very little and others may play in long streaks. Unfortunately, we do not have the ability to control for this individual-level variation. Indeed the R-squared of the regression is only 0.012. In Column (2), we include pairwise interactions of the three variables to test for the sensitivity of the win-loss gap when player's time is expected to be less constrained. The win-loss gap increases on the weekend and during primetime hours, indicating that players' game-play is more elastic to the outcome of the match when they have more free time.

Table 2, Column 1 establishes the size of the win-loss gap for comparative purposes. Indeed, the effects of weekend and time of day are stark: the win-loss gap is not distinguishable from zero on weekdays outside of primetime, and it is approximately double the average effect (1.68 percentage points) on weekends during primetime. This suggests that constraints are probably the dominant factor in continued game-play, but the high significance of the win-loss gap shows that the game's outcome affects future game-play. We expect that the win-loss gap is an upper bound for the effect sizes we would expect from attributes of the belief path. Winning is the objective of the game and the most salient attribute of a completed match. Other features of the match, like lagging or surprise, are expected to be secondary.

Table 2: Benchmark Effects on Game-play — Weekend, Hour, and Win-Loss Gap

	(1)	(2)
Win	0.88 (0.02)	-0.03 (0.05)
Weekend	2.12 (0.03)	0.93 (0.05)
Primetime	1.68 (0.03)	0.89 (0.04)
Win \times Weekend		0.84 (0.05)
Win \times Primetime		0.88 (0.05)
Weekend \times Primetime		1.05 (0.06)
Mean of DV	59.20	59.20
SD of DV	49.15	49.15
R2	0.01	0.01
Clusters	2,870,490	2,870,490
N	27,898,907	27,898,907

Notes: All specifications are linear probability models using OLS with robust standard errors clustered at the match-team level. Primetime is defined as play between the hours of 7:00 pm to 7:00 am CST, where the empirical density of games is the highest. All specifications include linear controls for games played in the previous 2 hours, games played in the previous 2 hours with the same champion, log minutes from the previous game, log minutes from the previous game using the same champion, quadratic controls for player skill rating, and fixed effects for date, day of week, hour of day, team, role, champion, and game length.

4 Match-level Analysis

4.1 Main Results

Table 3 shows the results from an ordinary least squares regression of continued play on our features of the belief-path, following equation (4), using our aforementioned vector of control variables.¹⁵ Columns (1)-(3) use the sample of team matches that lost, and Columns (4)-(6) use the sample of team matches that won.

Table 3: Continued play as a function of time and \hat{p} , by winner.

	Losers			Winners		
	(1)	(2)	(3)	(4)	(5)	(6)
\hat{p} Average [†]	-9.30 (0.21)	-9.60 (0.22)	-9.61 (0.22)	-6.30 (0.21)	-6.55 (0.22)	-6.62 (0.18)
Surprise Average [†]	-6.69 (1.57)	-15.79 (2.45)	-15.80 (2.45)	-14.89 (1.59)	-7.16 (2.47)	
Surprise ² Average [†]		26.98 (5.66)	26.90 (5.66)		-22.91 (5.69)	-34.91 (3.54)
Suspense Average [†]	18.31 (1.69)	22.41 (1.88)	22.53 (1.88)	8.12 (1.71)	4.64 (1.90)	
\hat{p}_0	10.74 (3.00)	10.90 (3.00)	4.63 (0.28)	29.03 (3.17)	28.98 (3.17)	29.08 (3.17)
\hat{p}_0^2	-6.31 (3.05)	-6.42 (3.05)		-25.30 (3.09)	-25.20 (3.09)	-25.43 (3.09)
Peak \hat{p}_0	0.85	0.85		0.57	0.57	0.57
Mean of DV	58.76	58.76	58.76	59.64	59.64	59.64
SD of DV	49	49	49	49	49	49
R2	.013	.013	.013	.012	.012	.012
Clusters	2,805,461	2,805,461	2,805,461	2,805,336	2,805,336	2,805,336
N	13,949,001	13,949,001	13,949,001	13,949,906	13,949,906	13,949,906

Notes: All specifications are OLS with robust standard errors clustered at the match-team level. Continued play is multiplied by 100 so that coefficients can be interpreted in percentage points. All specifications include linear controls for games played in the previous 2 hours, games played in the previous 2 hours with the same champion, log minutes from the previous game, log minutes from the previous game using the same champion, quadratic controls for player skill rating, and fixed effects for date, day of week, hour of day, team, role, champion, and game length.

[†] denotes that the variable definitions are averaged over all minutes, excluding minute zero.

In Columns (1) and (4) we regress continued play on average \hat{p} , average surprise, average suspense, and a quadratic specification in \hat{p}_0 . Columns (1) and (4) estimate the average effect of beliefs, surprise, suspense, and flow. The association between games played and \hat{p} is negative and highly significant for both winners and losers. This negative effect of average beliefs is consistent

¹⁵Note that our standard errors treat \hat{p} as data. Since \hat{p} are not directly observed but are estimated using machine learning, \hat{p} may suffer from at least some measurement error.

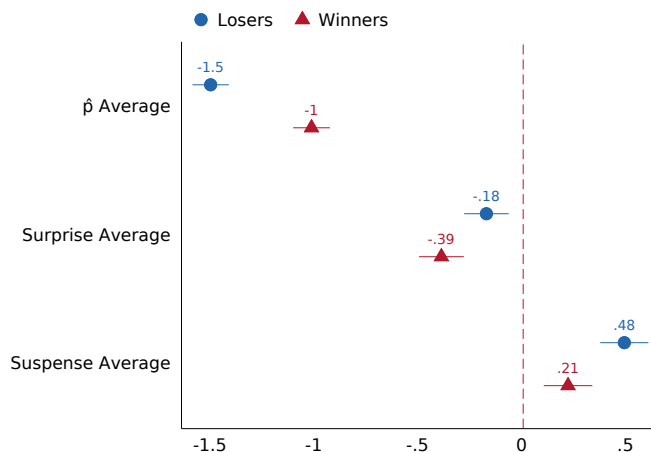
with prospect theory in Proposition 2 and inconsistent with anticipatory utility in Proposition 1. To the extent that there may be anticipatory utility, the effect of prospect theory appears to dominate. Lagging behind leads to lower reference points, increasing the enjoyment of a win and mitigating the pain of a defeat. There is a larger effect for losers – nearly twice the magnitude – as predicted in Proposition 2, induced by loss aversion. The difference between the coefficients on \hat{p} Average in Columns (1)-(3), versus (4)-(6) are significant at the $p < 10^{-6}$ level. These results are consistent with existing research suggesting the coefficient of loss aversion is approximately two (for example Tversky and Kahneman, 1992).

In Columns (1) and (4), we include surprise and suspense linearly to facilitate interpretation. Surprise reduces continued play for losers and winners, the opposite of a “preference for surprise” in Definition 1. Suspense, however, increases continued play for losers and winners consistent with a “preference for suspense” in Definition 2. The negative coefficient on surprise is twice as large for winners, whereas the positive coefficient on suspense is about twice as large for losers. Though average surprise and average suspense are correlated, they can vary within a match. From a game-design perspective, increasing suspense does not necessarily increase expected surprise. Keeping average surprise fixed and making expected surprise more clumpy will increase the suspense in a match. The results here suggest a preference for exactly that kind of information revelation.

For flow theory, in Definition 2.1 we posit that game-play will have a unique interior maximum as a function of \hat{p}_0 , which we operationalize as the challenge level of a match. We attempt to parametrically estimate this inverted-U pattern using a quadratic specification of \hat{p}_0 . Since we are using \hat{p}_0 directly in these specifications, \hat{p}_0 should not be implicitly included in the definitions of our other explanatory variables. For that reason the other variables are averaged over all minutes excluding minute zero. For losers, the coefficient on the linear term is statistically significant at $10^{-2} > p > 10^{-3}$, which is relatively noisy given the large sample size. The quadratic term is only marginally significant at $p = 0.04$. When we estimate the peak \hat{p}_0 , it’s at 0.85, which is essentially out of the support of \hat{p}_0 (it is about 5 standard deviations away from the mean). We can conclude that our hypothesis of a continued play being single-peaked in initial challenge is rejected for losers, as within the distribution of \hat{p}_0 games-played is strictly increasing in \hat{p}_0 . For winners, both the linear term and the quadratic terms are significant and the unique maximizing value \hat{p}_0 is 0.57, consistent with single-peaked preferences. This fits the description of flow theory. People want to be challenged, but not too much.

Figure 4 displays the coefficients from Table 3, Columns (1) and (4), but re-scaled and expressed as proportions of the win effect from Table 2. The figure shows that the most prominent effect is that of average \hat{p} . A standard deviation *decrease* in average \hat{p} for a losing team increases play as if the team won $1.5\times$ over. A standard deviation *decrease* in average \hat{p} for a winning team increases play by $1\times$ the win-loss gap. The effects of surprise and suspense are comparatively weaker but still meaningful. A standard deviation decrease in average surprise leads to a 0.18 and 0.39 fractions of the win-loss gap for losers and winners respectively. A standard deviation increase in average suspense leads to a 0.48 and 0.21 fractions of the win-loss gap for losers and winners respectively.

Figure 4: Standardized effects of \hat{p} Average, Surprise Average, and Suspense Average expressed as a proportion of the win-loss gap



Notes: Specifications come from Table 3, Columns (1) and (4). Explanatory variables are standardized. The outcome is continued play divided by the “win effect” as measured in Table 2. Therefore coefficients are the proportion of the win effect from a standard deviation increase in the explanatory variable. Lines show the 99% confidence interval.

As a form of visual non-parametric robustness, we plot continued play on the independent variables from the specification in Columns (1) and (4). In Figure 5, we regress continued play on all control variables from the specification and then plot the residuals as a function of time and \hat{p} . In the heat map, a darker color represents a lower chance of continued play. The color gets darker with the probability of winning for both losers and winners, indicating that low expectation is associated with playing *more* matches. The effect of \hat{p} appears to be qualitatively similar across time.

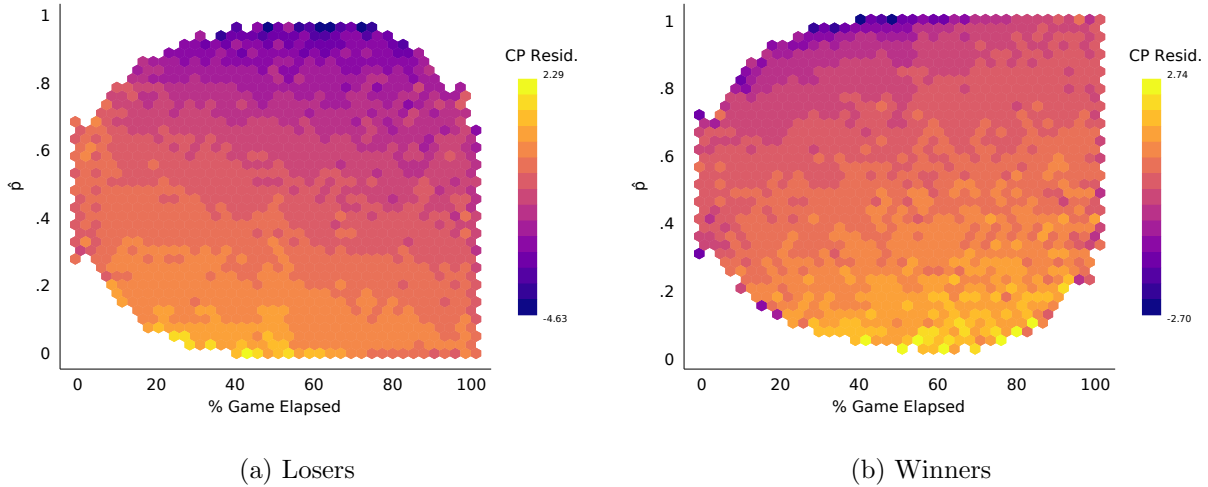
Figure 6 visualizes analogous relationships between continued play and surprise and continued play and suspense. We regress continued play on all the variables in the specification in Columns (1) and (4) excluding average surprise or average suspense, and we then binscatter the residuals on surprise and suspense respectively. The figure shows that continued play appears to be decreasing in average surprise and increasing in average suspense. The linear approximation appears well-suited for both average surprise and average suspense.

Figure 7 displays a binned scatterplot of continued play as a function of \hat{p}_0 , controlling for all other variables in Table 3, Columns (1) and (4).¹⁶ The relationship for winners appears to be increasing and flatten out, and then to possibly decrease. This is consistent with the single-peaked preferences we estimated in the quadratic model. In contrast, losers exhibit an increasing relationship in \hat{p}_0 .

Columns (2) and (5) introduce Surprise² Average to test whether players have a preference for clumpy or smooth surprise. If the coefficients on Surprise² Average are positive, it implies a preference for clumpy surprise. If, on the other hand, the coefficients are negative it implies a

¹⁶The averaged variables here exclude minute zero in order to not use controls partially derived from \hat{p}_0 .

Figure 5: Continued play as a function of time and \hat{p} , by winner.

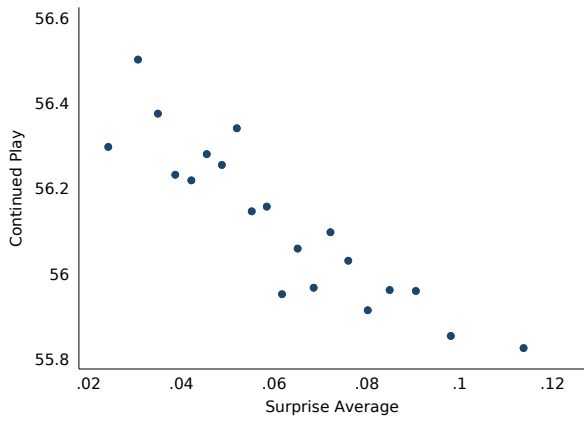


Notes: The continued play residuals are generated from the specification in Columns (1) and (4) of Table 3, excluding average \hat{p} . Cells that have less than 5,000 observations are left empty.

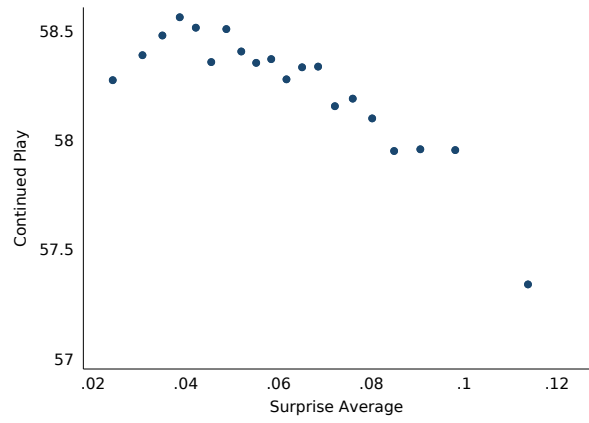
preference for smoothed surprise. For losers we find the coefficient on Surprise² Average is positive and significant. This implies that conditional on losing, players have a preference for clumpy surprise. For winners, in contrast, we find the coefficient on Surprise² Average is negative and significant, suggesting that they prefer surprise smoothed. Columns (2) and (5) include features of the belief path for which the coefficients are non-significant. In Section 5 we wish to use an estimated objective function for our structural analysis. In Columns (3) and (6) we iteratively eliminate the features of the belief-path for which the coefficients are not significant at the $p < 0.01$ level. This eliminates \hat{p}_0^2 for losers, and Surprise Average and Suspense Average for winners (Surprise² is included). All remaining coefficients are now highly significant ($p < 10^{-4}$) and qualitatively similar. Column (3) now makes clear that loser's continued play is strictly increasing in \hat{p}_0 . We use the coefficients in Columns (3) and (6) for our structural analysis.

The preference conditional on losing is consistent with expectations-based reference dependence (Kőszegi and Rabin, 2009) from Proposition 3, which predicts a preference for clumpy surprise. However, the preferences for winners seems to be the opposite. Winners both dislike surprise and prefer it smoothed out. What might explain this? We speculate that diminishing sensitivity might play a role. Most surprises will be good news, conditional on winning, and diminishing sensitivity suggests that many small gains are valued more than fewer equivalent large gains Thaler (1999).

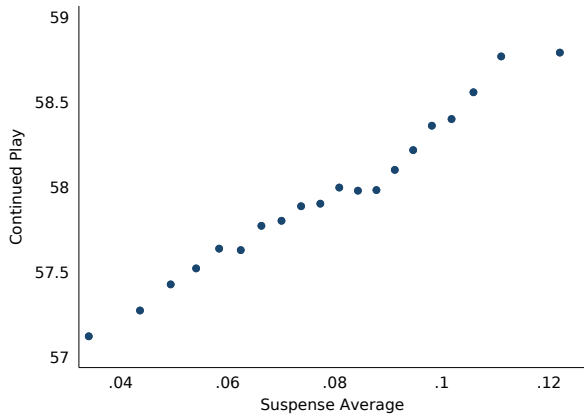
Figure 6: Continued play as a function of average surprise and as a function of average suspense, by winner.



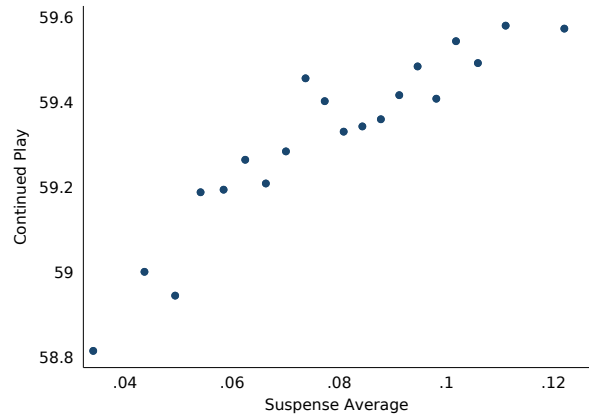
(a) Surprise Losers



(b) Surprise Winners

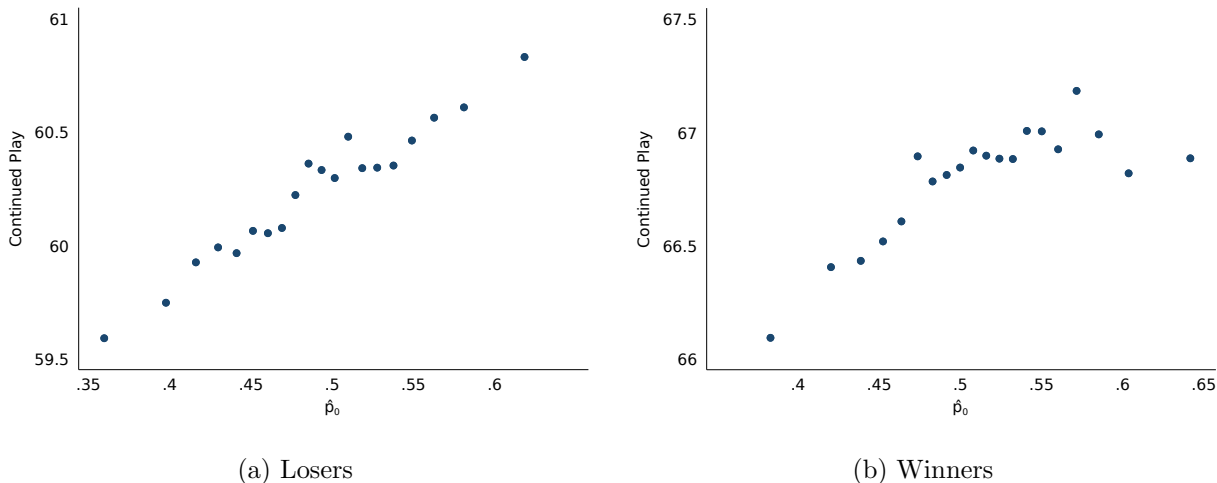


(c) Suspense Losers



(d) Suspense Winners

Figure 7: Continued play as a function of \hat{p}_0 , by winner



Notes: This binned scatterplot controls for the following variables averaged over all but the initial minute: \hat{p}_t , surprise, and suspense. It also controls for match-length fixed effects.

5 Structural Game Optimization

Our previous analysis provided insights into what makes a compelling match for players. However, the production of these belief-paths are subject to constraints. First, most obviously, as a competitive game, for every winner, there is a loser, and for every belief path, there is an inverted belief path experienced by the opposing team. Second, rational beliefs must obey the Martingale property; that is, the expected belief in the next period must equal the belief in the current period, $E_t[p_{t+1}] = p_t$. This has important implications for the distribution of feasible belief paths. For example, dramatic comebacks from the brink of defeat are necessarily uncommon, because having a belief on the verge of defeat – say a 90% chance of defeat – must lead to defeat 90% of the time. These constraints make optimizing the data-generating process that we call the “game” a very different task from comparing or ranking matches.

In this section, we move from match analysis to analyzing the game itself. The game can be modeled as a data-generating process that produces belief paths. Our estimates in Table 3 provide a straightforward way to evaluate expected engagement produced by a single match. We extend this to the game level by evaluating a game G based on the average engagement of the matches it produces:

$$U(G) := \int u(p)dG(p), \quad (7)$$

where $u(p)$ is the engagement of the match implied by Columns (3) and (6) of Table 3. We then seek to optimize the data-generating process for engagement in the local vicinity of the *League of Legends* data-generating process. We summarize this process below and provide additional details in Appendix C.

Our first step is to provide a representation of the *League of Legends* game which we define as

G_{LoL} . Empirically, G_{LoL} is easy to construct. To characterize the data-generating process, we must express the probability that p_t transitions to a belief p_{t+1} . We discretize the probability space into 20 interior bins (i.e., $(0, 0.05], (0.05, 0.1], \dots, (.95, 1)$) and 2 “absorbing state” bins (i.e. $\{0\}$ and $\{1\}$). This allows us to have a finite number of transition probabilities. For a given belief bin $b_{n,t}$ where $n \in \{0, \dots, 21\}$ at time t , there are 22 transition probabilities that represent the probability that $p_{t+1} \in b_{n,t+1}$. Each bin at time t has a 22-element vector of these transition probabilities. The transition probabilities are then simply conditional frequencies. For example, to compute the transition probability at minute 8 from bin $(0.65, 0.7]$ to bin $(0.4, 0.45]$ we simply take all matches that have $\hat{p}_8 \in (0.65, 0.7]$ and measure the fraction of those that have $\hat{p}_9 \in (0.4, 0.45]$.

Figure 8 displays the game *League of Legends*, G_{LoL} . Time is plotted on the horizontal axis, and the probability of winning is plotted on the vertical axis. The color indicates the degree of information revelation, as measured by expected surprise, for that cell, with more blue having more expected information revelation. To show the density of matches, we then overlay the plot with 100 random matches with jitter (random displacement so that multiple observations at the same point can be visualized). As one can see, greater information revelation occurs when \hat{p} is near 0.5, and information revelation, conditional on a bin, increases over time.

Now that we have estimated an objective function and have operationalized the game G , the next step is to choose the optimal game G^* from a subset of possible games \mathcal{G} :

$$G^* := \arg \max_{G \in \mathcal{G}} U(G). \tag{8}$$

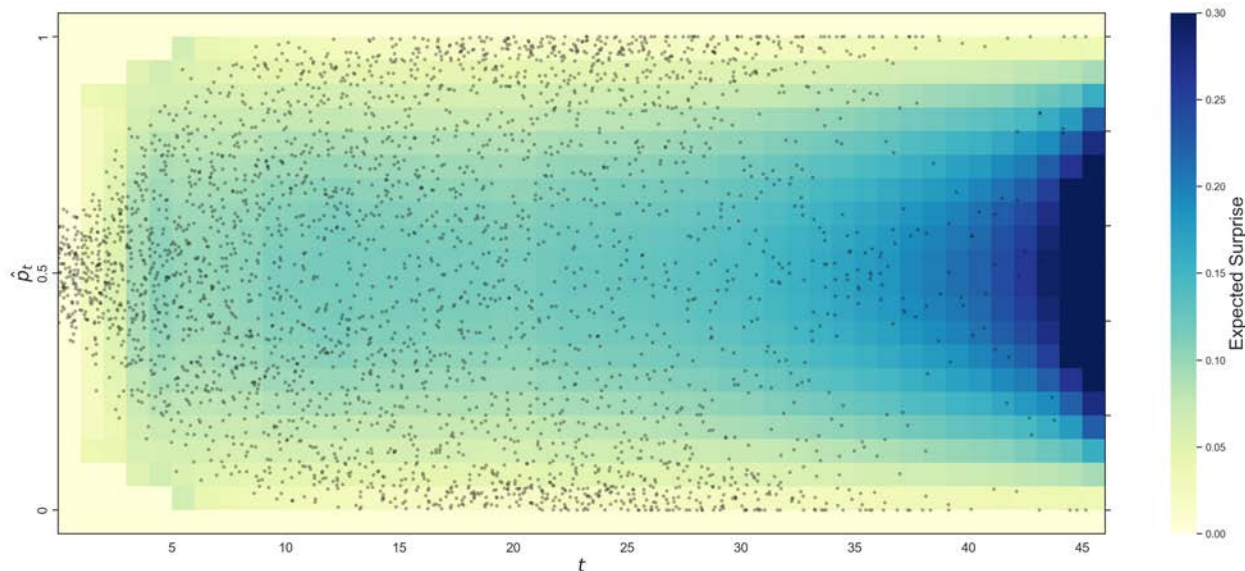
We restrict the set \mathcal{G} in two important ways. First, we require that all $G \in \mathcal{G}$ satisfy the martingale property: $E_t[p_{t+1}] = p_t$. This ensures that each game G does not violate the rules of probability and can be represented by an information process.¹⁷ Second, we constrain $G \in \mathcal{G}$ to resemble the current *League of Legends* game except in its pacing. Formally, we require that $G \in \mathcal{G}$ be a modification of G_{LoL} that is achievable by a sequence of mean-preserving spreads or contractions of the transition probabilities. These spreads and contractions are applied to all the transition probabilities in a given minute, or all the transition probabilities in a given belief bin (i.e. columns and rows of Figure 8). Third, we require that $G \in \mathcal{G}$ produces matches with the same average length as G_{LoL} to ensure that any gains in engagement are due to substantive changes in the information process rather than simply changing the length of the match. Further details of this optimization are provided in Appendix C.2.

5.1 Results

G^* shows considerable improvements over G_{LoL} . G^* improves the game experience substantially for losers but at a cost for winners: the probability of continued play changes by 1.70pp for losers, -0.95pp for winners, and overall by 0.38pp. The win-loss gap estimated in Table 2 is 0.88pp to benchmark these effects. This implies the improvement is roughly $1.93 \times$ win-loss gap for losers

¹⁷See Lemma 1 of Ely, Frankel and Kamenica (2015).

Figure 8: Expected surprise in the original game *League of Legends*, G_{LoL} , as a function of time and \hat{p}_t .



Notes: The color in each cell represents the expected surprise as computed from the transition vector. The dots represent 100 matches chosen at random and all the minutes of those matches. The dots are jittered with random displacement to visualize multiple observations at the same point.

with a disimprovement of $1.08\times$ win-loss gap for winners.¹⁸ The overall effect is an improvement equivalent to roughly $0.43\times$ the win-loss gap.

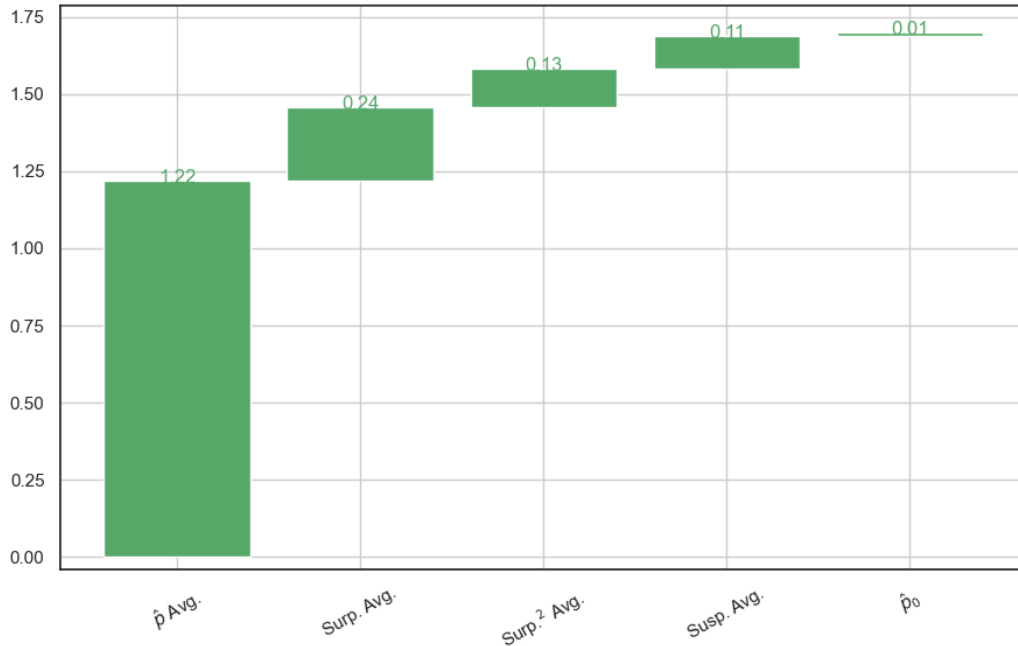
In Figure 9 we disaggregate the improvement of G^* over G_{LoL} , by feature of the belief-path. Each bar measures the amount of improvement measured in increased percentage point chance of continued play that can be attributed to that variable. Panel (a) shows the disaggregation for losers and (b) shows the disaggregation for winners. The figure makes clear that the majority of the improvement for losers can be attributed to letting them discover who they are sooner, thereby mitigating loss aversion. Losers also benefit slightly from increased surprise, more clumped surprise, and suspense. Winners, on the other hand, suffer primarily by discovering who they are sooner, thereby removing the gain of victory. The cost to them, however is vastly outweighed by the gain to losers. Winners also suffer from more clumped surprise. All of this harm is ever so slightly alleviated by benefits stemming from a more ideal ex ante challenge level.

These results are important for several reason. First, our optimized game achieved a substantial increase in predicted engagement. This implies that the information structure of the game is both important and that there may be room for improvement, even for a nearly 2 billion-dollar revenue product such as *League of Legends*.

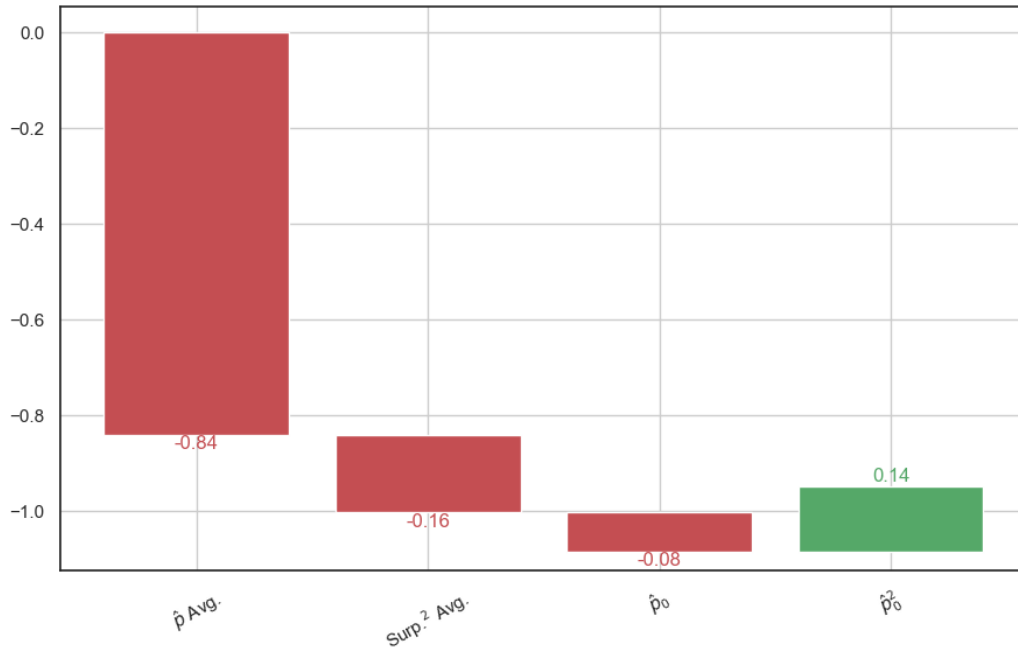
We now unpack the optimized game in Figure 10. Panel (a) shows a heatmap of the expected surprise over time and over \hat{p} . The more blue a cell, the greater the information revelation in that

¹⁸While this may appear to imply that G^* losers play more matches than winners, it does not. The win-loss gap estimated in Table 2 does not condition on the belief path; it is the unconditioned average gap in continued play between winners and losers.

Figure 9: Disaggregating the improvement from G^*



(a) Losers



(b) Winners

Notes: Improvement is measured in terms of predicted increased continued play measured in percentage points.

belief bin as measured by expected surprise. The game looks noticeably different from G_{LoL} in Figure 8. Panel (b) shows the change in expected surprise relative to G_{LoL} . Red represents regions that experienced mutations with increased information revelation, and blue represents regions that experienced mutations with decreased information revelation. It is apparent that mean-preserving spreads are used in the first 11 minutes and mean-preserving contractions are used thereon. The summary statistics for G_{LoL} and G^* are provided in Table 4. G^* exhibits greater average certainty about victory and less average surprise.

Table 4: Summary statistics for G_{LoL} and G^* conditioned on winning.

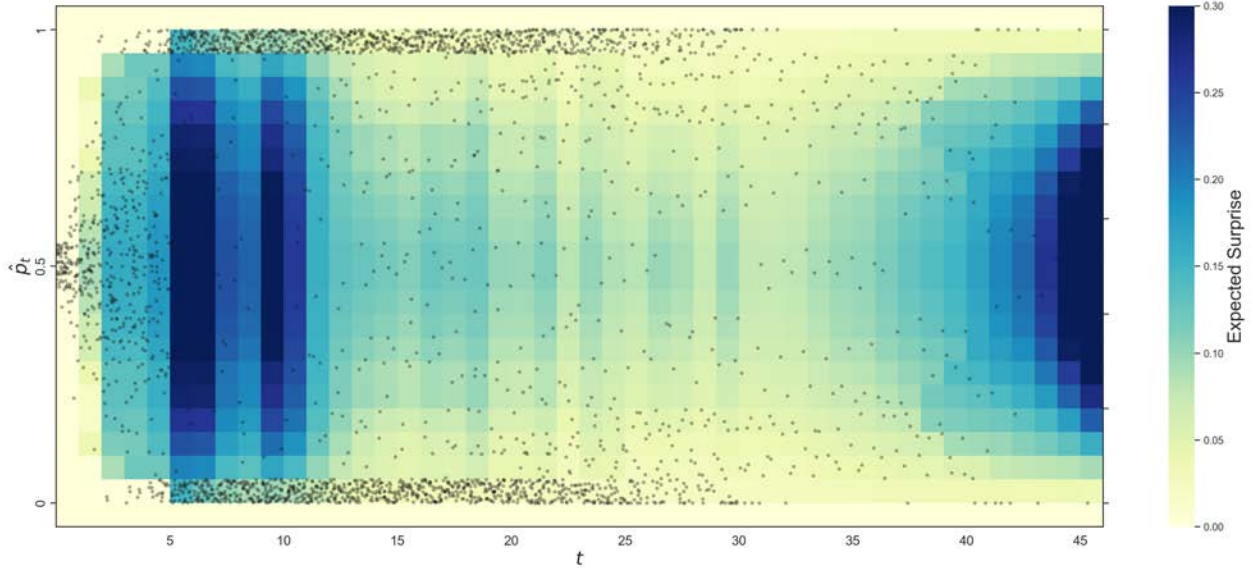
	G_{LoL}		G^*	
	Mean	SD	Mean	SD
\hat{p} Avg.	0.70	0.13	0.83	0.13
Surp. Avg.	0.06	0.02	0.04	0.02
Surp. ² Avg.	0.01	0.00	0.01	0.01
Susp. Avg.	0.08	0.02	0.08	0.02
\hat{p}_0	0.51	0.07	0.50	0.04
\hat{p}_0^2	0.26	0.07	0.25	0.04
Game Minutes	30.68	6.03	29.64	6.86

Whereas the heat-maps of expected surprise in Figures 8 and 10 show expected surprise *conditional* on reaching that cell, Figure 11 shows the unconditional average surprise over time, computed through simulating 30,000 matches. As a metaphor, the heatmaps in Figures 8 and 10 show the informational terrain should a belief path make its way there, and Figure 11 shows the average informational terrain traversed. G_{LoL} is the dotted red line, and G^* is in blue. There are several differences as previously noted, but the differences are perhaps more transparent in this view. G^* has considerably more surprise than G_{LoL} , for the first 5 minutes of the game, and then has less information revelation for the rest of the match.

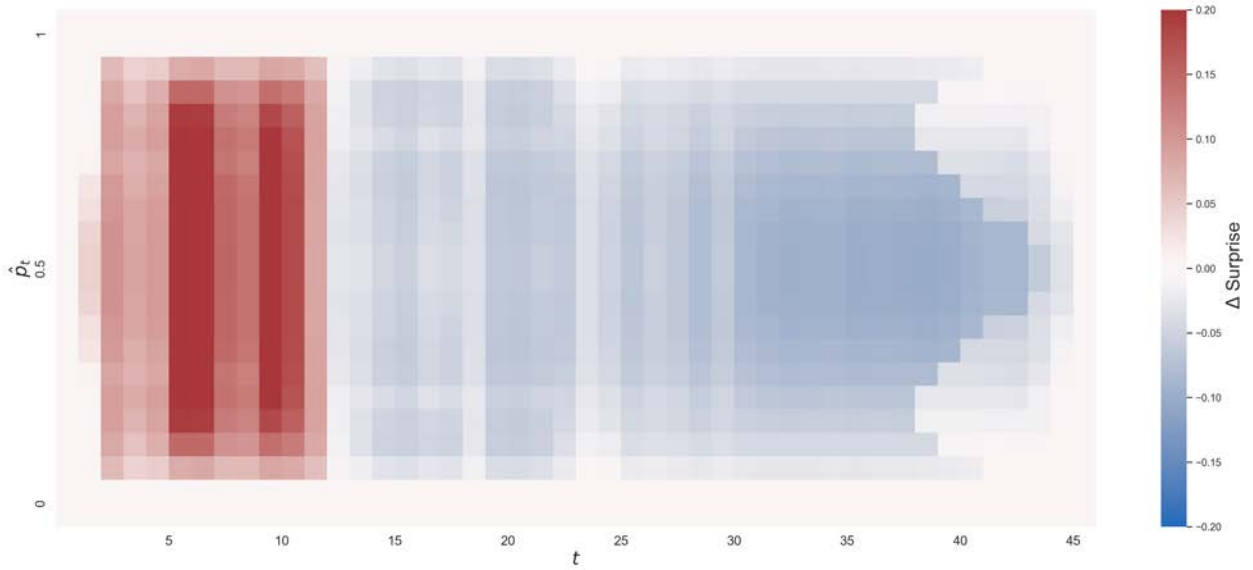
Our structural estimation implies that substantial gains in player engagement could be achieved through optimized information revelation that considers players' preferences to believe they are behind, and their preferences for surprise, suspense, and flow. But perhaps there are other features of the belief path that players also value, features considered by other theories or theories not yet discovered. Given the large number of observations, we can exploit the high-dimensionality of our data (wide and long) using high-dimensional non-parametric estimation methods and use a non-parametric objective functions instead of the OLS-estimated one. Specifically, we use gradient-boosted trees, as we did in our construction of \hat{p} and suspense, to estimate continued play as a function of features of the belief path.

The gradient-boosted optimized game delivers improvements over G_{LoL} that exceed G^* . We find that the boosted game increases continued play by 0.79pp for losers, 1.03pp for winners, and 0.91pp overall, on average. This represents improvements on the level of $0.90\times$, $1.17\times$, and $1.03\times$ the win-

Figure 10: Expected surprise as a function of time and \hat{p}_t for G^* .



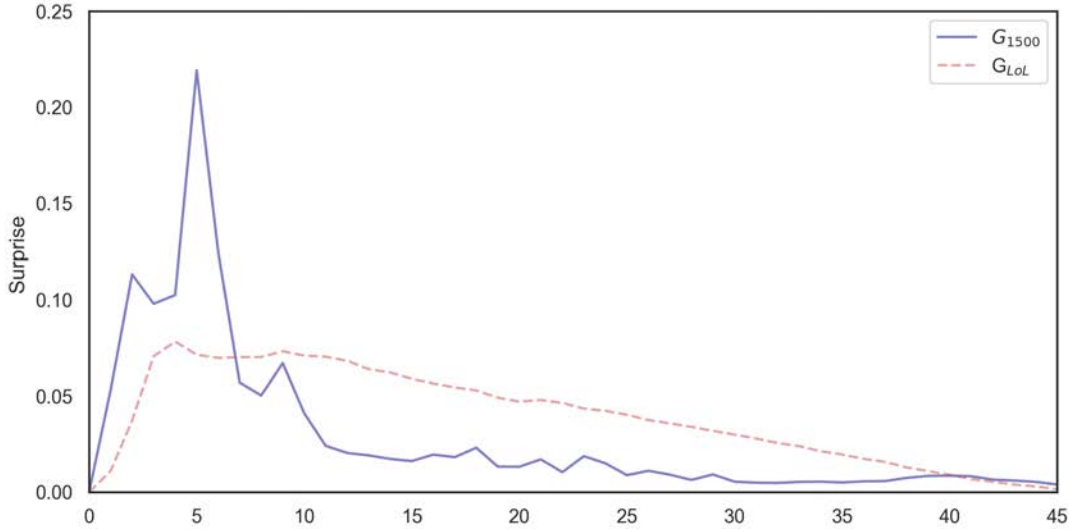
(a) Expected surprise in G^* .



(b) Δ expected surprise between G^* and G_{LoL} .

Panel (a) The color in each cell represents suspense as computed from the transition vector. The dots represent 100 matches chosen at random and all the minutes of those matches. The dots are jittered with random displacement so that multiple observations at the same point can be visualized. Panel (b) The color in each cell represents the change in suspense as computed from the transition vector.

Figure 11: Expected surprise over time for G_{LoL} and G^* .



loss gap, and amounts to about twice the improvement of G^* . These improvements accrue to both the losers and the winners rather than just the losers at expense of the winners. Qualitatively, the boosted game similarly exhibits increased surprise over G_{LoL} in the first two minutes of the match. This is followed by mostly reduced surprise that’s clumped into specific minutes with a final spike of surprise at minute 32.¹⁹

These results suggest that there may be other features of the belief path that play an important role in players’ experiences of the game. There appears to be room for novel theories to develop hypotheses regarding new features of the belief path. However, we do advise caution. We found, for example, that our results were sensitive to including additional features, and these additional features could inflate out-of-sample MSE. Without the guidance of theory or a reliable method of feature selection it is difficult to know what feature set to include.

6 Discussion

6.1 The Use of Structural Analysis in Game Design

Our evolutionary optimization of the information structure of *League of Legends* yields substantial improvements to the existing game. According to our step-wise objective function, G^* is estimated to increase continued play by 0.38pp, or 43% of the win-loss gap, with all the gains going to the losers at the expense of winners. Allowing for a more flexible atheoretical objective function based on the full belief-path achieves even greater improvement: 0.91pp or 103% of the win-loss gap with gains going to both the losers and the winners.

¹⁹We constrain the average match length to be within 1 minute of 31 minutes.

Taking a step back, one of the paper’s contributions is our structural analysis approach. Specifically, we use an evolutionary algorithm to optimize a high-dimensional information structure. As far as we know, our approach is novel in its application to information design and game design. Traditional structural approaches rely on convex optimization techniques. To do so would require greatly simplifying the player’s objective function or the information structure of the game. This would eliminate much of the real-world richness of the game or people’s preferences. Instead, we allow the information structure of the game to be complex with thousands of parameters. And we allow player’s preferences to be a function of expectation, surprise, suspense, clumpiness, and challenge level.

The main result of our structural analysis is that if *League of Legends* increased the information revelation, that is the average surprise of the early game, particularly the first 5 minutes, they could achieve meaningful increases in player engagement. Given that our data is from mid-2018, we wondered whether the designers had adjusted in this direction. We found several modifications in their patch notes that suggest they did just this.²⁰ When we presented our results to Tom Cadwell, the Chief Design Officer and Executive Vice President of R&D at Riot, he said, “Around the era this data was taken (2018), the LoL team did feel the early game had a lower level of action, and we felt this was making the early game experience less engaging than it could be. To correct this and in turn create higher action/more early fights, we used a variety of tools including changes to objectives, phase timings, and rewards.”

Though we believe that significant improvements to game design can be achieved through our structural approach, we view it as a complement to experimentation. The exact game design predictions of our evolutionary optimization can vary depending on what features of the belief-path one wants to include in the objective function. The lower the confidence that the included feature set is the ground-truth feature set, the less confidence one should have in the predictions. Lack of confidence could come from focusing on a narrow set of theories, or from including many features without having reliable feature-selection algorithms. We believe G^* may be vulnerable to the first criticism, and our gradient-boosted optimization may be vulnerable to the second criticism. However, our structural approach can be a powerful tool in game design when complemented with experimentation and a feedback cycle that improves the structural model and the game. For example, model m recommends changing x , enhancing some objective outcomes for y . Designers can then implement x and evaluate y and use this to update m . This leads to a new model m' that makes recommendation x' , and so on. To be clear, Riot has not changed the game based on this research. This is a suggestion for what firms that design information products *could* do.

²⁰The most prominent change we found on early game acceleration since 2018 is that champions spawn with a speed boost that helps them move into position. Minions also spawn 5 seconds earlier from 1:10 min in 2018, to 1:05 at the time of this writing in 2023.

6.2 Implications

Optimizing the information revelation in the game leads to a 0.38pp additional likelihood to play another match. This represents 0.62% increase in player engagement or a lower estimate of 5.9 million hours of more game-play globally per month.²¹ By way of comparison, a human lifespan of 77 years amounts to 674,520 hours. So this represents an increase of game time equal to 9 human lifespans per month. If we instead use the gradient-boosted model we conclude that the improvement amounts to a 1.53% increase in player engagement, and 14.7 million hours more played per month (about 21 lifespans). *League of Legends*' revenue model is based on in-game purchases as the game is free to play. Revenue comes from micro-payments to expand the players' options, such as more champion options, or to have skins and other cosmetic enhancements to their champions. If revenue increases linearly in player engagement,²² this increase implies a $0.64\% \times \$ 1.8 \text{ billion} = \11 million increase in annual revenue (\$28 million based on the gradient-boosted model).²³ These are benefits to the supply side of the market. If players choose to play more matches than implied by classical consumer theory, demand has shifted out and consumer surplus has increased as well. A better product for the same costs is an unambiguous increase in welfare.

These calculations ignore any other intertemporal substitution effects and there are at least two other effects that suggest that this understates the benefit. Our analysis shows that the optimized game yields more engaging matches that lead to an increase in continued play via adjacent complementarity, *ex post*. But it is also plausible that if players are receiving a more desirable distribution of belief-paths, they will have higher expected utility from playing, *ex ante*. As players discover that *League of Legends* is noticeably improved, they should reallocate planned time, as opposed to discretionary time, away from other near substitutes (other games, watching TV, and other leisurely activities) towards *League of Legends*. Second, improving the experience for losers may be more valuable than improving the experience for winners. The analysis showed that early information revelation improves the experience for losers at the expense for winners, but the net effect is positive. There may be additional benefits unaccounted for by our analysis. When do players leave a game for long stretches of time or permanently? We speculate that these long-term exits occur after a frustrating string of losses. If so, improving the experience, specifically for losers could increase player retention.

Furthermore, there are likely unaccounted social network benefits. *League of Legends* is fundamentally a team-based game with player devotion based, in-part, on playing with friends on the

²¹There are no published statistics on *League of Legends*' total monthly player-hours. The website <https://activeplayer.io/league-of-legends/> estimates as of August 4, 2023, that there are 1.331 million active hourly players. If we multiply $1.331 \text{ million} \times 24 \text{ hours} \times 30 \text{ days}$ we get a rough estimate of 950 million monthly hours of the game played per month. We multiply this time 0.62% to get 5.9 million additional hours.

²²On Riot Games' [statement of values](#), accessed on August 4, 2023, they state that they prioritize "player experience first". Our interpretation of their model is that providing gamers with good experiences leads to more revenue, with engagement as an observable proxy.

²³Approximate revenue comes from <https://mobilemarketingreads.com/league-of-legends-revenue-and-user-stats/> on April 5, 2023, which reports an estimate of \$1.8 billion in 2022. Another article <https://www.reuters.com/article/esports-lol-revenue-idUSFLM2vzDZL> accessed on August 3, 2023 reports a similar estimate of \$1.75 billion in 2020. Riot is a subsidiary of a publicly traded company and exact revenue statistics are not publicly available.

platform. This is a classic network externality and thus any marginal increase in engagement is likely to yield additional engagement through players’ social networks.

The approach presented in this paper could be easily adapted to other games. Any game that takes place over time, and has a winner and a loser, could be assessed in this way. It would be interesting to explore whether there are general preferences for beliefs across games. There may be some universals, but we think it is unlikely that preferences over all features of the belief-paths will be consistent across games. People play different games for different reasons. Some games serve as an uninvolved momentary distraction, others are strategic and deeply involved. In the same way that people seek different narrative elements in different story genres – for example, surprise endings in mysteries and well-choreographed violence in action movies – players may seek different belief-paths for different game genres. Future work can explore the commonalities and differences by genre.

6.3 Limitations

There are several limitations to our analysis. The first is with regards to the interpretation of the outcome, continued play within the next 60 minutes. From the firm’s perspective, engagement is a proxy that designers use as a performance metric, but revenue is not determined by engagement. Though it is likely that greater engagement leads to more purchases, we have no direct evidence on this point. From the player’s perspective, greater engagement may be linked with higher consumer surplus but the relationship is not certain. In favor of this relationship is the evidence that wins, which we presume are preferred over losses, result in higher engagement. Future research could improve on this point by including additional outcome variables such as purchases and consumer satisfaction surveys.

In this paper, we assume players have preferences over the experienced belief-paths, but they could possibly have preferences that are affected by the distribution of belief-paths. This is an important distinction, because our game optimization exercise would be invalid in this latter case. This case can be conceptualized in a few ways. First, preferences may have habit formation in the sense of [Becker and Murphy \(1988\)](#). The key concept in the classic paper is of temporal adjacent complementarity. One’s consumption of a good at time t could affect preferences for consumption at some later date $\tau > t$. Whereas in our current formulation utility for match m is a function of the belief-path (beliefs and suspense), $u(p_m, s_m)$, a Becker-Murphy approach would also include the history of past experienced matches, $u(p_m, s_m | h_{m-1})$, where $h_{m-1} = [(p_0, s_0), \dots, (p_{m-1}, s_{m-1})]$ is the history. In our context, an experienced belief-path of a particular kind might affect the preferences in the subsequent match. For example, a player who has a preference for variety and who just experienced a quick dash to victory might subsequently have greater demand for a drawn-out match with slow information revelation. A second conceptualization is that players preferences for an experienced match may be a function of the distribution of those matches, $u(p_m, s_m | F(p, s))$, where $F(p, s)$ is the distribution. For example, if players value paths more the rarer they are, a quick dash to victory is more highly valued under a late information-resolution game than in an

early information-resolution game.

These kinds of preferences would invalidate our game optimization exercise because we kept the objective functions over matches fixed as we changed the distribution of belief-paths. We suspect that preferences of this nature may exist but we think that they may be weak relative to the preferences over the outcome (win versus loss) and the direct preferences over the belief-path.

7 Conclusion

In this paper, we estimate the effect of features of the belief-paths in the game *League of Legends* on continued game-play. We find that both losers and winners play more if they were lagging behind, and the effect is bigger for losers. This is consistent with prospect theory and loss aversion. We find that surprise has a negative effect on continued play but suspense has a positive effect. The preference for surprise is non-linear in that losers prefer their surprise clumped and winners prefer surprise smoothed. Furthermore, we find that consistent with flow theory, winners play more when the last match is at an intermediate challenge level. This is the first paper (to our knowledge) to test for preferences over average beliefs and the clumpiness of information revelation in an entertainment product.

We then use an evolutionary algorithm to optimize an information product according to a theoretically motivated objective function. The method yields a predicted improvement to the likelihood of continued play, on the order of $0.43 \times$ win-loss gap or around 0.64% increased likelihood to play another game within 60 minutes. This corresponds to an increase in monthly engagement of 5.9 million hours. Our high-dimensional estimation suggests that even larger improvements are possible. We believe this method and similar approaches will be helpful in many future applications that involve structural analysis with high-dimensional information structures.

There are several avenues for future research. Why do people have a preference for suspense but seem to dislike surprise? Whereas suspense may allow for hopeful feelings and increased attention from both sides, surprise is good news for some, bad news for others, and may be on net negative due to loss aversion. Perhaps there are ways to rationalize the distaste in surprise from models of reference-dependent utility. The gradient-boosted optimized game achieved improvements above the game optimized by our theoretically driven objective function. This suggests that there may be features of beliefs that are highly motivating to players that remain beyond our theoretically motivated analysis. What future theories might rationalize these preferences? This project presumes that the player’s objective is victory. However, other outcomes, such as individual-level performance, may be important to players. Do players have similar preferences over beliefs for their individual-level performance, and how do they value them relative to the belief path of victory? These are some questions that remain to be explored.

References

Abuhamdeh, Sami and Mihaly Csikszentmihalyi, “The importance of challenge for the

- enjoyment of intrinsically motivated, goal-directed activities,” *Personality and Social Psychology Bulletin*, 2012, *38* (3), 317–330.
- Acland, Daniel**, “An Investigation of Flow Theory in an Online Game,” *Review of Behavioral Economics*, 2020, *7*, 317–336.
- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis**, “Recent developments in the economics of time use,” *Annual Review of Economics*, 2012, *4* (1), 373–397.
- Bartling, Björn, Leif Brandes, and Daniel Schunk**, “Expectations as reference points: Field evidence from professional soccer,” *Management Science*, 2015, *61* (11), 2646–2661.
- Becker, Gary S and Kevin M Murphy**, “A theory of rational addiction,” *Journal of Political Economy*, 1988, *96* (4), 675–700.
- Bizzozero, Paolo, Raphael Flepp, and Egon Franck**, “The importance of suspense and surprise in entertainment demand: Evidence from Wimbledon,” *Journal of Economic Behavior & Organization*, 2016, *130*, 47–63.
- Buraimo, Babatunde, David Forrest, Ian G McHale, and Juan de Dios Tena**, “Unscripted drama: Soccer audience response to suspense, surprise, and shock,” *Economic Inquiry*, 2020, *58* (2), 881–896.
- Caplin, Andrew and John Leahy**, “Psychological Expected Utility Theory and Anticipatory Feelings,” *The Quarterly Journal of Economics*, 2001, *116* (1), 55–79.
- Coates, Dennis, Brad R Humphreys, and Li Zhou**, “Reference-dependent preferences, loss aversion, and live game attendance,” *Economic Inquiry*, 2014, *52* (3), 959–973.
- Cowley, Ben, Darryl Charles, Michaela Black, and Ray Hickey**, “Toward an understanding of flow in video games,” *Computers in Entertainment (CIE)*, 2008, *6* (2), 1–27.
- Csikszentmihalyi, Mihaly and Reed Larson**, “Validity and reliability of the experience-sampling method,” in “Flow and the foundations of positive psychology,” Springer, 2014, pp. 35–54.
- Ely, Jeffrey, Alexander Frankel, and Emir Kamenica**, “Suspense and surprise,” *Journal of Political Economy*, 2015, *123* (1), 215–260.
- Friedman, Jerome H**, “Greedy function approximation: a gradient boosting machine,” *Annals of statistics*, 2001, pp. 1189–1232.
- Green, Brett and Jeffrey Zwiebel**, “The hot-hand fallacy: Cognitive mistakes or equilibrium adjustments? Evidence from major league baseball,” *Management Science*, 2018, *64* (11), 5315–5348.
- Hoffman, Bobby and Louis Nadelson**, “Motivational engagement and video gaming: A mixed methods study,” *Educational Technology Research and Development*, 2010, *58*, 245–270.
- Huang, Yan, Stefanus Jasin, and Puneet Manchanda**, ““Level up”: Leveraging skill and engagement to maximize player game-play in online video games,” *Information Systems Research*, 2019, *30* (3), 927–947.
- Kahneman, Daniel and Amos Tversky**, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 1979, *47* (2), 263–291.
- Kaplan, Scott**, “Entertainment Utility from Skill and Thrill,” 2020. Working Paper.
- Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu**, “Lightgbm: A highly efficient gradient boosting decision tree,” *Advances in neural information processing systems*, 2017, *30*.
- Kőszegi, Botond and Matthew Rabin**, “A model of reference-dependent preferences,” *The Quarterly Journal of Economics*, 2006, *121* (4), 1133–1165.
- and –, “Reference-dependent risk attitudes,” *American Economic Review*, 2007, *97* (4), 1047–

1073.

- and —, “Reference-dependent consumption plans,” *American Economic Review*, 2009, *99* (3), 909–36.
- Liaw, Richard, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica**, “Tune: A research platform for distributed model selection and training,” *arXiv preprint arXiv:1807.05118*, 2018.
- Liu, Xiao, Matthew Shum, and Kosuke Uetake**, “Attention and Inattention to Baseball Telecasts: Implications for Content (Re-)Design,” 2020. Working Paper.
- Loewenstein, George**, “Anticipation and the Valuation of Delayed Consumption,” *Economic Journal*, 1987, *97* (387), 666–684.
- Malone, Thomas W**, “Toward a theory of intrinsically motivating instruction,” *Cognitive Science*, 1981, *5* (4), 333–369.
- Nakamura, Jeanne and Mihaly Csikszentmihalyi**, “Flow Theory and Research,” *Handbook of Positive Psychology*, 2009, *195*, 206.
- Nguyen, C Thi**, *Games: Agency as Art*, Oxford University Press, USA, 2020.
- Olson, Jarrod and Daniel F Stone**, “Suspense-optimal college football play-offs,” *Journal of Sports Economics*, 2014, *15* (5), 519–540.
- 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), 129–157.
- Przybylski, Andrew K, C Scott Rigby, and Richard M Ryan**, “A motivational model of video game engagement,” *Review of General Psychology*, 2010, *14* (2), 154–166.
- Ryan, Richard M, C Scott Rigby, and Andrew Przybylski**, “The motivational pull of video games: A self-determination theory approach,” *Motivation and Emotion*, 2006, *30*, 344–360.
- Simonov, Andrey, Raluca Ursu, and Carolina Zheng**, “Do Suspense and Surprise Drive Entertainment Demand? Evidence from Twitch.tv,” 2020. Working Paper.
- Tang, Tang, Jake Kucek, and Steven Toepfer**, “Active within structures: Predictors of esports gameplay and spectatorship,” *Communication & Sport*, 2022, *10* (2), 195–215.
- Thaler, Richard H**, “Mental Accounting Matters,” *Journal of Behavioral Decision Making*, 1999, *12* (3), 183–206.
- Tversky, Amos and Daniel Kahneman**, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, 1992, *5*, 297–323.

Online Appendix

Beliefs that Entertain

A Supplementary Proofs

Proof of Proposition 3. Suppose $\gamma^t \approx \gamma^{t+1} \approx 1$, the expected gain-loss utility when $F_{\tilde{p}_{t+1}} = F_{\tilde{p}_t}$ is zero and the expected gain-loss utility when $F_{\tilde{p}_{t+2}} \neq F_{\tilde{p}_t}$ is

$$\eta(\lambda - 1)E_t\left[\int_{\tilde{p}_{t+2}}^{\infty} n(u(\tilde{p}_{t+2}) - u(r))dF_{\tilde{p}_{t+2}}(r)\right]$$

whereas the expected gain-loss utility when $F_{p_{t+1}} \neq F_{p_t}$ and $F_{\tilde{p}_{t+2}} \neq F_{\tilde{p}_t}$ is

$$\eta(\lambda - 1)E_t\left[\int_{p_{t+2}}^{\infty} n(u(p_{t+2}) - u(r))dF_{p_{t+2}}(r) + \eta(\lambda - 1)E_t\left[\int_c^{\infty} n(u(p_{t+1}) - u(r))dF_{p_{t+1}}(r)\right]\right]$$

now the (ex ante) volatility in $p_{t+1}, \tilde{p}_{t+1}, p_{t+2}, \tilde{p}_{t+2}$ is denoted by σ_{t+1} and σ_{t+2} and suppose that $\sigma_{t+1}, \sigma_{t+2} \rightarrow 0$ if $t \rightarrow 0$ then:

$$\eta(\lambda - 1)E_t\left[\int_{\tilde{p}_{t+2}}^{\infty} n(u(\tilde{p}_{t+2}) - u(r))dF_{\tilde{p}_{t+2}}(r)\right]$$

versus

$$\eta(\lambda - 1)E_t\left[\int_{p_{t+2}}^{\infty} n(u(p_{t+2}) - u(r))dF_{p_{t+2}}(r) + \eta(\lambda - 1)E_t\left[\int_c^{\infty} n(u(p_{t+1}) - u(r))dF_{p_{t+1}}(r)\right]\right]$$

now, each of these components are negative and proportional to σ_t

$$\int \int (u(\sigma y) - u(\sigma y'))dF(y')dF(y) = -\frac{1}{2}\eta(\lambda - 1) \int \int (u(\sigma \max(y, y')) - u(\sigma \min(y, y')))dF(y')dF(y)$$

because $\sigma_{t+2,t} = \sqrt{\sigma_{t+1}^2 + \sigma_{t+2}^2} < \sigma_{t+1} + \sigma_{t+2}$, the agent prefers the uncertainty to resolve clumped not piecewise. \square

Proof of Proposition 4. For ease of notation, g^t map $p = (p_1, p_2, \dots, p_{t-1}, p_t, p_{t+1}, \dots, p_{T-1}, p_T)$ to $p' = (p_1, p_2, \dots, p_{t-1}, p_{t+1}, p_{t+1}, \dots, p_{T-1}, p_T)$. Consider the difference in players' expected utilities between games G and G^t :

$$\begin{aligned} \tilde{U}^{PT}(G^t) - \tilde{U}^{PT}(G) &= \int U^{PT}(g^t(p)) - U^{PT}(p)dG(p) \\ &= \int v\left(wz - z\frac{1}{T}\sum_{\tau} g^t(p_{\tau})\right) - v\left(wz - z\frac{1}{T}\sum_{\tau} p_{\tau}\right) dG(p) \\ &= \int v\left(wz - z\frac{\sum_{\tau \neq t} p_{\tau} + p_{t+1}}{T}\right) - v\left(wz - z\frac{1}{T}\sum_{\tau} p_{\tau}\right) dG(p). \end{aligned}$$

Now consider the relative utilities of these games conditioning on being a winner:

$$\begin{aligned}
\tilde{U}^{PT}(G^t|w=1) - \tilde{U}^{PT}(G|w=1) &= \int U^{PT}(g^t(p)) - U^{PT}(p) dG(p|p_T=1) \\
&= \int \eta \left(z - z \frac{\sum_{\tau \neq t} p_\tau + p_{t+1}}{T} \right) - \eta \left(z - z \frac{1}{T} \sum_{\tau} p_\tau \right) dG(p|w=1) \\
&= \frac{\eta z}{T} \int p_t - p_{t+1} dG(p|w=1) \leq 0.
\end{aligned}$$

Thus, winners prefer G to G^t . The inequality above follows from the fact that $E[p_{t+1} - p_t | w = 1] \geq 0$ or else winners are expected to mis-learn on average at minute t . (For proof, see Proposition 5 and Corollary 0.1.) The argument that losers prefer G^t to G is similar. \square

Proposition 5 (At any point in a match, beliefs are expected to move towards the realized outcome.). *Denote the match history at time t by h_t and the probability of winning the match given history h_t by $p_t := E[w | h_t]$. A slightly stronger claim than necessary for Proposition 4 is that beliefs are always expected to update in the correct direction between t and $t+1$. If the eventual outcome will be a win, then this translates to:*

$$E[p_{t+1} | h_t, w = 1] \geq E[p_t | h_t, w = 1]. \quad (9)$$

Proof. For ease of notation, define W to be state of winning the match (i.e. $w = 1$), and let everything that follows be conditioned on h_t so that we are only interested in how the event e happening at time t is expected to change beliefs.²⁴ Using this notation, the proposition is:²⁵

$$E[P(W | e) | W] \geq P(W), \quad (11)$$

Observe that:

$$E[P(W | e) | W] := \int P(W | e) dP(e | W); \quad P(W) := \int P(W | e) dP(e) \quad (12)$$

Then, the result follows from Bayes rule:

$$P(e | W) := \frac{P(W | e) P(e)}{P(W)} \geq P(e) \iff P(W | e) \geq P(W). \quad (13)$$

Relative to integrating with respect to measure $P(e)$, integrating with respect to $P(e | W)$ increases weight on events that increase the posterior probability of winning and decreases weight on events

²⁴In other words: $h_{t+1} := (h_t, e)$. This could equivalently be written as:

$$E[P(W | h_{t+1}) | W, h_t] := \int P(W | h_{t+1}) dP(h_{t+1} | W, h_t) \geq P(W | h_t) := \int P(W | h_{t+1}) dP(h_{t+1} | h_t). \quad (10)$$

²⁵Note that the right hand side of the inequality follows from the fact that $E[P(W)|W] = P(W)$. This is because the expectation integrates with respect to e , which does not affect $P(W)$. (Recall that $P(W)$ is being used as shorthand for $P(W|h_t)$, so more formally, $P(W|h_t)$ is not a function of e .)

that decrease the probability of winning. \square

Corollary 0.1. *For any game G and minute t , beliefs are expected to move towards the realized outcome.*

Proof. Proposition 5 shows that the desired inequality holds given any history h_t at time t . The result follows from integrating over the distribution of histories at t generated by game G . \square

B Estimation of the Belief-Paths

B.1 Estimation of \hat{p}

We approximate p_{mjt} — the probability of team j ²⁶ winning match m at minute t — using gradient boosted decision trees (GBDT) via LightGBM (Ke et al., 2017). GBDT (Friedman, 2001) builds on standard boosting approaches, where an ensemble of “weak” learners—in this case, decision trees—are aggregated into a single “strong” model. In the GBDT algorithm, weak learners are iteratively trained to predict the residual errors of the strong model, then incorporated into the strong model. Formally, at iteration i , the weak learner f_i predicts the residuals of strong model F_i , and the strong model is updated as $F_{i+1} = F_i - f_i$. At the end of training, we have a resulting strong model F^* . LightGBM modifies the standard GBDT algorithm with efficiency and scalability improvements, while maintaining high performance relative to the state-of-the-art. Full details are described in detail by Ke et al. (2017).

The predictive model learns the relationship between the state s of match m at minute t and the probability of a team j winning that match. The objective function used in training is to minimize the mean squared error, given by:

$$MSE := N^{-1} \sum_m \sum_{t=0} (w_{mj} - \hat{p}_{mjt})^2. \quad (14)$$

where $\hat{p}_{mjt} = F_i(s_{mjt})$ and $w_{mj} \in \{0, 1\}$ denotes whether or not team j won match m and N is the number of observations in the training set.

The input features to the predictive model are described, at a high level, by Table B1. In total, there are 146 independent variables provided as input to the model. We use Ray Tune (Liaw et al., 2018) to select the hyperparameters of the LightGBM model using k -fold cross-validation.

Figure A1 shows the distribution of \hat{p} at 5 minute intervals. The mass of the distribution is centered around 0.5 for both losers and winners at the beginning of the game, and progressively shifts to the extremes of 0 and 1, for losers and winners respectively. There are a few observations worth noting. First, the shifting of the probability mass visualizes the steady resolution of uncertainty over time. Second, though the mass is centered at 0.5, there is still heterogeneity at minute

²⁶Note that, in practice, we approximate \hat{p}_{mjt} for one of two teams, using $1 - \hat{p}_{mjt}$ as the estimate for the second team.

Table B1: Variables Used to Predict the Probability of Winning

Variables	Level of Description
Time	Game level
Queue	Game level
Dragon Kills	Team level
Tower Kills	Team level
Inhibitors	Team level
Base Position	Team level
Lane Position	Team level
Champions	Player level
Kills	Player level
Deaths	Player level
Gold	Team and Player level
Experience Points	Team and Player level
Damage to Champions	Team and Player level
Health	Team and Player level
Power	Team and Player level

0, suggesting that after match assignment and the champion selection process, matches can start with the scales tilting in a team’s favor.

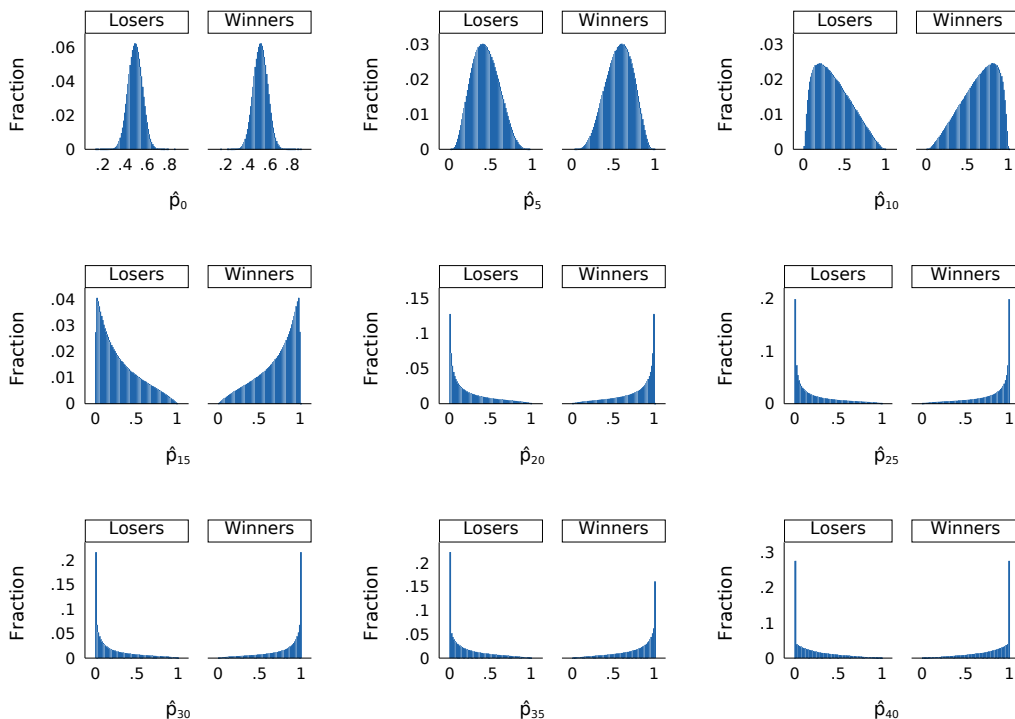
B.2 Estimation of Suspense

Our estimation of suspense follows a nearly identical procedure to that of \hat{p} . We use LightGBM with the same set of independent variables described in Table B1, with the addition of \hat{p} . Again, we use k -fold cross validation to select the hyperparameters of the model. The primary difference is in the objective function, where the model learns to predict the suspense by minimizing the MSE objective of

$$MSE := N^{-1} \sum_m \sum_t^{T-1} (\text{suspense}_t - F_i(s'_{mjt}))^2 \quad (15)$$

where s'_{mjt} denotes the state of the match with the addition of \hat{p} .

Figure A1: Distribution of \hat{p} at 5 minute intervals.



C Supplemental Information on the Game Optimization

C.1 Operationalizing the Game

We begin by defining the game. We abstract away all aspects of the game aside from belief-path generation. To characterize the data-generating process, we need to express the probability that p_t transitions to a belief p_{t+1} . We discretize the probability space into 20 interior bins (i.e., $(0, 0.05]$, $(0.05, 0.1]$, \dots , $(.95, 1)$) and 2 “absorbing state” bins (i.e. $\{0\}$ and $\{1\}$). This allows us to have a finite number of transition probabilities. For a given belief bin $b_{n,t}$ where $n \in \{0, \dots, 21\}$ at time t , there are 22 transition probabilities that represent the probability that $p_{t+1} \in b_{n,t+1}$. Each bin at time t has a 22-element vector of these transition probabilities. Let M_t be a transition matrix composed of each of the 20 transition vectors (one for each interior bin) for time t . The game is thus defined as

$$G := \langle F(p_0), \{M_t\}_{t=0}^{t=45} \rangle$$

where $F(p_0)$ represents the CDF of initial beliefs before a match begins.²⁷ To reduce the computational burden, we assume all matches end within 45 minutes (this represents 97% of the sample). G characterizes the distribution of belief-paths. It is worth noting how the laws of probability constrain G . First, $F(p_0)$ must be symmetrically distributed around 0.5. Second, for each transition vector for bin n at time t there must be an inverted transition vector for bin $21 - n$. Third, transition probabilities must be such that $E_t[p_{t+1}] = p_t$.

Our game optimization procedure relies on several primary components: mean-preserving spreads and mean-preserving contractions, and an evolutionary algorithm. In this section, we provide full details on both.

C.2 Evolutionary Algorithm

Now that we have estimated an objective function and have operationalized the game G , the next step is to choose a G from a subset of possible games to maximize the objective function. Because G is high-dimensional with thousands of parameters, we take an evolutionary approach, searching for local optima in the vicinity of G_0 . Because game length directly decreases the probability of playing another match through mechanical means, we constrain our search to local optima that are approximately of the same average game length. We use the same algorithm to optimize with the OLS objective function.

The evolutionary algorithm is detailed in Algorithm 2. At a high level, at each iteration i , we apply bin- or minute-level spreads on the transition matrix G_i , the procedure of which is described in Algorithm 1. If the expected match length is greater than our fixed threshold, we only consider spread perturbations that decrease match length. If below the threshold, we only consider contractions to increase length. After creating mutants of G_i at each minute and bin, we score them using our scoring function—e.g., a LightGBM model trained to predict continued play from the belief path or our OLS regression coefficients—then select the mutant with the maximum score. If no mutant yields an improvement over the current G_i , we increase the number of perturbations applied at each bin and minute and repeat the process, either until we’ve reached a specified number of perturbation iterations or an improvement is reached over G_i .

C.2.1 Mutations

A transition vector is a probability distribution, and creating variation in those vectors is required for the evolutionary process. We create variation in these transition vectors by applying mean-preserving contractions or mean-preserving spreads. Every cell in Figure 8 is represented by a transition vector, and can be altered in this way. This yields a total of 45 minutes \times 20 bins

²⁷From the econometrician or designer’s perspective there are two reasons that p_0 need not always be 0.5. First, ratings are visible to the designer (and to us), and one can use this to see whether one team has an advantage over the other. Moreover, the variance of this advantage is a design choice. Second, there are pre-match decisions that affect the probability of winning, most notably players choose their champions. Though the designers aim for game balance, some champions may be slightly stronger than others, and certain team compositions may be particularly effective at countering other team compositions.

Algorithm 1 Game Mutations

```
procedure MUTATIONBYTIME( $G, t, \text{method}, n = 1$ )  
  for  $i \in \{1, \dots, n\}$  do  
    for  $b \in \{1, \dots, 20\}$  do  
      if  $\text{method} == \text{contraction}$  then  
         $G \leftarrow \text{MEANPRESERVINGCONTRACTION}(G, t, b, p = 0.05)$   
      else if  $\text{method} == \text{spread}$  then  
         $G \leftarrow \text{MEANPRESERVINGSPREAD}(G, t, b, p = 0.1)$   
      end if  
    end for  
  end for  
  return  $G$   
end procedure
```

```
procedure MUTATIONBYBIN( $G, b, \text{method}$ )  
  for  $i \in \{1, \dots, n\}$  do  
    for  $t \in \{0, \dots, 44\}$  do  
      if  $\text{method} == \text{contraction}$  then  
         $G \leftarrow \text{MEANPRESERVINGCONTRACTION}(G, t, b, p = 0.05)$   
      else if  $\text{method} == \text{spread}$  then  
         $G \leftarrow \text{MEANPRESERVINGSPREAD}(G, t, b, p = 0.1)$   
      end if  
    end for  
  end for  
  return  $G$   
end procedure
```

Algorithm 2 Evolutionary Algorithm

Require: G **Require:** $f := \text{PATHEVALUATIONFN}$ **Require:** $\text{maxPerturbationIterations} \leftarrow 50$ **Require:** $\text{iterationIncrement} \leftarrow 3$ **Require:** $N \leftarrow 30,000$ **Require:** $\text{lengthThreshold} \leftarrow 1$ $n \leftarrow 1$ $i \leftarrow 0$ $G_i \leftarrow G$ $L \leftarrow \text{EXPECTEDMATCHLENGTH}(G)$ $\text{curScore} \leftarrow N^{-1} \sum^N f(g \sim G)$ $\text{improving} \leftarrow \text{TRUE}$ **while** improving **do** **if** $\text{EXPECTEDMATCHLENGTH}(G_i) \geq L + \text{lengthThreshold}$ **then** $\text{methods} \leftarrow \{\text{spread}\}$ **else if** $\text{EXPECTEDMATCHLENGTH}(G_i) \leq L - \text{lengthThreshold}$ **then** $\text{methods} \leftarrow \{\text{contraction}\}$ **else** $\text{methods} \leftarrow \{\text{spread}, \text{contraction}\}$ **end if** $\text{mutants} \leftarrow \{\}$ **for** $\text{method} \in \text{methods}$ **do** **for** $t \in \{0, \dots, 44\}$ **do** $\text{mutants} \leftarrow \text{mutants} \cup \text{MUTATIONBYTIME}(G_i, t, \text{method}, n)$ **end for** **for** $b \in \{1, \dots, 20\}$ **do** $\text{mutants} \leftarrow \text{mutants} \cup \text{MUTATIONBYBIN}(G_i, b, \text{method}, n)$ **end for** **end for** $\text{scores} \leftarrow \left\{ N^{-1} \sum^N f(g \sim \text{mutants}_j) \right\}_{j=1}^{|\text{mutants}|}$ $j^* \leftarrow \arg \max_j \text{scores}_j$ $\text{bestScore} \leftarrow \text{scores}_{j^*}$ **if** $\text{bestScore} > \text{curScore}$ **then** $i \leftarrow i + 1$ $\text{curScore} \leftarrow \text{bestScore}$ $G_i \leftarrow \text{mutants}_{j^*}$ $n \leftarrow 1$ **else if** $n \leq \text{maxPerturbationIterations}$ **then** $n \leftarrow n + \text{iterationIncrement}$ **else** $\text{improving} \leftarrow \text{FALSE}$ **end if****end while**

$\times 2$ (contract or spread) = 1,800 potential mutations. To keep the number of potential mutants manageable, we mutate not by minute-bin, the cells of Figure 8, but by the entire minute or the entire bin, the columns and rows of Figure 8 respectively. We create minute-mutations by selecting a single minute t and apply a mean-preserving spread to all transition vectors for that minute; or in other words, all the cells in that column. We create bin-mutations by selecting a bin n and its inverse bin $21 - n$ and apply a mean-preserving spread to all transition vectors for those two bins. Likewise, we do the same for mean-preserving contractions. This produces 90 minute-mutations and 18 bin-mutations.²⁸ Each mutation is a slightly altered version of G_i .

We define mean-preserving spreads and contractions in Algorithm 3. Spreads effectively take mass from the center of the distribution and iteratively spread mass out to the endpoints. Beginning at the central bins 10 and 11, we take $d\%$ of each bin, giving percentage $\frac{d}{2}\%$ to each of the neighboring bins. From the center two bins, this process is repeated until we reach bins 2 and 19, this means the mass of all bins will change except for terminal bins 0 and 21.

Contractions are defined in an alternative way to spreads. Rather than the same iterative procedure,²⁹ we take the approach of decreasing the mass in all non-terminal bins other than the mean bin of the distribution. Formally, we want to decrease the total mass of these non-mean, non-terminal bins by proportion d while ensuring that we maintain the mean. To do so, we want to take proportion x_1 of the mass from bins below the mean bin and proportion x_2 of the mass above the mean bin and transfer that mass to the mean bin. We calculate x_1 and x_2 by solving a system of two linear equations. Using the notation in Algorithm 3, the first equation is given by

$$(1 - x_1)m_l + (1 - x_2)m_u = (1 - d)(m_l + m_u) \quad (16)$$

and the second, mean-preserving equation is

$$\begin{aligned} \mathbb{E}[p|G_{t,b,0} = 1]G_{t,b,0} + (1 - x_1)w_l + (G_{t,b,\bar{b}} + x_1m_l + a_2m_u) \mathbb{E}[p|G_{t,b,\bar{b}=1}] \\ + (1 - a_2)w_u + \mathbb{E}[d|G_{t,b,21} = 1]G_{t,b,21} = \bar{G}_{t,b} \end{aligned} \quad (17)$$

C.2.2 Selection Dynamics

To evaluate each mutant, we simulate 30,000 belief-paths and score them based on our regression model for continued play separately for winners and losers. Specifically, we use our stepwise regression coefficients from Columns (3) and (6) of Table 3. The mutant that scores the highest becomes the parent for the next generation. This best-performing mutant becomes G_{i+1} . We then repeat the process: G_{i+1} produces 108 mutant offspring, the best-performer is selected to be the parent of the next generation, and so on. The evolution continues until there is no improvement across a

²⁸We do not include mutants that mutate the bins 1 and 20 (the two non-absorbing state corner bins: (0,0.05] and (0.95,1)). Our mean-preserving spread and contraction algorithms shift mass symmetrically to neighboring bins, and as these are on the corners they do not have sufficient neighbors for the algorithm.

²⁹Preliminary optimizations showed that using a similar iterative procedure as in the spreads yielded excessively “spiky” distributions.

Algorithm 3 Mean Preserving Perturbations

procedure MEANPRESERVINGSPREAD($G, t, b, p = 0.1$)

$l \leftarrow 10$

$h \leftarrow 11$

while $l > 1$ and $h < 20$ **do**

$m_l \leftarrow p \cdot G_{t,b,l}$

$m_h \leftarrow p \cdot G_{t,b,h}$

$G_{t,b,l} \leftarrow G_{t,b,l} - m_l$

$G_{t,b,l-1} \leftarrow G_{t,b,l-1} - \frac{m_l}{2}$

$G_{t,b,l+1} \leftarrow G_{t,b,l+1} - \frac{m_l}{2}$

$G_{t,b,h} \leftarrow G_{t,b,h} - m_h$

$G_{t,b,h-1} \leftarrow G_{t,b,h-1} - \frac{m_h}{2}$

$G_{t,b,h+1} \leftarrow G_{t,b,h+1} - \frac{m_h}{2}$

$l \leftarrow l - 1$

$h \leftarrow h + 1$

end while

return G

end procedure

procedure MEANPRESERVINGCONTRACTION($G, t, b, p = 0.05$)

$\bar{b} \leftarrow \lceil \frac{\sum_{b'} G_{t,b,b'} E[p|b']}{0.05} \rceil$

▷ Determine expected bin given the distribution $G_{t,b}$

$m_l \leftarrow \sum_{b'=1}^{\bar{b}-1} G_{t,b,b'}$

▷ Mass of lower bins

$m_u \leftarrow \sum_{b'=\bar{b}+1}^{20} G_{t,b,b'}$

▷ Mass of upper bins

$w_l \leftarrow \sum_{b'=1}^{\bar{b}-1} G_{t,b,b'} \mathbb{E}[p|b']$

▷ Weighted sum of lower bins

$w_u \leftarrow \sum_{b'=\bar{b}+1}^{20} G_{t,b,b'} \mathbb{E}[p|b']$

▷ Weighted sum of upper bins

$\bar{G}_{t,b} \leftarrow \sum_{b'} G_{t,b,b'} E[p|b']$

$\mathbf{A} \leftarrow \begin{bmatrix} m_l & m_u \\ \mathbb{E}[p|\bar{b}]m_l - w_l & \mathbb{E}[p|\bar{b}]m_u - w_u \end{bmatrix}$

$\vec{b} \leftarrow \begin{bmatrix} p(m_l + m_u) \\ \bar{G}_{t,b} - \mathbb{E}[p|bin=0]G_{t,b,0} - w_l - w_u - \mathbb{E}[p|bin=21]G_{t,b,21} - G_{t,b,\bar{b}}E[p|bin=\bar{b}] \end{bmatrix}$

$\vec{x} \leftarrow$ Solve system of linear equations $\mathbf{A}\vec{x} = \vec{b}$

if $|x_1| > 1$ **then**

$x_1 \leftarrow 1.0$

$x_2 \leftarrow \frac{\sum_{b'} G_{t,b,b'} E[p|b'] - \mathbb{E}[p|bin=0]G_{t,b,0} - \mathbb{E}[p|bin=21]G_{t,b,21} - w_u - G_{t,b,\bar{b}}E[p|bin=\bar{b}] - m_l E[p|bin=\bar{b}]}{m_u E[p|bin=\bar{b}] - w_u}$

else if $|x_2| > 1$ **then**

$x_2 \leftarrow 1.0$

$x_1 \leftarrow \frac{\sum_{b'} G_{t,b,b'} E[p|b'] - \mathbb{E}[p|bin=0]G_{t,b,0} - \mathbb{E}[p|bin=21]G_{t,b,21} - w_l - G_{t,b,\bar{b}}E[p|bin=\bar{b}] - m_u E[p|bin=\bar{b}]}{m_l E[p|bin=\bar{b}] - w_l}$

end if

▷ Update $G_{t,b}$

$G_{t,b,\bar{b}} \leftarrow G_{t,b,\bar{b}} + x_1 m_l + x_2 m_u$

$G_{t,b,1:\bar{b}-1} \leftarrow (1 - x_1)G_{t,b,1:\bar{b}-1}$

$G_{t,b,\bar{b}+1:20} \leftarrow (1 - x_2)G_{t,b,1:\bar{b}+1:20}$

return G

end procedure

generation.

We wish to keep our analysis in the vicinity of G_{LoL} . Exploration of the game spaces far from *League of Legends* are more likely to produce unrealistic solutions, as there are many aspects of game design that are not included in our model. An unconstrained optimization could lead to average match lengths that are extremely short or extremely long, fundamentally changing the character of the game. To limit changes to average match length we constrain which mutants can be selected. We allow the evolved game G_i to produce an average match length within 1 minute of the average produced by G_{LoL} , 30.65 minutes. Once G_i has produced an average match length shorter than 29.65 minutes, only mean-preserving contraction mutants are considered for the next generation.³⁰ Likewise, if average match length exceeds 31.65 minutes, only mean-preserving spread mutants are considered for the next generation.

C.3 Evaluation and Starting Points

We need a way to evaluate whether these evolved games improve upon G_{LoL} . We simulate 30,000 matches and take the average predicted continued-play, using our stepwise regression coefficients from Columns (3) and (6) of Table 3, just as we did for evaluating the mutants at each evolutionary iteration.

The evolutionary search is myopic and local in nature, and therefore it is quite possible that the output is a local optimum. Different starting points in the vicinity of G_{LoL} may lead to different outcomes. We therefore choose 27 starting points as G_0 , and run the optimization algorithm from each point. The highest predicted continued-play game is our optimized game G^* . To create the 27 starting points we break up time into blocks of $t \in [1, 10]$, $t \in (10, 20]$, and $t > 20$. For each block we either apply 25 spreads to every minute, ten contractions to every minute, or do nothing. The 3^3 factorial design give us the 27 starting points. Our G^* started from the G^0 that was spread-spread-nothing.

We found that we can reach qualitatively similar results more efficiently by increasing the number of spreads applied to mutants. In our 27 starting points, we use 7 as the starting number of spreads: $n = 7$ in the notation of Algorithm 2.

C.4 Game Optimization with the Objective Function Estimated by Gradient-Booted Trees

We use gradient-boostered regression trees with MSE-minimization to estimate an objective function for our evolutionary optimization. However, the process is not as straightforward as with the OLS counterpart. Linear regression has the property that the effects of the variables are additively separable, but this is not generally the case with regression trees. In order to isolate the effects of the belief-path from the effects of the control variables, we first run LightGBM using only the controls. These controls include rating, matches played in the previous 2 hours, matches played in

³⁰Mean-preserving contraction mutations increase average match length and mean-preserving spread mutations decrease average match length.

the previous 2 hours with the same champion, log minutes from the previous match, log minutes from the previous match with the same champion, and fixed effects for day, hour, and date. We then take the residuals from this regression and then run a second LightGBM regression on the belief-path. We used the same features as in our step-wise OLS specification.

One of the most important determinants of continued play is match length. Players are much more likely to play another match if the last match was short. This says little about player engagement or enjoyment of the game as the relationship could be mechanical. A player who has a fixed amount of play-time can fit in more matches if those matches are shorter. Thus, unconstrained optimization for maximizing continued play will drive the game to be very short. We want to control for the match-length effect. We have two approaches for solving this problem. The first is the buffer method that we described above.

The second approach ensures limited changes to average match length by modifying the objective function. We modify our objective function to have a marginal penalty for shorter matches and the same marginal bonus for longer matches. This, in principle, will keep the evolutionary process on track, selecting mutations for their increased engagement without altering average match length. We find the marginal incentive by treating the penalty as a parameter and performing a grid search over optimizations that result in the lowest change in game length over the first 100 iterations. We also layer in the buffer approach as before as an additional control to ensure the average match length stays within the 1-minute bound.³¹ We then evaluate the optimized game based on our estimated objective function.

D Additional Results

Tables D2 and D3 show alternative specifications of OLS regressions of continued play on features of the belief path. Interestingly, suspense average has a negative coefficient for losers and surprise average has a positive coefficient for winners contrary to our main results in Table 3. This is because surprise average and suspense average are correlated. This shows the importance of including both terms simultaneously.

³¹When the length-incentive is included in the objective function, match length never reaches the bounds of the buffer, as one should expect. The buffer is not entirely redundant, however, as it plays an important role when we conduct exercises in which the initial match length begins outside of the buffer.

Table D2: Effect of the Belief-Path on Continued Play (Losers)

	(1)	(2)	(3)	(4)	(5)
\hat{p} Average [†]	-7.50 (0.13)				
Surprise Average [†]		-31.53 (0.77)			
Surprise ² Average			-97.46 (2.63)		
Suspense Average [†]				-30.07 (0.82)	
\hat{p}_0					9.28 (3.00)
\hat{p}_0^2					-8.84 (3.05)
Peak \hat{p}_0					0.53
Mean of DV	58.76	58.76	58.76	58.76	58.76
SD of DV	49.23	49.23	49.23	49.23	49.23
R2	.013	.013	.013	.013	.013
Clusters	2,805,461	2,805,461	2,805,461	2,805,461	2,805,461
N	13,949,001	13,949,001	13,949,001	13,949,001	13,949,001

Notes: All specifications are OLS with robust standard errors clustered at the match-team level. Continued play is multiplied by 100 so that coefficients can be interpreted in percentage points. All specifications include linear controls for games played in the previous 2 hours, games played in the previous 2 hours with the same champion, minutes from the previous game, minutes from the previous game using the same champion, quadratic controls for player skill rating, and fixed effects for date, day of week, hour of day, team, role, champion, and game length.

[†] denotes that the variable definition varies by column: in (1), (2), (4) and (5) explanatory variables are averaged over all minutes, and in (3) and (6) explanatory variables are averaged over all minutes excluding minute zero.

Table D3: Effect of the Belief-Path on Continued Play (Winners)

	(1)	(2)	(3)	(4)	(5)
\hat{p} Average [†]	-5.03 (0.13)				
Surprise Average [†]		18.32 (0.78)			
Surprise ² Average			52.74 (2.65)		
Suspense Average [†]				23.14 (0.83)	
\hat{p}_0					25.72 (3.17)
\hat{p}_0^2					-24.68 (3.09)
Peak \hat{p}_0					0.52
Mean of DV	59.64	59.64	59.64	59.64	59.64
SD of DV	49	49	49	49	49
R2	.012	.012	.012	.012	.012
Clusters	2,805,336	2,805,336	2,805,336	2,805,336	2,805,336
N	13,949,906	13,949,906	13,949,906	13,949,906	13,949,906

Notes: All specifications are OLS with robust standard errors clustered at the match-team level. Continued play is multiplied by 100 so that coefficients can be interpreted in percentage points. All specifications include linear controls for games played in the previous 2 hours, games played in the previous 2 hours with the same champion, minutes from the previous game, minutes from the previous game using the same champion, quadratic controls for player skill rating, and fixed effects for date, day of week, hour of day, team, role, champion, and game length.

[†] denotes that the variable definition varies by column: in (1), (2), (4) and (5) explanatory variables are averaged over all minutes, and in (3) and (6) explanatory variables are averaged over all minutes excluding minute zero.

* $p < 1e - 2$, ** $p < 1e - 4$, *** $p < 1e - 6$