Experience-Driven Peer Effects: Evidence from a Large Natural Experiment

Anonymous Authors

Abstract

Social interactions between people are a central mechanism by which behavior spreads. Several field studies have shown how observing peer behavior affects ones own behavior in a wide range of domains including health, information diffusion, advertisement, and education. However, the role of observing peer experience or outcome—the reward for the observed peer behavior—has largely gone uninvestigated. Here we examine evidence from a large-scale online setting, game play records from League of Legends match-ups, where we are able to disentangle the effects of observing peer behavior from observing peer experience. We find that in addition to positive peer effects from observing behavior, the effect is accentuated by observed experience, with a large positive effect when observing a good outcome and a small (but still positive) effect when observing a bad outcome. We further find that this experience driven effect is moderated by skill, becoming more pronounced among highly skilled subjects. Our findings demonstrate the importance of the role of observed peer experience beyond peer behavior. We anticipate this result to be of use to both practitioners and theoreticians in the space of social influence. For example, online platforms may wish to broadcast the positive outcomes of peers, more than mere behaviors, when a user performs a behavior desirable to the platform. Furthermore, efforts aimed at maximizing the social spread of a product may benefit from modeling experience-driven peer effects as part of the spreading process.

Introduction

A wide variety of human behaviors spread via peer to peer social interactions. These peer effects have been shown to cause people to exercise (Aral and Nicolaides 2017; Zhang et al. 2015), interact with advertisements (Bakshy et al. 2012a), learn (Hoxby 2000), and share information (Bakshy et al. 2012b). In addition, a separate vein of both empirical and theoretical work explores how network structure affects the manner by which behaviors spread, including ideas such as complex contagion where multiple peers must adopt for a behavior to spread (Centola and Macy 2007). Furthermore, extensive investigation has gone into how to seed intervegntions as to maximize the spread of a behavior (Centola 2010; Kempe, Kleinberg, and Tardos 2003).

An understanding of peer effects is especially relevant to the design of web platforms, as their design inherently dictates how peer effects unfold and information spreads, influencing how users behave on the platform. Furthermore, because platforms can control what peer behavior a user sees, they offer the ability for researchers to randomize peer observation, allowing for accurate measurements of peer effects in various settings. The detailed trace data which comes along with these online observations also signals an opportunity to understand in detail the mechanisms by which peer influence occurs. Analogous to (Anderson, Kleinberg, and Mullainathan 2017), who use a large scale chess database to closely examine what factors lead humans to make errors in decision making, we can study what factors, including observed experience, attenuate or amplify peer influence.

Peer influence could plausibly affect many separate parts of a decision making process. Consider a decision maker who is trying to choose a behavior from a set of options and observes a peer perform a certain behavior. The observation may influence the decision maker in the following ways:

- Information-Driven: If a decision maker does not know a behavior is possible, observing a peer do that behavior will modify the set of choices that the decision maker considers.
- Behavior-Driven: Observing a peer choose a particular behavior over other behaviors may increase the agent's posterior estimate of the payoff of the behavior, as they may infer the peer believes it is superior to alternative behaviors.
- Experience-Driven: If payoffs are observable by the agent, they may update their posterior belief about the payoff of an observed behavior with the observed payoff.

In this paper we do not study information-driven effects because users choose from a fixed set of displayed choices, but instead focus on disentangling behavior-driven and experience-driven peer effects. Prior empirical work has not typically focused on the role the observed outcome of a behavior plays in contagion, instead only investigating whether the behavior was observed. In reality, whether a behavior spreads may strongly depend on the observed outcome of peers. For a patient in a medical setting, even if many of their peers use a drug, they may not adopt it if they observe many of their peers suffer from adverse effects. One can easily imagine analogues in other domains: a user may re-share a news story a peer posts more readily (or only) if it receives many likes, or a dieter may adopt a diet more readily (or only) if they observe a peer on the diet lose a certain amount of weight. The idea that good payouts should be influential on adopting a behavior is consistent with a Bayesian updating model (Gallagher 2014; Banerjee 1992) where individuals use observed experiences of their peers to learn the payoff of various behaviors, and choose the behavior which maximizes expected payoff.

On the other hand, theoretical models of social influence often touch upon one of experience-driven or behavior-driven but not the other. For instance, much of the literature on social herding (Banerjee 1992) constructs models based on the experience-driven mechanism, updating posteriors from payoffs, but ignores that observing peer behavior independently of payout may affect decision making. On the other hand, many cascade models consider the behavior-driven mechanism, where nodes are "activated" into behaviors depending on the number or proportion of neighbors who are "activated", regardless of observed payoffs. The theoretical literature on on influence maximization allows for some heterogeneity in the strength of the tie/influence (Aral and Dhillon 2018), but largely uses simplifications and does not consider experience-driven and behavior-driven effects separately.

A recent exception to the neglect of the role of observed experience is a line of work examining the question of whether observed payouts affect behavior adoption beyond observed behavior in weather insurance adoption patterns. In rainfall index insurance, farmers receive a payout on their insurance policy if rainfall at a measured station is outside an acceptable range for growing crops, insuring them against catastrophic weather. A field experiment with farmers in rural India finds that farmers are more likely to adopt if policies purchased in the previous year by fellow villagers had high returns (Cole, Stein, and Tobacman 2014), suggesting experience-driven peer effects at an aggregate, village level. A separate field experiment with Chinese farmers (Cai and Song 2012) finds that observing peer farmers purchase index insurance without knowing previous payoffs has no effect on whether farmers adopt, although they do find spillover effects from information diffusion, suggesting that there there is not much behavior-driven peer influence. Together, these studies tentatively suggest that the main mechanisms by which farmers are influenced by peers into adopting rainfall insurance are information and payoff based, and that there may be payoff-based peer influence even in settings where there is no behavior-based peer influence, further emphasizing the importance of disentangling the two effects.

One heterogeneity in peer influence that researchers have thought about are the disparate impacts of observing different types of peers. For example, Granovetter makes the point that strong ties are individually stronger than weak ties, but the latter are more numerous, and thus more likely to provide job opportunities in aggregate(Granovetter 1977). Similarly, observing two peers who are not socially connected adopt a behavior may be less (or more) influential than observing two peers from the same social circle do a behavior (Ugander et al. 2012). Many of these can be thought of as positions of alters in the network, but independently from that, the social status of the observed peer may also play a role in the strength of the peer effect (Paluck, Shepherd, and Aronow 2016). While interesting, we do not study the role these structural factors play in this paper.

Understanding the role of payoffs in behavior contagion informs many decisions made by practitioners in the space. For example, if observing positive payoffs is the central mechanism mediating peer effects, the decision of whether a social intervention will work may be strongly informed by whether the intervention comes with a positive payoff, and further whether peers can easily observe those payoffs. It may also inform influence maximization efforts – if both observed payoff and behavior are important, jointly modeling both payoff and behavior may result in better performance compared to a contagion model based on only one mechanism. Instead of solely optimizing based on the network structure, this could result in seeding which more heavily prioritizes individuals who are predicted to have highly visible positive payoffs. In particular, strong payoff effects suggest that influence maximization may be more useful in settings where the network is strongly homophilous in expected payoff of the intervention.

In this paper, we identify a large-scale online setting where we can investigate in detail the roles that both behavior and payoff play in behavior adoption. We carry out our investigation on the decision of champion (character) selection in League of Legends (LoL), the most popular video game in the world, with over 100 million monthly and 8 million daily active users (Volk 2016) (Goslin 2019) in 2019. Online settings are widely used to study social influence due to their large size and detailed trace data. Furthermore, games such as Second Life (Bakshy, Karrer, and Adamic 2009) have been used to study social influence in settings where users organically make decisions, and LoL itself has been used to study collective intelligence within teams (Kim et al. 2017). We applied and were granted access to the official LoL API, which we used to examine the match history of a month of games for 20,000 players, observing which champions each player and their counterpart on the opposing team chose, giving us the observed *behavior*. We further recorded whether their counterpart won or lost the game, giving us the observed *experience*. We then ask: how does observing behavior and experience affect which champions users select? To answer this question, we define a test statistic measuring how often players choose the champion they just observed, compute its empirical value, and generate its distribution under a range of null hypotheses where we vary whether there are experience-driven and behavior-driven peer effects. We find that there are both behavior-driven and experience-driven peer effects in LoL champion selection, where players are more likely to choose a champion if they observe it in their previous game. Further, we find that the effect is diminished (but still positive) when the champion is observed losing, and augmented when the champion is observed winning.

League of Legends

There are three stages to playing a game of LoL: Matchmaking, Champion Selection, and Multiplayer Online Battle Arena. First, players enter the matchmaking queue either by themself or with a single friend. Each player chooses a primary and secondary position out of the five total positions. From all the players in the queue, the matchmaking algorithm attempts to choose 10 players with similar Elo, and assigns the chosen 10 players to two teams of five (Isto 2013). It also assigns the players to the 5 unique positions per team based on their expressed preferences. We take advantage of the randomization offered by matchmaking: controlling for time and skill, users are randomly assigned to games, and thus, peers to observe. Further, LoL's matchmaking system highly values balancing the two teams in a game, preferring to 'wait a little longer in the queue to get a fairer match' (Isto 2013).

In the player selection phase, each player simultaneously is given the option to choose one champion to ban. Then, the players take turns choosing characters out of the pool of 131 remaining champions, of which each champion may be chosen at most once. This champion selection decision is our main behavior of interest. In particular, we focus on the champion chosen by the opponent with the same position as our player of focus, and see if our player adopts that champion by selecting it in the following game. We also focus our analysis on games played in one out of the five positions, top lane, which spends the most time during the game fighting one-on-one with the opposing counterpart, due to differential champion popularity per position[1].

In the Multiplayer Online Battle Arena stage, all 10 players control the character they chose in a shared Battle Arena. Simply put, members of each of the two teams work with their teammates and attempt to destroy the other team's base while preventing the other team from destroying theirs. Games typically last for around 30 minutes. This is the phase where the player of focus observes a payoff from their opponent's choice. In particular, we use the simplest notion of payoff, whether their opponent wins the game.

Research Questions

We are interested in whether there are peer effects for champion adoption: does observing an opponent pick a champion increase the probability that a player chooses that champion in their subsequent game? In other words, are there behavior-driven peer effects? Further, we are interested in the role of observed experience: if the player observes their opponent having a good experience with their champion selection, are they more likely to adopt? Here we measure experience by observing if their opponent won the game. In addition to studying payoff, our main construct of interest, we also examine the role skill plays in moderating peer effects. In particular, we want to understand if increasing sophistication with a system interacts with the role of peer effects, which occurs for other decision-making heuristics such as loss aversion (Haigh and List 2005). Additionally, scraping users of varying rank allows us to check for heterogeneous effects of payoff by sophistication. The Bayesian learning interpretation of peer effects suggests this may be the case because more sophisticated users may have stronger priors on payoffs of choices, and their posterior beliefs may not update as much from a single instance of data compared to a less sophisticated user.

Data

We examine the match histories for one month of play (September 2018) for the players ranked 1-10,000 and 100,000-110,000 on the ranked ladder for the North American LoL server (at the start of the month)[2]. We use the LoL API to scrape their match history, obtaining records of 468,112 games by 13,190 players, excluding the $\sim 5\%$ of games where the Riot API failed to record the position of the opponent. For each match, we observe the champion c the player chose, the champion o the opponent chose (again, focusing only on players and opponents in the "top lane" position), the champion s the player chose in their subsequent game, k which denotes whether the player is among our high skilled (rank 1-10,000) or low skilled (100,000 - 110,000) sample, and the observed payoff e_i (whether the opponent won or lost). Thus, each point of data is a tuple of form $(c_i, o_i, s_i, k_i, e_i)$. We call the set of all observed games $\mathcal{D} = \{(c_i, o_i, s_i, k_i, e_i)\}$. To avoid merely detecting when an opponent chooses the character that our player wanted to play, we filter our data to only contain games where the user picks first: that is, c is picked before o and thus there is exogenous variation in o among users who choose the same c. Table 1 and figure 1 summarize our filtered data, showing the number of scraped players, how many games they played in our month of data in the top lane where they chose first, and how often each champion was played.

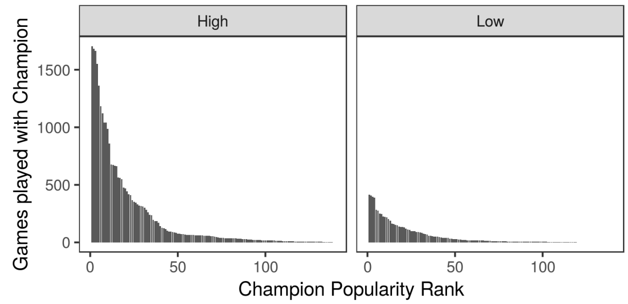

Figure 1: The number of games users played with each champion by skill level, with champions organized by popularity, showing the heterogeneity in champion popularity.

Statistic	High Skill	Low Skill
Total Games	28817	7433
# Active Players	3705	2060
Mean GPP	7.78	3.61
SD GPP	10.84	5.08

Table 1: Summary statistics of scraped games where our scraped player chose first and played the "top lane" position, where GPP stands for games per player.

Analysis

After playing a game where they've observed the champion their opponents picked and the outcome of that game, we are interested in whether the players in the dataset choose, in their subsequent game, the same champion: in other words, whether $s_i = o_i$. To this end, we define a test statistic m, the proportion of games per player where $s_i = o_i$, averaged

[1]League of Legends teams have 5 members, each playing a different position: top lane, mid lane, bottom lane, support, and jungle. Each position corresponds to different parts of the map and different champions are popular for different positions.

[2]Upon publication the authors will make the full dataset and code publicly available.

over all players, which will be central to the analysis:

$$m(\mathcal{D}) = \frac{1}{P}\sum_{p=1}^{P}\frac{1}{n_p}\sum_{i=1}^{n_p} 1_{s_{p,i}=o_{p,i}}.$$

Within the full data $\mathcal{D}$, we have P unique players and n_p games played by player p, where $s_{p,i}$ and $o_{p,i}$ are the champions that player p chose in their subsequent game and observed in their current game respectively. Informally we call m the *match proportion*.

Randomization Inference

In order to understand if we are truly observing any sort of peer effect in the observed data, we compare the observed value of m to its distribution under the null hypothesis that there are no effects of peer observation or observed payoff on champion selection. To compute such a null distribution, we can take the observed data $\mathcal{D}$ and create randomized dataset $\mathcal{D}'$. In $\mathcal{D}'$, we resample the champions and payoffs that each player observes. That is, for each game $g_i = (c_i, o_i, s_i, k_i, e_i)$ in $\mathcal{D}$, we generate a corresponding tuple in $\mathcal{D}'$ as $g'_i = (c_i, o'_i, s_i, k_i, e'_i)$. Then, we compute the distribution of $m(\mathcal{D}')$ over many instances of $\mathcal{D}'$. In other words, we ask the question: what would the distribution of the test statistic be if there were no peer effects and payoffs did not affect behavior adoption?

The remaining technical question is how to sample o'_i and e'_i. First, we resample o'_i from the distribution of observed champions from games played at the same skill level in the same week. This design avoids selection effects caused by differences in the distribution of observed champions by skill or trending popularity of champions over the month of observations. Then, we resample e'_i from the empirical distribution of e_i in games played at the same skill level in the same week where the player observed champion o'_i. That is, we sample[3]

$$e'_i \sim \text{Bern}\left(\frac{\sum_{i=1}^{n} e_i 1_{o_i=o'_i}}{\sum_{i=1}^{n} 1_{o_i=o'_i}}\right)$$

In addition to the *full null* hypothesis, described above, where neither the observed champion nor payoff matters, we consider two weaker null hypotheses:

1. Behavior-driven null: The observed champion matters, but not the payoff.
2. Change-driven null: The payoff matters but not the observed champion.

The behavior-driven null corresponds to a purely *behavior-driven* peer effect, where peer effects are driven solely by observing the behavior of peers (i.e. the users are influenced by what champions their peers chose, but not whether they observed the peers winning with them). The second hypothesis, the change-driven null, corresponds to a model where user behavior changes according to whether they win or lose, irrespective of which champion they observe. The purpose of this null is to disentangle two payoff effects which are seemingly conjoined in this setting. Because one team wins and the other team loses in LoL, whenever a player observes the opposing champion win, that player must have lost the game. This latter null captures the possibility that players may simply get frustrated when they lose and pick a different champion than the one they played, but which champion they lose to does not affect their choice. We might expect m to be higher under losses under the change-driven null than under the full null hypothesis because champions are selected without replacement, so the champion the player observes cannot be the same as the one they were originally playing. Therefore, if players tend to change champions after a loss, even without a predilection for the champion they observed over other champions, they might still be more likely to choose the observed champion because it's not the one they were originally playing. Under this null, there is still a payoff-driven peer effect – it's just that the impetus is to simply change one's behavior (thus the name *change-driven*), not to adopt the observed behavior.

To simulate m under these two hypotheses, we change how we resample o'_i and e'_i from the full null. Under the behavior-driven hypothesis, we simply let $o'_i = o_i$, as we want to preserve the original observed champion, but resample e_i in the same manner as before. On the other hand, for the change-driven hypothesis we resample o'_i as before, but do not resample e_i. In this manner, we preserve the original payoff observed but modify the observed champion, in line with the hypothesis that the observed champion does not matter since the player is only being influenced by losing.

In addition to varying the null hypothesis under which we simulate the distribution of m, we also vary the set of games $\mathcal{D}$ we compute m over. For instance, we can separate games played by low skill and high skill players, $\mathcal{D}_l, \mathcal{D}_h \subset \mathcal{D}$, games where users observed a positive or negative experience (win or loss), $\mathcal{D}_1, \mathcal{D}_0 \subset \mathcal{D}$, or even $\mathcal{D}_c$ for games where a certain champion is observed. Importantly, the randomization is always done at the level of $\mathcal{D}$, the full dataset, and then $\mathcal{D}'$ is subsetted to match the subset of interest.

Consequentially, some nulls only affect the distribution of m over certain subsets. First, $m(\mathcal{D}_x) = m(\mathcal{D}'_x)$ under the behavior-driven null unless x includes a subset by observed experience. This is because under the behavior-driven null, $1_{s_{p,i}=o_{p,i}}$ remains the same for all games, so m only changes if the set of games it is computed over changes. Then, m will only vary if resampling the observed experience will change the games included in the subset of interest. For instance, whether a game goes into $\mathcal{D}'_1$ and $\mathcal{D}'_0$ depends on which experience is sampled for that game. Thus, we only consider the behavior-driven null when we subset the data by experience. Similarly, the change-driven null will not yield a different distribution of m compared to the full null unless we subset by experience: in both nulls we resample champions, and then as per above further resampling experience (as in the full null) only has an effect if we subset by experience.

[3]Sampling e'_i from a Bernoulli distribution parametrized by the mean observed experience can be thought of as a parametric bootstrap. We prefer this over shuffling because of the small number of games where rare champion are observed, where shuffling may not adequately express the true randomness of observed experience.

Peer Effect

First we compute the test statistic m on the entirety of the dataset and its distribution under the null hypotheses. The observed $m(\mathcal{D}) = 0.0210$ and the 95% CI is $[0.0134, 0.0177]$ under the full and change-driven nulls (recall these yield the same distribution of m). We find that the observed m is higher than expected under these nulls, meaning that a user observing a peer play a champion increases the chance that in the following game the user chooses that champion.

We also compute m along with its null distributions for the data subsetted by skill (low skill vs. high skill). We compute $m(\mathcal{D}_h) = 0.0206$ with a null 95% CI of $[0.0124, 0.0172]$ under the full and change-driven nulls and $m(\mathcal{D}_l) = 0.0217$ with a null 95% CI of $[0.0127, 0.0212]$. We find that the observed m is higher than expected under the null for both skill levels. The difference between m and its null mean is higher for high skill users, as shown in table 2, although the difference in differences is not statistically significant.

Next, we compute m along with its null distributions for the data subsetted by payoff (observed win vs observed loss), shown in figure 2. We find the observed m is higher than expected under both the change driven and full null hypotheses, meaning that we observe peer influence for champions observed both winning and losing, suggesting that observed payoff is not the only mechanism by which peer influence occurs in this setting. However, we note that the difference is far larger for games where the user observes a win than in those where the user observes a loss: this difference is quantified in the "Observed win vs loss" comparison in table 2 and is statistically significant. We can further test this line of inquiry by comparing m to its distribution under the behavior-driven null hypothesis, where we only resample whether the user observed a win or a loss. The observed m for games where the user observed a champion win is higher than than it is under the behavior-driven null and m for games with observed losses lower, meaning that users who observe a champion win are more likely to choose the champion in their following game than those who observe it lose.

Finally, we compute m for each of the four conditions of (high skill, low skill) $\times$ (observe opponent win, observe opponent loss). We compare these values with the distribution of m under all three null hypotheses. Figure 3 shows the results of this analysis. We first consider the change-driven and full nulls. We find that directionally, all point estimates for m are higher than their distributions under these null, although not all differences for games with observed losses are statistically significant at the $p < .05$ level. This again suggests that there is peer influence in this system regardless of observed payout. For both skill levels there is a larger difference between the observed and null estimates of m for the observed wins than for the observed losses (quantified in table 2), although the difference is only statistically significant for the high skill users. Finally, we can again examine the behavior driven null to look for evidence of experience-driven peer influence. In both skill levels we directionally find that observing a win leads to higher peer influence while observing a loss leads to lower peer influence, although again only in the high skill domain is it statistically significant.

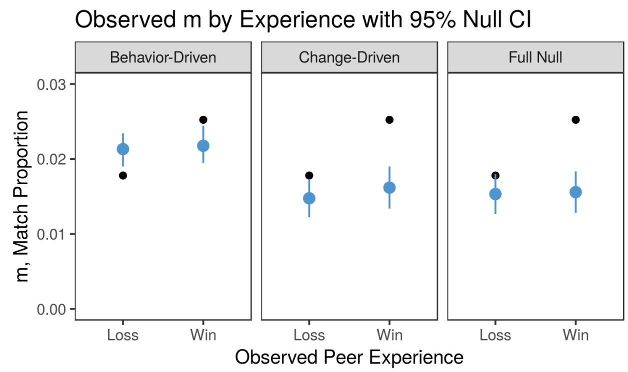

Figure 2: The match proportion m by observed payoff, with 95% null confidence intervals in blue. The point estimates are above the CIs for the change-driven (middle) and full null (right) hypotheses, indicating peer effects in all skill $\times$ payoff conditions and that payoff is not the only mechanism at work here, although the higher difference in the observed win conditions suggests that observing a positive payoff increases the strength. This is further confirmed by examining the relation of the point estimates to the null CIs from the behavior driven hypothesis (left), which shows that the point estimate for m is higher under wins and lower under losses than it would be if payoffs did not affect m.

By Champion

In addition to aggregating the strength of peer effects over all champions, we can also do the analysis on individual champions. In particular, we are curious whether peer effects vary in some manner depending on the rarity of the champion. In figure 4 we plot the difference between the observed m and its mean under the full null distribution for all champions. We find that in general, more popular champions have larger boosts in match proportion compared to less popular champions - that is, the difference between the observed m and the m under the full null for games where a certain champion was observed was higher for already popular champions. Taking the point of view of a random utility model, it could be that these champions were already close to the threshold for being chosen. As popular champions, we might expect many people to already have high utilities for them, which were pushed beyond the choice threshold by the peer observation. This phenomenon suggests a "rich get richer" model for the role of peer effects in this setting, where champions who are already more popular and thus observed more also get a larger boost from peer effect per observation because there are more users on the margin for being converted.

Conclusion

In this paper we investigate in-depth the roles that user skill and observed experience play in peer effects in the choice of champion selection in League of Legends. We find positive

	Null	Delta	SE	p-value	Comparison	Skill	Observed Payoff	Signif
1	Behavior-Driven	-0.00001	0.00218	0.49859	High to Low skill			
2	Change-Driven	0.00100	0.00339	0.38412	High to Low skill			
3	Full Null	0.00111	0.00331	0.36834	High to Low skill			
4	Behavior-Driven	0.00689	0.00242	0.00220	Observed win vs loss			**
5	Change-Driven	0.00591	0.00260	0.01162	Observed win vs loss			*
6	Full Null	0.00708	0.00262	0.00344	Observed win vs loss			**
7	Behavior-Driven	0.00873	0.00271	0.00064	Observed win vs loss	High		***
8	Change-Driven	0.00772	0.00297	0.00468	Observed win vs loss	High		**
9	Full Null	0.00889	0.00299	0.00146	Observed win vs loss	High		**
10	Behavior-Driven	0.00311	0.00473	0.25516	Observed win vs loss	Low		
11	Change-Driven	0.00221	0.00516	0.33434	Observed win vs loss	Low		
12	Full Null	0.00338	0.00518	0.25709	Observed win vs loss	Low		
13	Behavior-Driven	-0.00281	0.00364	0.22000	High to Low Skill		Loss	
14	Change-Driven	-0.00213	0.00397	0.29557	High to Low Skill		Loss	
15	Full Null	-0.00197	0.00406	0.31385	High to Low Skill		Loss	
16	Behavior-Driven	0.00280	0.00406	0.24483	High to Low Skill		Win	
17	Change-Driven	0.00337	0.00444	0.22358	High to Low Skill		Win	
18	Full Null	0.00354	0.00439	0.21004	High to Low Skill		Win	

Table 2: Comparisons of how the observed value of m compares to its mean under the null distribution, for each null and comparison. We find that the difference between m and its null expectation is significantly higher in games where the user observes a win rather than a loss for all three nulls. Further, we find that this difference is mainly driven by high skill games, although the results are directionally the same in low skill games (but not statistically significant). We find no significant differences in the difference between m and its null expectation between high skill and low skill games, even when separately considering games with an observed wins and losses.

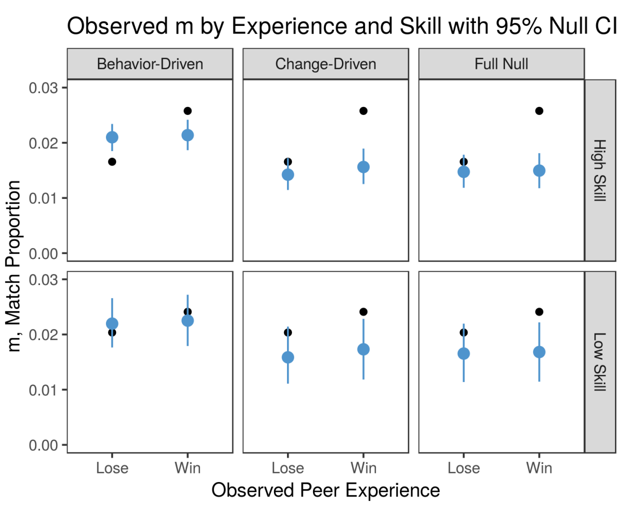

Figure 3: The match proportion m for each of the 4 skill $\times$ payoff conditions, with 95% null confidence intervals in blue. Comparing the PEs to the behavior driven null CIs, we find that observing a win increases the peer effect whereas observing a loss decreases it, although the result is only statistically significant for high skill users, suggesting experience-driven peer effects. We find that directionally m is higher than we would expect under the full or change-driven null, although not all results are statistically significant, suggesting the presence of peer effects not driven by experience as well.

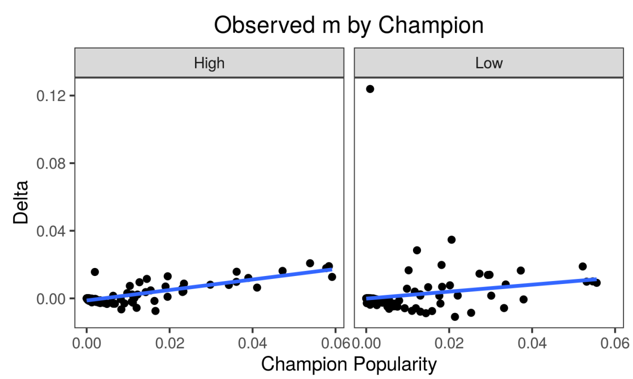

Figure 4: Difference between observed m and its mean (labeled Delta) under the full null, by champion popularity. We find a positive correlation between champion popularity and m, suggesting that the most popular champions benefit the most from the peer effect. The outliers and steadily decreasing points for popularity range 0.00 to 0.02 are artifacts resulting from unpopular champions being observed infrequently and thus having small denominators in inference: for example, one outlier has an observed m of 1/8. Correspondingly, many observed m are 0. However, under the null, champions have monotonically decreasing null m as popularity increases. Thus, below popularity 0.02, many champions have empirical m of 0 but a steadily decreasing null m, corresponding to steadily decreasing points. Inferences on individual unpopular champions have low power but are not heavily weighted in aggregate effects since unpopular champions are rarely observed.

evidence of peer effects in this setting, and that the peer effect is stronger when observing a win than a loss, although still positive in both instances. We rule out the possibility that the effects are caused by losing the game versus observing a win by comparing the observed test statistic to that of the change-driven null, instead attributing them to observed experience . We find that the average strength of the peer effect remains roughly the same between users of varying skill level, although the data suggests that users of high skill may care more about the observed experience than low skill users. Returning to the breakdown of the possible effects observing a peer could have, the attention, behavior, and experience-driven hypotheses, we find strong evidence of experience-driven peer effects which are responsible for almost the entire peer effect for high skilled users and roughly half the peer effect for low skilled users. The remaining peer effect could be attributed to either the attention or behavior-driven hypotheses, but the data leans towards the behavior-driven hypothesis. Evidence of this comes from observing the peer effect by rarity of champion, where we find that for rarely played champion there are less peer effects than for more popular champions. Under the attention driven hypothesis, we would expect to see higher peer effect for rare champions, as those are called into attention less frequently. The finding that observing a popular behavior is more likely to marginally tip an individual into performing the behavior than observing an unpopular behavior, combined with the tautology that more popular behaviors are observed more, suggest a "rich get richer" model of peer effects, where peer influence mainly serves to make already-popular behaviors get more popular.

We believe that this work informs efforts in influence maximization and social interventions. The importance of observing positive experiences suggests efforts to maximize influence of interventions should focus on seeding users with strongly visible positive experiences, especially ones connected to other such individuals. For example, a medical intervention might be best seeded in a cluster of individuals who will likely exhibit positive health outcomes to their peers. Furthermore, understanding the role observed experience plays allows us to understand when social interventions may or may not be expected to work. If experiences are not easily observed (for example, in single-instance decisions like whether to apply for college, when most social ties are between individuals who are deciding at the same time), social interventions may be less successful because they can only appeal to behavior-driven effects, which may be weaker than experience-driven effects. We hope that this paper inspires further work in identifying what factors, including status or foreignness of the behavior, correspond to individuals adopting peer behavior, and how they interact with observed experience.

References

Anderson, A.; Kleinberg, J.; and Mullainathan, S. 2017. Assessing human error against a benchmark of perfection. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 11(4):45.

Aral, S., and Dhillon, P. S. 2018. Social influence maximization under empirical influence models. *Nature human behaviour* 2(6):375.

Aral, S., and Nicolaides, C. 2017. Exercise contagion in a global social network. *Nature communications* 8:14753.

Bakshy, E.; Eckles, D.; Yan, R.; and Rosenn, I. 2012a. Social influence in social advertising: evidence from field experiments. In *Proceedings of the 13th ACM conference on electronic commerce*, 146–161. ACM.

Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012b. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, 519–528. ACM.

Bakshy, E.; Karrer, B.; and Adamic, L. A. 2009. Social influence and the diffusion of user-created content. In *Proceedings of the 10th ACM conference on Electronic commerce*, 325–334. ACM.

Banerjee, A. V. 1992. A simple model of herd behavior. *The quarterly journal of economics* 107(3):797–817.

Cai, J., and Song, C. 2012. Insurance take-up in rural china: Learning from hypothetical experience. *Available at SSRN 2161649*.

Centola, D., and Macy, M. 2007. Complex contagions and the weakness of long ties. *American journal of Sociology* 113(3):702–734.

Centola, D. 2010. The spread of behavior in an online social network experiment. *science* 329(5996):1194–1197.

Cole, S.; Stein, D.; and Tobacman, J. 2014. Dynamics of demand for index insurance: Evidence from a long-run field experiment. *American Economic Review* 104(5):284–90.

Gallagher, J. 2014. Learning about an infrequent event: evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics* 206–233.

Goslin, A. 2019. League of Legends has nearly 8 million peak daily concurrent players globally, says Riot. https://www.riftherald.com/2019/9/17/20870382/league-of-legends-player-numbers-active-peak-concurrent. [Online; accessed 8-October-2019].

Granovetter, M. S. 1977. The strength of weak ties. In *Social networks*. Elsevier. 347–367.

Haigh, M. S., and List, J. A. 2005. Do professional traders exhibit myopic loss aversion? an experimental analysis. *The Journal of Finance* 60(1):523–534.

Hoxby, C. 2000. Peer effects in the classroom: Learning from gender and race variation. Technical report, National Bureau of Economic Research.

Isto. 2013. Matchmaking Guide. https://support.riotgames.com/hc/en-us/articles/201752954-Matchmaking-Guide. [Online; accessed 8-October-2019].

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146. ACM.

Kim, Y. J.; Engel, D.; Woolley, A. W.; Lin, J. Y.-T.; McArthur, N.; and Malone, T. W. 2017. What makes a strong team?: Using collective intelligence to predict team performance in league of legends. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2316–2329. ACM.

Paluck, E. L.; Shepherd, H.; and Aronow, P. M. 2016. Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences* 113(3):566–571.

Ugander, J.; Backstrom, L.; Marlow, C.; and Kleinberg, J. 2012. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences* 109(16):5962–5966.

Volk, P. 2016. League of Legends now boasts over 100 million monthly active players worldwide. https://www.riftherald.com/2016/9/13/12865314/monthly-lol-players-2016-active-worldwide. [Online; accessed 8-October-2019].

Zhang, J.; Brackbill, D.; Yang, S.; and Centola, D. 2015. Efficacy and causal mechanism of an online social media intervention to increase physical activity: results of a randomized controlled trial. *Preventive medicine reports* 2:651–657.