Social Influence in Online Reviews: Evidence from the Steam Store*

Adam Di Lizia †

March 19, 2025

Abstract

How does social influence affect consumer ratings? Using a dataset from the popular *Steam* gaming platform I investigate how quality judgements depend on pre-existing consumer assessments. In 2019, *Steam* introduced a new review system which decreased the exposure of users to previous ratings. Firstly, I find that user ratings are dependent on average ratings. The result is not due to selection, and is robust to a wide range of alternative specifications. Secondly, the effect is heavily asymmetric: individual reviewers are more negative when exposed to a lower average rating, but do not respond to a higher one. This negativity compounds and inflates the gap between lower rated and higher rated games. Overall, these social influence effects are driven by less experienced users on the platform. I rationalize these results with a model of reviewer behavior. Finally, using owner data, I estimate a structural model of game choice. A 1% increase in rating is equivalent to a 2.5 dollar price reduction. This suggests social influence has large implications for buyers and sellers, and that it reduces efficiency across online platforms.

*I would like to thank Mirko Draca, Ao Wang, Lily Shevchenko, Manuel Bagues, Carlo Schwarz, Chris Roth, Dan Bernhardt, Stefano Caria, Nathan Canen and fellow PhD Students for their advice and feedback. I am grateful to the ESRC for helping fund this research agenda. All mistakes are my own.

†Department of Economics, University of Warwick

1 Introduction

User reviews are an essential part of the online experience, and vital in understanding how purchasing decisions are made. A platform's reviews are crucial to users choosing between products, and are arguably one of the biggest benefits to shopping online. It has been shown that these reviews have a large, causal impact on purchase decisions (Anderson and Magruder 2012; Reinstein and Snyder 2005), and that firms care about these reviews enough to attempt to manipulate them (Mayzlin, Dover, and Chevalier 2014). The online market is both large and growing[1] and therefore understanding how reviews are generated is important for understanding the decisions made by the modern consumer.

In theory, these reviews should be representative and informative about a product's quality. This does not often hold in practice. On many platforms, users are exposed to the ratings of others when leaving their review. The issue explored in this paper is whether those opinions influence reviewers. This "social influence" could persist across many platforms and harm the ability of a consumer to accurately gauge product quality.[2]

Examining these issues empirically is difficult. This paper's key contribution is leveraging an exogenous shock in the information available to users to causally estimate social influence effects. I also analyse who is most susceptible to social influence, and how social influence affects sales.

I use data from *Steam*, a popular gaming platform and the main retailer for PC games, to identify how the average rating affects user reviews. *Steam* is a large platform that distributes video games over the internet. It is both a substantial market in its own right and is typical of the review systems across the internet. Further description of *Steam*, its structure and content can be found in the next section and Appendix A. Importantly, only those who own the game can leave a review.

An update to the platform in 2019 resulted in reviewers no longer seeing the average rating by default when leaving their own review. Before users reviewed via the store page and thus were exposed to the average rating. Since the update, users are able to review directly, without visiting the store, avoiding the average rating.[3]

I show that this exogenous change in the information set has substantial impacts on reviewing behavior. A 10 percent[4] increase in average rating increases the probability a review is positive by 5.4% before the policy change, but only by 2.8% after. This suggests user ratings do depend on the average rating. A game rated at 60-70% positive is on average 13% more likely to be rated positively after the policy, but there is very little change for any game rated above 80% positive. Therefore compounding negativity is the primary cause of social influence bias in this context. In a simple structural model I show that this change to rating affects sales. Consumers value a 1% increase in rating equivalently to a 2.5 dollar decrease in price, showing that social influence does affect product choice.

It is perhaps unsurprising that exposure to information would lead consumers to take this information into account. However, a straightforward herding story does not suffice here for two reasons. Firstly, we consistently see users leave reviews dif-

[1] The fraction of all purchases made that are online is around 16%

[2] Selection bias is also a key issue. This will be the topic of future research.

[3] Care will be taken to ensure selection of different types of reviews do not cause the results in section 5

[4] In absolute terms as average rating is measured as a percent in this setting

ferent to an average rating that is generated by potentially millions of users. Neither does herding explain the strong asymmetry of negativity compounding. I rationalise user behavior with a simple model where consumers are unsure of the true quality of the game when they go to review. Average ratings are unlikely to sway users who strongly feel they had a good or bad experience, but might be useful for reviewers who are on the border between whether to recommend a product or not. The asymmetry then follows immediately from lower ratings being rarer and thus surprising reviewers substantially more. Alternatively, users may be risk averse in their product recommendations, happy to go against the crowd not to recommend a product but unwilling to recommend something many users did not like.

Related literature This paper adds to the literature examining the importance of social influence across online platforms. Muchnik, Aral, and Taylor (2013) experimentally up-vote or down-vote comments on a news aggregator site and find an initial positive vote leads to a higher probability the comment is voted positively by others. Aral and Nicolaides (2017) use data from the *Strava* running app and exogenous variation in weather to show that users exercise more when their peers do. More general herding effects across online reviews have been examined in the quantitative marketing literature. Lee, Hosanagar, and Tan (2015) and Sunder, Kim, and Yorkston (2019) both show that user rating and the rating of friends consistently move together across a movie website and board game website respectively. However, this could be driven by the similarity between users[5]. In my setting this is controlled for allowing the precise social influence effect to be identified.

This work sits as part of the larger peer effects literature. Manski (1993) shows the challenges of identification in this setting, which this paper will deal with through quasi-experimental variation. More specifically, this work belongs to the reduced-form style as described by Sacerdote (2014). This type of work often looks at peer effects and student performance, such as J. D. Angrist and Lang (2004) or Imberman, Kugler, and Sacerdote (2012), using natural experiments for identification. However, a key difference is that rather than looking at exogenous assignment to groups, my work analyse the exogenous treatment effect of a policy which removes information about a users reference group. The most similar type of paper is Bursztyn and Jensen (2015) who leverage a digital policy change across online courses where students were suddenly exposed to a leader-board, and they show this information worsens student performance by quartile.

It is important to stress the distinction here, as Angrist (2014) shows that the usual linear-in-means peer effects estimator often has no real interpretation whether it is identified or not. While here the average rating is a (slightly different) mean of user rating, it is a real variable and the variation is in whether it is displayed to users. This avoids the main issues mentioned by Angrist (2014) as I am estimating the treatment effect of the policy as this average rating varies, rather than a specific peer effect.

Structure. The paper is structured as follows. Section 2 describes the background of Steam as a platform and the policy change. Section 3 onwards describes the data, framework and results, and Section 8 concludes.

[5]Even when exposure to the average rating is removed in my setting, there is still a large positive correlation between average rating and user rating suggesting that similarity between users is a key reason for this correlation. See Table 1.

2 Background

2.1 Broad Overview

Steam was founded by *Valve*, a gaming company, in order to automatically provide updates for their products over the internet. Eventually this evolved into a service where entire games could be sold without the need for a physical copy. Initially, only Valve-made games were offered on the platform, but by 2005 other game companies started distributing their products on *Steam*.[6] This made it easier to buy and sell games, and has become the main way to distribute PC games.

While competitors do exist, such as *Epic Games*, *Steam* enjoys an almost monopolistic status with around 75% of the market for PC game distribution. The value of *Steam* has expanded rapidly in the last decade as video games have become more digital. The service generates around $6 billion in revenue and has 132 million active users circa 2020. Thus while this research aims to provide a more general understanding of online reviews, *Steam* is an incredibly valuable platform in its own right. *Grand Theft Auto 5*, a popular game sold on Steam, is the highest grossing media product of all time[7]. For a further discussion on the types of games on *Steam* and *Valve*'s goals as a platform, see Appendix A.

2.2 Steam Reviews

User reviews were introduced to the platform in 2011.[8] Reviews are binary: users can either recommend or not recommend a game. Users can also write text explaining their rating. This information is aggregated to the consumer as an average rating (% who recommended) across all reviews and reviews left in the last 30 days (since 2016). Consumers can also search through the history of all posted reviews. These reviews display the reviewer's username, the number of reviews they have left, number of games they own, playtime at review and the text of the review. Users must own the game in order to leave a review.[9] A valuable feature of the platform is that games are tied to an individual's account, so the average user is unlikely to switch between accounts.

2.3 Policy Change

The biggest change to this review system occured as part of the "Library Update" on October 31, 2019. Before the update, the reviewer was directed to leave a review

[6] In return they give Valve a 30% cut of all sales. This is 20% for larger games.

[7] https://www.marketwatch.com/story/this-violent-videogame-has-made-more-money-than-any-movie-ever-2018-04-06.

[8] Originally, only the metacritic score - calculated by taking the average across all verified game critic reviews for a game- was displayed. It is unchanging once all critics have reviewed a given game, usually in the first week or two after release.

[9] This can occur by purchasing the product directly from Steam, or purchasing a Steam Key - a code giving them access. Only verified Steam purchases are aggregated into the average rating. In addition Steam for some games chooses not to incorporate reviews as part of review bombs (off topic review activity rating a game as negative for actions of the publisher or developer rather than game quality) into the average rating. These reviews can be identified while scraping to best recreate the closest rating measure. Nevertheless, as shown in Figure 3 and Table C.2, the effect is not sensitive to small changes in how rating is measured.

via the store page[10]. This resulted in the user seeing the average rating.[11] After the update, when exiting a game played for more than a few hours users are given the option on the page from which the game was launched to leave a review. Users can post a review entirely within a pop-up box without being exposed to the average rating.[12] The old system is still present, but is more difficult to use. Even if some reviewers still use the old system, or seek out the average rating regardless, it unlikely to be common behavior and at worst would attenuate the results.

It appears the aim of the update was to increase the sheer number of reviews. *Steam* notes in a blog post how the daily number of reviews has tripled after the update.[13] However, this also provides an excellent setting for exploring how reviews are generated and how platforms may influence this process.

3 Data

3.1 Review Data

Review data is scraped directly from the public *Steam* store pages of individual games. It includes user rating, text of review, time of review, username and other supplementary information. I collect all English language reviews for the top 250 games by their number of reviews in May 2019.[14] I obtain the history of reviews from a game's release to November 2023. The top 250 games represent around 20 % of all reviews left on the platform as a whole.[15] The raw data is around 20 million unique reviews, and the main regressions are run on approximately 6.5 million reviews after cleaning and limiting the sample to those present before and after the policy.

3.2 Average Rating

The rating at time of review is not observed, but can be recovered as *Steam* would have calculated and presented it by averaging over all previous reviews from *Steam* purchases at the time of review. This is an imperfect measure due to some reviews being updated[16] and a restriction to English language reviews. However, both are a small fraction of total reviews. In addition, for games released after 2016, I have daily ratings from *SteamSpy*. These ratings include non-English and unedited reviews and using them does not substantially change the results.

3.3 Text Analysis

Additionally, I calculate review length and sentiment of review text. Review length is the number of characters in a review. The sentiment measure, *Spacy text blob*, is taken

[10]See https://store.steampowered.com/reviews/ as the initial announcement directing users to do this

[11]See Figure A.2

[12]See Figure A.3

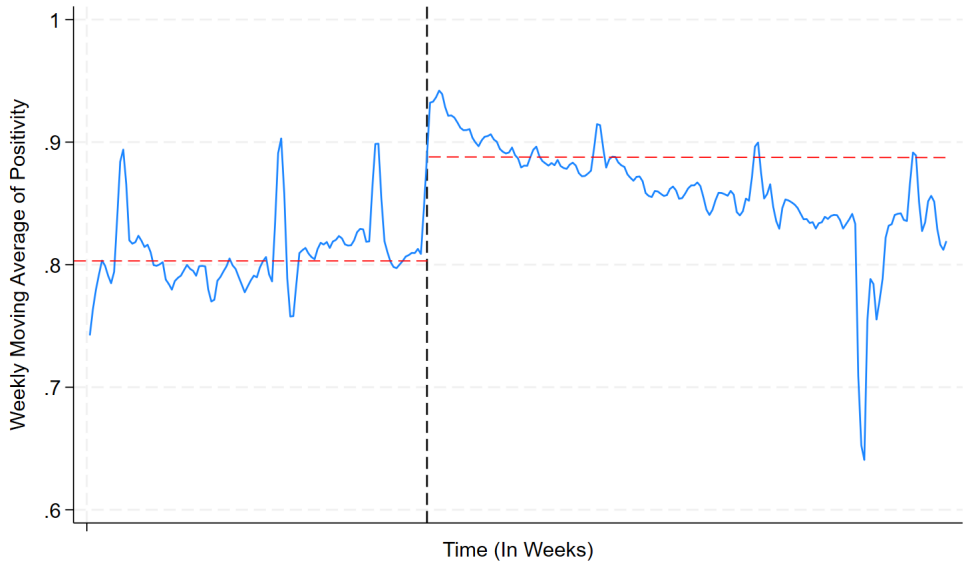
[13]See https://store.steampowered.com/news/group/4145017/view/1697229969004830849

[14]Due to one of these games being re-branded as a 'game of the year edition" the final number of games in my sample is 249

[15]Calculated as a percentage of all reviews in November 2023

[16]Steam users can update their review after leaving it. When scraping reviews, only the most recent versions of reviews are seen. Since the observed characteristics of the review occured at the update, I treat updated reviews as being created at their date of update, but results are robust to alternate codes or dropping all updated reviews from the data set.

Figure 1: Change in Positivity



Change in weekly average positivity across all games in the sample around the policy change, denoted by the dashed black line. This is the effect only on reviewers present before the policy change and after so is not driven by selection or attrition. This is a 3 period weekly moving average of the positivity of reviews across all games in a given week. The large spikes are related to sales or new game releases. In the full specification, these will all be controlled for.

directly from the popular *Spacy NLP library*, and is a measure of how positive the text is. It can differentiate between negation i.e "not good" is correctly recognized as a negative phrase. Sentiment supplements the binary rating with a continuous measure of like and dislike.

3.4 Price, Player Number and Sales Data

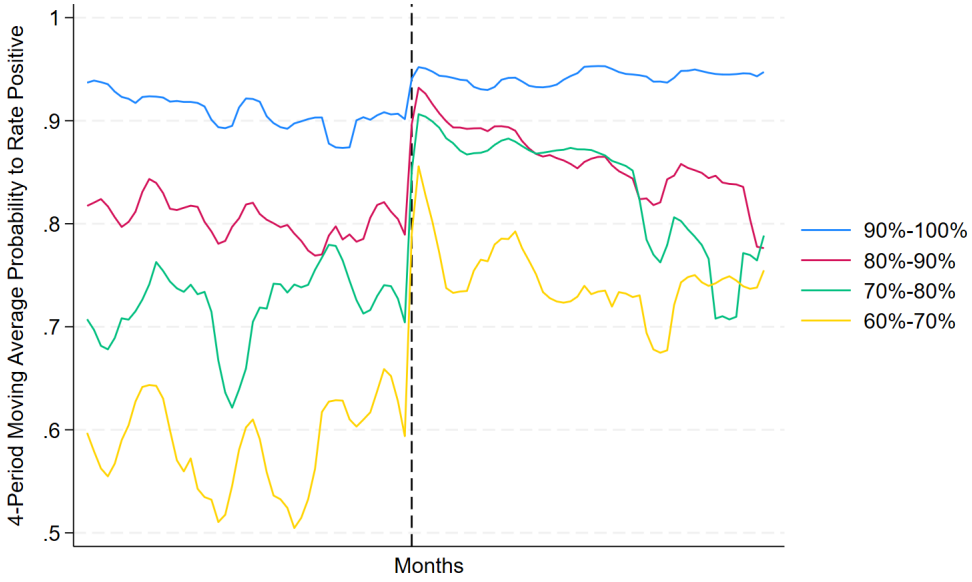
Separately, the full history of a game's price, sales, player numbers and discounts at the time of review since 2015 is scraped from *SteamSpy.com*[17] These will be useful controls; price and sale data will also be useful for running a simple structural model in Section 7. See Appendix B for further comments on this data.

3.5 Descriptive Statistics

Descriptive statistics and simple correlations for the variables are shown in Table B.1. As shown in Figure 1, the update resulted in a large change to the average positivity of a review. Figure 1 includes only reviewers present before and after the policy change to avoid impact from user selection. If there is indeed some effect of average rating on user rating, we should see this more clearly by splitting Figure 1 into different rating groups. Figure 2 shows the effect of the policy across 4 different rating categories. Clearly, there is a large effect of the policy, which differs by rating. The worst rated games see the largest increase, with smaller effects for the better rated games. This does not control for sales, game releases or trends in a game's life cycle and so it is still quite noisy. These results will be further examined in Section 5.

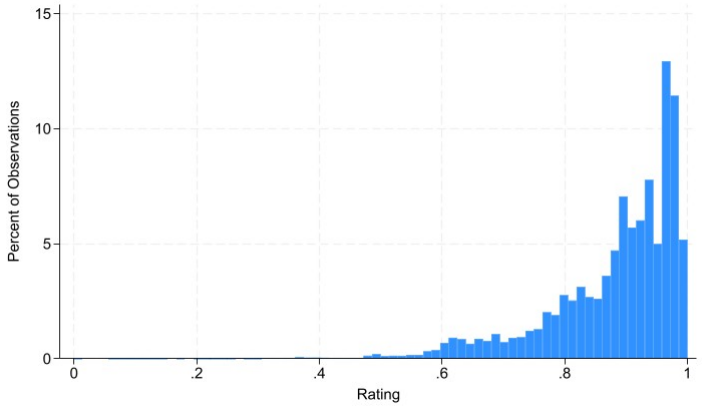
[17]See https://steamspy.com/ for the data and a wider discussion on owner data estimates.

Figure 2: Group Level Responses in Raw Data



Similar to Figure 1, the raw average positivity is displayed except across different rating bins. Data is aggregated monthly and a 4-period moving average is taken. The y-axis is the average probability to rate positive which simply takes the mean positivity of all user ratings in the bin for a given month. The x-axis is time in months. The dotted line marks the policy change. The graph shows that worse rated games do substantially better after the policy, with less change to the lowest rated games. This will be further explored in the later sections but is strong evidence for the main effect of interest.

Figure 3: Ratings Exposed to Users at Time of Review



Histogram of the average rating for a given user at time of given review. This is the distribution of average ratings displayed to users, across all games and time. The y-axis represents the percentage of all posted reviews that at time of writing were exposed to a rating within that bin.

Figure 3 shows how a game's average rating varies at time of review for each user review in the dataset. Most reviews are left when the average game rating is upwards of 60%. Similar to platforms like Uber, where 5 stars is the modal rating, games can reasonably be 100% positively rated. Therefore, games rated at 70% are not necessarily well rated, as this reflects that 30% of all users would not recommend the experience. I explore how this distribution of game ratings relates to social influence effects in section 6.

Basic correlations of characteristics within reviews can also be seen in Table B.2. Negative reviews tend to be longer, and be left by users who own more games.[18] The average rating correlates positively with user rating.

4 Empirical Analysis

To identify the effect of the average rating on user behaviour, I run the following regression:

$$y_{rijt} = \alpha_i + \mu_j + \tau_t + \beta \bar{y}_{jt} \times D_{jt} + \kappa D_{jt} + \lambda \bar{y}_{jt} + X'\gamma + \epsilon_{rijt} \tag{1}$$

In equation (1) y_{rijt} is the binary rating r left by individual i for game j at review cycle time t. Review cycle time is calculated as the number of weeks since a game's release. $\bar{y}_{jt}$ is the average of all previous user ratings for game j at review cycle time t. This is the information displayed to the reviewer before the policy change but not afterwards. $\bar{y}_{jt}$ will be different for each new reviewer leading to variation within game. D_{jt} is a dummy for treatment which is equal to 1 if game j at review cycle time t is reviewed using the new system and 0 otherwise. X' is a vector of controls containing price, player numbers, playtime at review, and whether the game has had a recent update. α_i are the individual fixed effects of a given user. μ_j are game fixed effects. τ_t are the time fixed effects for each week since a game's release. ϵ_{rijt} is a stochastic error term. All games are treated, though at different times in their life-cycle, and this specification compares the average probability to rate a game positive pre and post treatment by rating. Implicitly I assume that removing the exposure to average rating should only have an effect as a function of the average rating. Therefore, $\kappa + \beta \bar{y}_{jt}$ is the expected increase in probability to rate a game positively conditional on rating. I also consider alternate specifications that split rating into 4 groups ($60\% - 70\%$, $70\% - 80\%$, $80\% - 90\%$ and $90\% - 100\%$) or allow rating to vary up to the 4th power.

The main issue with attempting to identify the effect of the average rating on individual user ratings is the co-movement between three different types of effects, as described by Manski (1993). This co-movement is known as the *reflection problem*. The first type, *endogenous effects*, refer to whether seeing a different average rating changes an individual reviewer's own behaviour. These are the main effects of interest. The second type, *exogenous effects*, refer to whether the average characteristics of a reviewer's reference group change their behaviour. Finally, *correlated effects* refer to the effects driven by the similarity between group members. Similar people generate similar responses but this is not due to interaction between them.

Furthermore, even if games were randomly assigned or there was no correlation between players of a game, Angrist (2014) shows that in many cases the usual peer

[18] For further empirical observations, see Lin et al. (2019) which looks at these correlations within review.

effects estimator does not reflect any true relationship whether it exists or not. It simply captures the difference between 2SLS and OLS using group dummies as instruments for the given peer values, and these could differ for many reasons.

Here, the identification of the social influence effect leverages the exogenous policy change. Therefore, the effect identified is the change in user rating behavior before and after the policy as a function of average rating. $\bar{y}$ is not a proxy for general interdependencies but a real observed variable. It changes constantly over time, and a new user is exposed only to the rating of those who came before them. These previous users were not exposed to that user and so this follows Angrist (2014) recommendations for a variable that will have interpretable meaning. Nonetheless, it is crucial to point out that the coefficient on $D_{jt} \times \bar{y}_{jt}$ is not interpreted as a "change in social influence." Equation (1) aims to recover the *treatment effect* of the policy by *different game ratings*. This is similar to Bursztyn and Jensen (2015) who identify the effect of peer pressure on student performance by comparing ability quartiles before and after a policy which exposed them to their fellow student's performance.

The regression is run only on users present both before and after the policy change. Identification comes from the within-reviewer, within-game analysis and the random treatment time within a game's life cycle. Equation (1) compares the probability to rate positive for a predicted user, rating, game and week since release both before and after the policy change. Standard errors are clustered at the game level. I control for price and player numbers, as well as include time fixed effects for each week since a game's release to ensure trends within a game's life cycle do not drive the results.

Crucially, the policy change is an exogenous treatment that reduces the information seen by a user at time of review. If this policy changes how users review, then the average rating causally impacts how an individual user reviews.

In the next section, I discuss the results as well as alternative mechanisms and potential concerns, such as selection. In short, the evidence strongly suggests a genuine change in the opinions formed by reviewers. In Section 6 I discuss the potential mechanisms driving the results.

5 Results

5.1 Baseline Results

Table 1 shows the change in probability a review is positive after the policy change. As expected, the probability of a positive review is highly correlated with average rating due to the reflection problem. Column 3 presents the main results. A 10 percent[19] increase in average rating increases the probability a review is positive by 5.4% before the policy change, but only by 2.8% after. This suggests that reducing the exposure of reviewers to the average rating reduces the correlation of their reviews with this rating.[20] On its own, this fall in correlation tells us little and so as discussed in Section 4 we will focus on the treatment effect by average rating.

The probability to leave a positive review increases by 26.6% after the policy. Removing exposure to the average rating increases the overall positivity in the review

[19] In absolute terms.

[20] There could be concern that updates might be driving the results if games change more after the policy. However there are on average fewer updates to the games in my sample after the policy change compared to an equivalent window before. Therefore at worst the coefficients would be underestimated. In addition Table C.5 runs the baseline including a unique Game × Time fixed effect for each game-week pair to attempt to fully capture updates, and the results are unchanged.

Table 1: Baseline Results

	(1) Positivity	(2) Positivity	(3) Positivity	(4) Sentiment
Policy × Rating	-0.179**	-0.156**	-0.269***	-0.058***
	(0.072)	(0.061)	(0.063)	(0.019)
Rating	0.900***	0.843***	0.544***	0.102***
	(0.047)	(0.046)	(0.080)	(0.033)
Policy	0.153**	0.120**	0.266***	0.059***
	(0.067)	(0.057)	(0.056)	(0.017)
Observations	6,510,809	6,403,981	6,403,981	6,403,981
Games (Clusters)	249	249	249	249
Controls	Yes	Yes	Yes	Yes
Individual FE	No	Yes	Yes	Yes
Game FE	No	No	Yes	Yes
Time FE	No	No	Yes	Yes

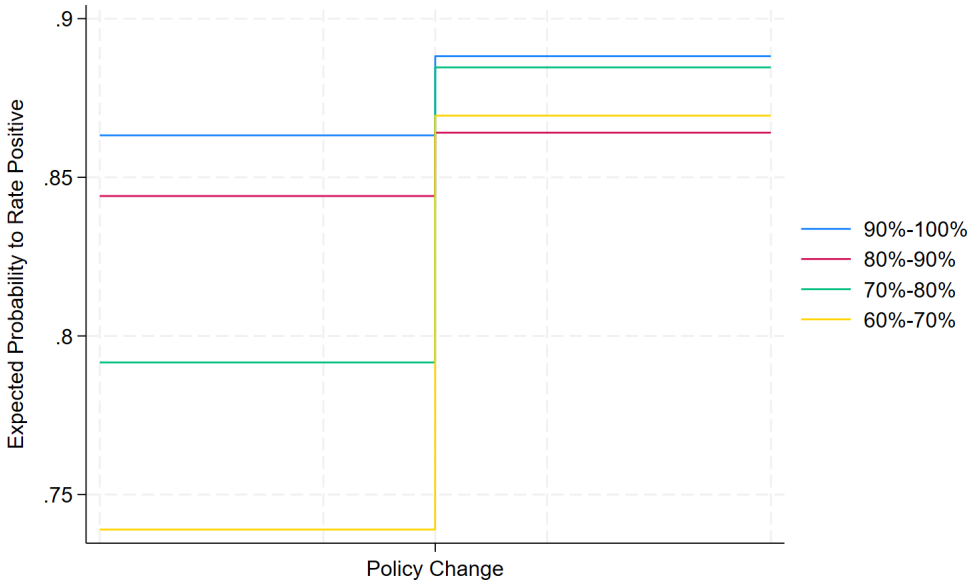
Results from the Baseline regression. Controls include price, player numbers, a dummy equal to 1 if any update has occured, review length and playtime at review. The sample is restricted only to users who left reviews before and after the policy change, which is 1,467,981 unique users leaving on average 5 reviews (on average 2.7 before and 2.3 after). See Figure C.1 for a frequency chart of total reviews left by those in the sample. See Table C.1 to see results are robust to excluding users above or below a certain number of reviews. There is a slight drop of observations when including fixed effects as while all users have left more than one review, some might have missing values for price if left before 2015. The outcome positivity refers to a binary variable which is 1 if the review was positive and 0 otherwise. Column 4 runs the same regression as column 3 but with sentiment as the dependent variable. This is an absolute measure between -1 and 1. Standard errors are clustered by game and in parentheses.

system. However, this is distributed unevenly. Worse rated games do much better after the policy change relative to better rated games. To illustrate, a 50% rated game has a 13.5% higher probability to be rated positively after the policy change. In contrast, a 100% rated game would see no change to the probability it is rated positively. The positive effect decreases to 0 as average rating increases. It seems that exposure to the average rating reinforces negative opinions across games.

Group Level Responses A straightforward way to see this asymmetry is to generate an indicator for each of the main rating groups in the sample: $60\% - 70\%$, $70\% - 80\%$, $80\% - 90\%$ and $90\% - 100\%$. These variables substitute $\bar{y_{jt}}$ in equation (1). The results are shown in Figure 4, which displays the probability that each group is rated positively. The lower the rating group, the greater the effect, with the $60\% - 70\%$ getting a 13% increase. $70\% - 80\%$ rated games receive a 10% increase, and $80\% - 90\%$ and $90\% - 100\%$ receive a 2% increase each. The result being so similar to the main specification even with such coarse aggregation implies that measurement errors in average rating are unlikely to affect the results. An alternate specification allowing the effect to be non-linear in average rating is shown in Figure C.2, and supports this mechanism. Furthermore, this effect is even visible in the raw data as previously shown in Figure 2. This strongly suggests that the average rating reinforces negativity in online review systems.

Alternate Rating Measures Results are robust to changing how the average rating is measured, as shown in Table C.2. Column 1 displays the results when considering reviews to be left when created rather than updated. Column 2 shows the results

Figure 4: Group-Level Results



A graph of the change in expected probability to rate positive by rating group. The rating at time of review is split into 4 dummies, from 60% and increasing by 10 to 100%. These are the coefficients from a regression identical to equation (1) but using four rating bins rather than a continuous rating. The dependent variable is again the binary rating of a user review. This compares the effect on a review for a rating in that bin before and after the policy. To see the change in probability a review is positive for a given group, see the jump after the policy. See Table C.3 for the companion regression and standard errors, and see Figure 2 for the effect in the raw data. Standard errors are clustered at the game level.

after dropping all updated reviews. In both cases the results are similar. In addition, results are robust to measuring average rating using *SteamSpy* data. Since *SteamSpy* only has data from mid 2015, the sample is restricted to games released from 2016. The effect is still present though smaller due to fewer low-rated games in this sample.

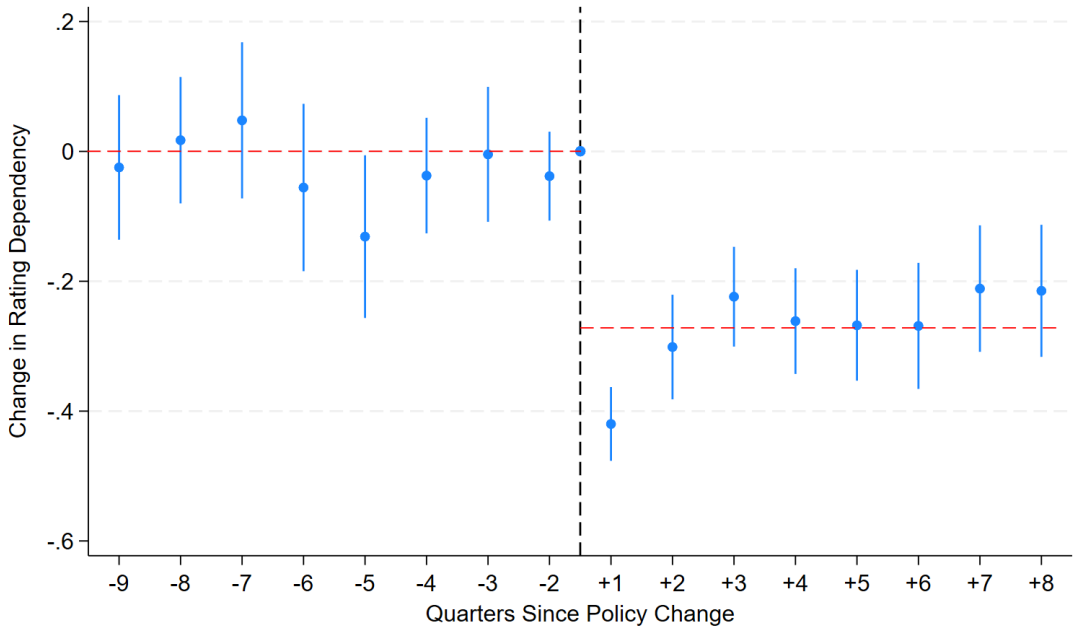
Sentiment Column 4 in Table 1 presents the results of estimating equation (1) using sentiment as the dependent variable. Sentiment varies from -1 to 1, and so to interpret the coefficient we compare the effect to the average absolute sentiment, which is 0.22. The coefficient on Policy reflects a 27% increase in positivity relative to the average before the policy change. The effect is again asymmetric: a game rated at 50% gets 13% more positively written reviews relative to the mean before the update, with no change for a 100% rated game. This exercise suggests social influence affects not just binary ratings but also the positivity of written text.

5.2 Selection Effects

Selection due to Change in Review Cost The main concern is that the results are driven by selection bias. Under the new system, the individual cost of leaving a review likely changes. This could in turn affect which types of reviews a user chooses to leave, making them more positive purely through selection. For this selection to be consistent with my results, reviewers would have to choose to post different types of reviews only for games rated poorly, and gradually change this selection rule in response to the average rating.

If this were the case, the effect would not persist. Average ratings would adjust

Figure 5: Variation Around the Policy Change



The quarterly change in rating dependency is plotted above, using the quarter before the policy change as a base month. The specification includes all users to allow for comparisons between quarters before and after the policy change. Sales, player numbers and their interactions are also controlled for in this specification. The Steam Summer Sale for the last 2 weeks of June 2019 has been specifically omitted due to the magnitude of its flash sales. All rating coefficients before the policy change are insignificant bar 5 quarters previous which is barely significant. There is a clear sustained drop before and after the policy. Confidence Intervals are at the 95% level and are displayed as error bars, clustered at the game level.

to become more positive over time and eventually average and user rating would correlate as they did before the update. If selection was driving these results we would expect a steady decay in the coefficient on rating after the policy. However, if social influence is present, the increase in positivity should persist. To address this, I run the main specification allowing the coefficients to vary quarterly.

These results are shown graphically in Figure 5. Each point represents the change in β for a given quarter relative to the last quarter before the policy change. The coefficients only become significant after the policy change. Indeed, the initially lower coefficient 1 period after the policy is likely picking up precisely this effect. Any selection after the policy would lead to a mechanical break in the rating dependency. However, it does not decay to pre-policy levels but stays consistently less predictive of review positivity after the policy change, suggesting the results are not driven purely by user selection.

Number of Reviews If results are driven by users selecting for more positive reviews on worse rated games, we should see no change in the daily number of negative reviews. If reviewers change positive reviews to negative reviews when exposed to the average rating, we would expect to see a decrease in the average number of negative reviews post policy and only on lower rated games. Table C.6 shows this occurs, with a 71% decrease in the daily number of negative reviews for the 60%-70% group relative to their pre-treatment mean. There is a 51% decrease in the daily number

of negative reviews for the 70%-80% group, and small increases for the highest rated games. This suggests a genuine reduction in negativity due to the policy. In addition, we want to ensure that selection in general is not greater on lower rated games, as this could drive the differential effect between the rating groups. If this is the case, then after the policy change we would observe a larger number of reviews specifically on lower rated games. Results are shown in column 3 of Table C.4. We find the opposite, and that there are proportionally more reviews posted on higher rated games.

Sentiment If the effect is driven by selection of a different type of review, we would expect the average review sentiment to change. Since the policy made it strictly easier to leave a review, we might expect that users post reviews for games they care less strongly about. This would be reflected in average sentiment moving closer to zero, that is a fall in the absolute value of sentiment. Column 1 of Table C.4 estimates equation (1) with the absolute value of sentiment as the dependent variable. The coefficient on rating is small and insignificant suggesting selection if present is not driving the results or again disproportionately affecting worse rated games. Since the effect is more positivity, column 2 looks at sentiment only for positive reviews and finds no significant effect. This suggests that the new positive reviews posted after the policy are not significantly different in their strength of opinion. Sentiment is not a perfect measure of opinion strength, but helps corroborate the presence of social influence.

5.3 Heterogeneity by Experience

To look at heterogeneity of results by user type, the regression is re-run interacting rating × policy with user's playtime at review[21] and lifetime number of reviews left. These are proxies of user experience within and across games. Of most interest is the user's playtime at review, as availability of the exact time spent with a product at time of review is a unique feature of the *Steam* platform.

As we can see from Table 2, experience reduces the susceptibility to social influence. An increase in the playtime at review for a given game or the lifetime number of reviews left both decrease the impact of the policy change. It makes sense that inexperienced users might be more likely to be swayed by the opinions of others compared to users who have left more reviews or played the game for longer. Less experienced users might supplement their lack of experience with the ratings of others. Reviewers try to be helpful to potential consumers, as it is purely a cost to the user to leave a review. Less experienced users who are unsure about the quality of a product may attempt to learn from the ratings of others when posting a review.

6 Mechanisms

I develop a simple theoretical model to help understand the results and the asymmetry of the effect. The first implication of my results is that they rule out the operation of a simple herding mechanism. Herding would imply that when exposed to the rating all reviewers would match with the average as best they could leading to many 0% and 100% rated games. Bayesian updating in order to learn the true quality of the

[21] Playtime at review is normalised by dividing by the mean playtime at review for a given game. This ensures that a low playtime at review reflects an objective lack of experience rather than a short game.

Table 2: Heterogeneous Effects

	(1) Positivity	(2) Positivity
Policy × Rating	-0.156**	-0.372***
	(0.061)	(0.072)
Policy × Rating × Log Playtime	0.040***	0.050***
	(0.014)	(0.014)
Policy × Rating × Log Reviews		0.087***
		(0.012)
Policy	0.153***	0.358***
	(0.054)	(0.063)
Policy × Log Playtime	-0.049***	-0.059***
	(0.012)	(0.012)
Policy × Log Reviews		-0.084***
		(0.011)
Observations	6,403,981	6,403,981
Games (Clusters)	249	249
Controls	Yes	Yes
Individual FE	Yes	Yes
Game FE	Yes	Yes
Time FE	Yes	Yes

The regression is re-run including a full set of interaction terms. The additional interactions with ratings are included in the controls but not displayed for ease, as the interest is on the total effect of the policy change. Column 1 is run only on playtime at review, though the results are similar including lifetime number of reviews as shown in column 2. To see the effect of interest, note how the Policy × Rating coefficient is negative, but the further interactions are positive. The same holds for the Policy coefficient but in reverse. As experience increases, this total effect of negativity compounding falls. Note that since playtime is normalised, log of playtime will be negative for anyone below the mean so the effect of the policy for someone with 1 game owned and a small amount of playtime will be much larger than the effect in Table 1. Standard errors are clustered by game and in parentheses.

game would suffer from a similar issue. Upon being exposed to the average rating reviewers would have no need for their own signal. in addition, both of these standard approaches do not explain the asymmetry of the effect.

My model below shows the results can be rationalized by a user who aims to leave a truthful review, but is uncertain about the precise quality of their gaming experience. Users are influenced by the average rating, but only in cases where their expectation of game quality is close to 0.

Empirical evidence shows that most games in my sample are highly rated. The asymmetry reflects how reviewers are substantially more surprised that a more positive signal comes from a "bad" game than a negative signal from a "good" game. Finally, I suggest an additional mechanism where reviewers are risk averse in their recommendations and would rather not recommend a game that they enjoy if it could lead to disutility for a future buyer. They prefer not recommending a game that many users like as if their review was pivotal the loss of utility is smaller.

6.1 Theoretical Framework

Consider a simplified version of the review system for illustrative purposes. A user i plays some game and decides to leave a review. For ease we assume that some users will review regardless. I assume that a game quality is observed as either good or bad from the average rating. More formally, $G \in \{Bad, Good\}$.[22] Their utility from leaving a review R_i is

$$U_i = -(R_i - q_i)^2 + k$$

Where $R_i \in \{-1, 1\}$ is a binary choice to recommend or not recommend, $q_i \in [-1, 1]$ is their idiosyncratic quality of the game which is drawn from some distribution $F_i(G)$ and k is some scalar > 0. Before they go to review they are exposed to the average rating of the game depending on whether the policy change is active. Intuitively, this user wants to leave a review that best describes their own personal experience with the product, and for now is indifferent to the "true" quality of the game. If q_i is known, then it is trivial to see that regardless of game quality reviewers recommend when idiosyncratic quality is greater than 0, and do not recommend when idiosyncratic quality is less than 0.

Unknown Quality Now, consider that the user does not know q_i exactly. Games are often substantive experiences with many parts. The exact level of quality that a game provides is unknown but rather reviewers are broadly aware of how much they liked something. One can imagine a user being able to confidently rate for example a movie or game as good or bad, but have more difficulty pinning down an exact rating between 0 and 100.

Rather than drawing some q_i from a distribution, they draw a uniform distribution which contains their true value of quality. $q_i \in S_i = Uniform(q_{i_{min}}, q_{i_{max}})$. For further simplification, assume that $S_i = Uniform(\mu_i - c, \mu_i + c)$ where now μ_i is drawn from some distribution $F_i(G)$. Note, μ_i is not their true signal here as receiving the full distribution would allow them to immediately back it out. Rather, the true quality

[22] This is done for ease of exposition. The average rating could be a noisy signal of a game's average quality, but given the large number of reviews it is likely quite clear what the average quality may be. The results only rely on a lower average rating being correlated with users having a worse experience with the game.

is somewhere uniformly in the interval, and we would expect that lower centered intervals are more likely for worse rated games. Utility from leaving a review is

$$U_i = -(R_i - E_[q_i|\Omega])^2 + k$$

Where $E[q_i|\Omega]$ is their expected idiosyncratic quality conditional on their information set. Note that for many users, their optimal choice does of R_i will not change. If the minimum and the maximum value of q_i are both positive, then they can still confidently recommend the experience. They would also get no benefit from learning more about the exact value of q_i. However, some users will be unsure whether the experience is worth recommending or not and this leads us to our first prediction of the model.

Outcome 1 If $q_{i_{min}} < 0 < q_{i_{max}}$ then exposure to the average rating can cause a reviewer to change their review. When a user observes if a game is *Good* or *Bad*, they can update on this information via Bayes' rule. Crucially, we assume that $P(G = Bad|q_i)$ is strictly decreasing in q_i and $P(G = Good|q_i)$ is strictly increasing in q_i.[23] Then, since q_i is uniform:

$$E[q_i] = \int_{q_{i_{min}}}^{q_{i_{max}}} qdq > \int_{q_{i_{min}}}^{q_{i_{max}}} q\frac{P(G = Bad|q_i)}{P(Bad)}dq = E[q_i|G = Bad]$$

From $P(G = Bad|q_i)$ strictly decreasing and with the reverse holding when seeing a game is good from $P(G = Good|q_i)$ strictly increasing. Intuitively if reviewers are unsure about their own quality, and they believe a games quality being good or bad is less likely when their own quality is lower of higher respectively, then the average rating allows them to update their expected quality of the game. A wider interval and strong learning from sharper likelihood functions $P(G|q_i)$ will result in more users changing their review in the presence of the average rating. A reviewer who is unsure about the exact quality of their experience learns whether their experience is better or worse from the average rating. Some users will change their review in the presence of the average rating, but crucially many will not. This also helps understand the heterogeneity discussed in Section 5.3, as we would expect experienced users to be more sure about the quality of their own experience.

6.2 Asymmetry

Asymmetry of Priors One empirical fact about the *Steam* platform is that most games can achieve a high average rating. This is shown in Figure C.3. As a result, $P(Game = Bad)$ is substantially smaller than $P(Game = Good)$. This immediately leads to the second outcome of the model.

Outcome 2 If most games are *Good*, users will update much stronger upon seeing a game is *Bad* than seeing a game is *Good*. This can be seen rewriting

$$E[q_i] = pE[q_i|G = Bad] + (1 - p)E[q_i|G = Good]$$

Where $p = P(Game = Bad)$. The absolute difference between the conditional and unconditional expectation by game quality is given by

[23] Intuitively this reflects that you would expect a game to be worse if your signal on its quality is lower.

$$|E[q_i] - E[q_i|g]| = \begin{cases} p(E[q_i|Good] - E[q_i|Bad]) & \text{if } G = Good \\ (1-p)(E[q_i|Good] - E[q_i|Bad]) & \text{if } G = Bad \end{cases}$$

And so as p gets smaller, the asymmetry grows larger. Intuitively, if I think a game is slightly good and am unsure, I believe with high confidence I simply had a worse time with a good game. When I see a low average rating and realize the game is bad, I update strongly that I had a worse experience with the product. On the contrary, if I felt I had a slightly bad experience with a game, I know that most games are good across *Steam* and believe I likely had a bad time with a good product. When that is confirmed, my posterior mean moves by substantially less as this is hardly a surprise. The distribution of game ratings is heavily skewed and with a high average of almost 90% so this seems likely to explain the extreme asymmetry. Only 20% of games in the sample are rated 60%-80%. Assuming these are the "bad" games, plugging 0.2 into the formula above leads to a 4 times larger negative effect than positive without any further calibrations or assumptions on other functions in the model.

Asymmetry of Recommendations Consider an additional factor in the reviewers' utility, now given by

$$U_i = \lambda(-(R_i - q_i)^2) + (1-\lambda)E[C_j(q_j)|G, R_i] + k$$

Where everything is as before, $\lambda \in [0, 1]$, but with the addition of $E[C_j(q_j)|G, R_i]$ which is the expected value of user j's consumption utility from the game at their own given quality assuming that user i's review is pivotal. Not only is the average rating observable, but on *Steam* individual reviews can be read. They can be rated as helpful and often the text is well written so in theory can be useful. Intuitively, if a reviewer is trying to be helpful they get utility from their review helping someone find a good game, and avoid a bad game.[24] A user wants to give a truthful review, but now they also have some preferences in trying to help match an incoming user with a better game. If a potential buyer does not purchase the game, they get an outside option of 0.

In this simple model a reviewer and potential purchaser immediately knows the exact overall quality of a game. This does not mean that an incoming user will dislike it, just that it is more likely they will have a bad experience. However, if the average harm of playing a bad game is worse than the average benefit of playing a good game, this leads to the next channel of asymmetry.

Outcome 3 Assume $C_j()$ is strictly concave. Also, assume $C_j(q_j)$ has the same sign as q_j. Then reviewers are less likely to recommend a bad game compared to not recommending a good game. Not recommending a good game leads to a prospective buyer getting 0 for sure, compare to some average quality. Not recommending a bad game leads to the prospective buyer getting 0 again, but recommending a bad game leads to a prospective buyer getting a much worse average loss from the concavity of $C_j()$.

Intuitively, if I am to recommend a game to someone I want to be honest about my experiences. However, I also know that if they buy a bad game because of me, they are

[24]The pivotal assumption is not necessary and could be some more complicated function of how likely their review is to be read but the core idea is brought across regardless.

going to have a low utility as they have wasted both their money and their time. In contrast, if they buy a good game because of me they get to pay for a good experience which is closer to 0 in absolute terms. This then helps explain the asymmetry of the effect. The model suggests that before the policy potentially positive reviews of a game were changed to negative. Reviewers did not want to risk recommending a product that they enjoyed if there was a higher likelihood of a bad experience. After the policy, they are no longer exposed to the average rating and so leave these positive reviews. This is seen in the increase in positivity and the reduction in the number of negative reviews on poorly rated games. This channel seems reasonable, though compared to asymmetry on priors it is harder to understand how pervasive this could be.

7 Consequences of Social Influence on Sales

7.1 Outcome on Market Shares

Another question is the extent to which this social influence might affect sales. In general, work such as Anderson and Magruder (2012) suggests that ratings do have causal impacts on purchase decisions. However, with *SteamSpy* approximations for owners of games, I can directly analyse the monetary impact of the policy change.

I run a simple multinomial logit BLP model in the style of the structural literature.[25] This allows aggregate market data to be used to make claims about an individual's product choice. *SteamSpy* data is available only from mid 2015, so only games released in 2016 or later are included. This then becomes a discrete choice problem across the 76 games in my sample and the outside option of some other game not included in my sample. Descriptive statistics for this data set are shown in Table B.3 and additional comments are made in appendix B. Following the literature[26] we assume the utility will take the form:

$$u_{ijt} = -\alpha p_{jt} + \beta x_j + \gamma r_{jt} + \xi_{jt} + \epsilon_{ijt} \tag{2}$$

Where p_{jt} are prices for product j at time t, x_j is the vector of observed product characteristics, r_{jt} is the rating of product j at time t, ξ_j is the unobserved product characteristics on utility, and ϵ_{ijt} is a Type-1 extreme value mean 0 error term. Product characteristics are taken from user submitted tags on games and then normalised. These tags describe the game, and are visible on the store page e.g single-player, adventure, action.

Care must be taken to ensure both rating and price are not endogenous. Price is surprisingly straightforward in this context. Games sold on *Steam* are digital copies and so have 0 marginal cost, and are in infinite supply. Prices do not have to adjust in response to demand to account for supply effects. This may not be true in the long run. For example, cheaper games might get updated more frequently[27] or games being kept at a low price might encourage the entry of other products. However, in the short run this is purely a demand shock, and common on *Steam* are many flash sales where products have their prices reduced for a short time. Often these sales are

[25] The framework will follow Nevo (2000). This model will be estimated by multinomial logit for simplicity as the goal is to identify the broad magnitude of the effect.

[26] See McFadden (1981)

[27] This is an important supply side consideration in Teng (2022) who looks at a structural model of the app store

Table 3: Multinomial Logit Results

	(1) Mkt. Share	(2) Mkt. Share	(3) Mkt. Share	(4) Mkt. Share
Rating	10.62***	10.88***	7.57***	11.19***
	(0.562)	(0.630)	(0.547)	(0.600)
Price	-0.041***	-0.041***	-0.058***	-0.021***
	(0.003)	(0.002)	(0.004)	(0.001)
Product Fixed Effects	Yes	Yes	No	Yes
BLP Instruments	Yes	No	Yes	No
Price Instruments	Yes	Yes	Yes	No

The results from a simple multinomial logit model on the aggregate market shares. The dependant variable is the total market share for product j in week w. The outside option is defined with an approximation on weekly aggregate Steam sales (around 40 million). Estimation occurs via linear IV-GMM, where price and rating are both endogenous except for column 4. Instruments for rating include Game × Policy, for price Game × Discount and for both the sum of characteristics of other products when BLP instruments are used. To ensure measurement issues are kept to a minimum, the product fixed effects or vector of characteristics are always interacted with a dummy which is 1 post April 2018. As discussed in appendix B, SteamSpy's owner estimation algorithm changed substantially at this time. When product fixed effects are not included, the instruments are instead Characteristics × Policy and Characteristics × Discount. Standard errors are in parenthesis

during certain times of the year such as Lunar New Year or Christmas, but games often go on sales semi-regularly at all times of the year. These short flash sales are a good instrument for game price.

Rating is also endogenous as choosing a given product means you value it, and are more likely to rate it highly. However, the policy change by game is a good instrument as it triggered an exogenous shift in rating. Furthermore, due to the nature of the update, the policy should not have affected sales in any other way. In addition, the classic BLP instruments of sums of product characteristics in the market will be used for both price and rating, though this will not be crucial for the results.

Results Table 3 shows the results of the simple multinomial logit estimation on the aggregate market data. Game specific dummies are used as suggested by Nevo (2000) and so product characteristics are absorbed.[28] We can see that rating does indeed have a positive effect on the market share for a given product. This is to be expected, but since the exogenous variation here is coming from the policy change itself this shows that policies reducing social influence do affect sales.

In terms of magnitudes, a 1% increase in rating[29] increases the market share relative to the outside option by 0.1%. A better understanding of the magnitude is to see that a 1% increase in rating has the same effect on market shares as a 2.5 dollar decrease in price. This reflects the importance to both sellers and consumers in reducing bias in online reviews. Consumers prefer to choose a product with a higher rating, and so a product where negativity has compounded due to social influence bias will receive fewer sales than in the counterfactual without this bias. This simple model supports the previous literature that suggests ratings affect purchase decisions.[30] This

[28]They can be recovered if we assume $E[\xi_{jt}|x] = 0$ but we are not particularly interested in these specific coefficients here.

[29]In absolute terms e.g 60% to 61% or 78% to 79%

[30]Simulating the policy is not necessarily revealing here. Ex-ante consumers would prefer the

suggests that platforms would benefit from aiming to reduce social influence bias, as if negativity compounds sales will be artificially depressed.

Efficiency It seems reasonable to suggest that removing exposure to the average rating increased efficiency of the game market on *Steam*. Games received higher ratings and sold better. In addition, the increase in the ratings is stable as discussed. We might expect if this policy led to worse matching between buyers and games that the increase would be eroded by then a sharp decrease in ratings. If users were elicited to buy games from artificially high ratings they should then perhaps leave negative reviews. This does not happen, and removing exposure to the average rating increases sales and user ratings consistently. This suggests it was good not only for the ex ante choice utility but broad efficiency in delivering consumers games. Consumers now buy games they would not have, and these consumers do not seem to have a negative experience.

8 Conclusion

Social influence is present in online review systems. I show that when a reviewer is exposed to the average rating of a product, their review is biased in line with that average rating. More importantly, this is driven by an asymmetrical effect. Negativity compounds in online reviews where the average rating is shown. Removing exposure to the average rating closes the gap between worse-rated and well-rated games. This effect is visible linearly, non-linearly and by rating group, and is robust to a battery of alternative measures. The effect is even visible in the raw data.

Many sites across the web show users the average rating before they leave a review. This paper suggests that this has serious outcomes for both product ratings and sales. An artificial gap between worse rated and better rated products is formed which not only makes the average rating a worse measure of quality, but likely reduces the sales of these products too. A simple theoretical model of reviewer behavior rationalizes this result as reviewers attempting to pin down the quality of their own experiences updating on the average rating. Asymmetries can arise from the relative rarity of negativity which leads to lower rated products being rated as substantially lower than they would be without exposure to the average rating.

I also show that this effect varies by user experience. Steam provides a unique opportunity to see how long a user spends with a product before they review. Social influence is driven by less experienced users on the platform. This has implications for products which might select for these types of users and could be more at risk to social influence.

Finally, using a simple structural model I show that the variation of the average rating due to the policy change has substantial impact on market share and consumer choice. Crucially, this bias affects sales. Games rated better after the policy sold better after the policy. The magnitudes suggest there is value in reducing this bias across platforms, especially if the asymmetrical compounding of negativity is common across review sites in general. Ultimately, platforms should take great care in how they display their average rating to incoming reviewers. It results in a social influence bias that affects both user rating and purchase decisions.

highest rating possible, even if this resulted in a worse matching probability. It is clear that varying ratings has monetary value, but a structural model in this context cannot make claims about whether rating increases are good for matching users to games.

References

[1] Michael Anderson and Jeremy Magruder. "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database". In: *The Economic Journal* 122.563 (2012), pp. 957–989.

[2] Angrist. "The perils of peer effects". In: *Labour Economics* 30 (2014). Special Section articles on "What determined the dynamics of labour economics research in the past 25 years? edited by Joop Hartog and and European Association of Labour Economists 25th Annual Conference, Turin, Italy, 19-21 September 2013 Edited by Michele Pellizzari, pp. 98–108. ISSN: 0927-5371. DOI: https://doi.org/10.1016/j.labeco.2014.05.008. URL: https://www.sciencedirect.com/science/article/pii/S0927537114000712.

[3] Joshua D Angrist and Kevin Lang. "Does school integration generate peer effects? Evidence from Boston's Metco Program". In: *American Economic Review* 94.5 (2004), pp. 1613–1634.

[4] Sinan Aral and Christos Nicolaides. "Exercise contagion in a global social network". In: *Nature Communications* 8 (Apr. 2017), p. 14753. DOI: 10.1038/ncomms14753.

[5] Leonardo Bursztyn and Robert Jensen. " How Does Peer Pressure Affect Educational Investments? *". In: *The Quarterly Journal of Economics* 130.3 (May 2015), pp. 1329–1367. ISSN: 0033-5533. DOI: 10.1093/qje/qjv021. eprint: https://academic.oup.com/qje/article-pdf/130/3/1329/17098054/qjv021.pdf. URL: https://doi.org/10.1093/qje/qjv021.

[6] Jiafeng Chen and Jonathan Roth. "Logs with zeros? Some problems and solutions". In: *The Quarterly Journal of Economics* (2023), qjad054.

[7] Scott A Imberman, Adriana D Kugler, and Bruce Sacerdote. "Katrina's children: Evidence on the structure of peer effects from hurricane evacuees". In: *American Economic Review* 102.5 (2012), pp. 2048–2082.

[8] Young-Jin Lee, Kartik Hosanagar, and Yong Tan. "Do I follow my friends or the crowd? Information cascades in online movie ratings". In: *Management Science* 61.9 (2015), pp. 2241–2258.

[9] Dayi Lin et al. "An empirical study of game reviews on the Steam platform". In: *Empirical Software Engineering* 24 (2019), pp. 170–207.

[10] Charles F Manski. "Identification of endogenous social effects: The reflection problem". In: *The review of economic studies* 60.3 (1993), pp. 531–542.

[11] Dina Mayzlin, Yaniv Dover, and Judith Chevalier. "Promotional reviews: An empirical investigation of online review manipulation". In: *American Economic Review* 104.8 (2014), pp. 2421–2455.

[12] Daniel McFadden. "Econometric models of probabilistic choice". In: *Structural analysis of discrete data with econometric applications* 198272 (1981).

[13] Lev Muchnik, Sinan Aral, and Sean J Taylor. "Social influence bias: A randomized experiment". In: *Science* 341.6146 (2013), pp. 647–651.

[14] Aviv Nevo. "A practitioner's guide to estimation of random-coefficients logit models of demand". In: *Journal of economics & management strategy* 9.4 (2000), pp. 513–548.

[15] David A Reinstein and Christopher M Snyder. "The influence of expert reviews on consumer demand for experience goods: A case study of movie critics". In: *The journal of industrial economics* 53.1 (2005), pp. 27–51.

[16] Bruce Sacerdote. "Experimental and Quasi-Experimental Analysis of Peer Effects: Two Steps Forward?" In: *Annual Review of Economics* 6.Volume 6, 2014 (2014), pp. 253–272. ISSN: 1941-1391. DOI: https://doi.org/10.1146/annurev-economics-071813-104217. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-economics-071813-104217.

[17] Sarang Sunder, Kihyun Hannah Kim, and Eric A Yorkston. "What drives herding behavior in online ratings? The role of rater experience, product portfolio, and diverging opinions". In: *Journal of Marketing* 83.6 (2019), pp. 93–112.

[18] Xuan Teng. "Self-Preferencing, Quality Provision, and Welfare in Mobile Application Markets". In: (2022).

Appendix A: Steam Background

Figure A.1: Example Review



An example of a simple review of the game 'Grand Theft Auto 5."

Steam Games

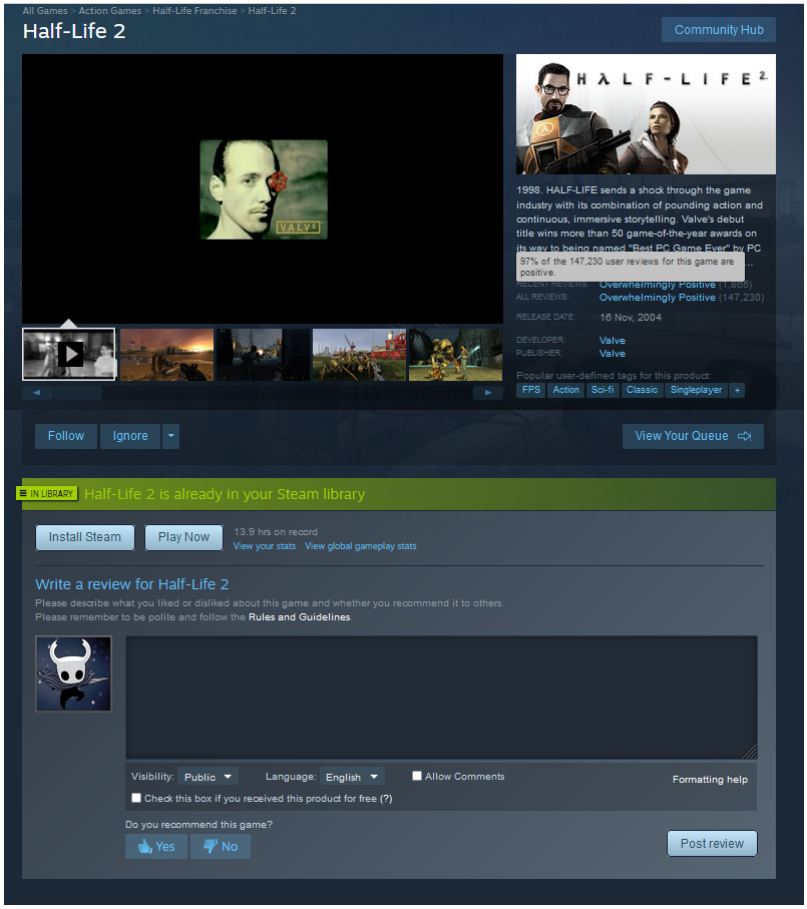
Games on Steam are often substantial experiences in line with games sold for consoles such as Xbox or Playstation. Most games available on those consoles since 2005 have been made available on the Steam store, except for some that are exclusive to a specific console. Some games are small experiences, cheap and short to finish, while some cost hundreds of millions to make and take many hours to complete. Some games have a multiplayer component and so are theoretically limitless in how much they can be played (some reviewers of online games have thousands of hours of playtime). One important aspect to keep in mind is that games, unlike movies or TV shows, can be updated or changed after the release. As a result, ratings can reasonably change over time and there can be large variation even within games. For example, often games can be released with glitches or 'bugs" that result in a poor user experience, but if the developers fix these issues then the game could be a much better experience. Steam has the largest selection of games across all platforms, including consoles, with around 30,000 products to choose from and so certainly represents well a wide variety of game types. Within my sample I have multi-player games, single-player games both short and long and even games made by a single person.

Valve's Goals

Valve has detailed in many blog posts [31] that their goal is to make a good review system that matches consumers to games well. Of course no company would admit to attempting to manipulate reviews to encourage purchase, but it does not seem that Valve is engaging in this type of behaviour. Valve has a voluntary refund policy whereby consumers who have played the game for less than 2 hours can get a guaranteed refund, so do not benefit selling games individuals do not enjoy. In addition, while Valve does not compete that closely with other online game distributors, they

[31] See https://steamcommunity.com/groups/steamworks/announcements/detail/1697229969000435735 for the policy change, user stats and https://store.steampowered.com/oldnews/24155 for a general discussion of making an accurate rating.

Figure A.2: Review System Before the Policy Change



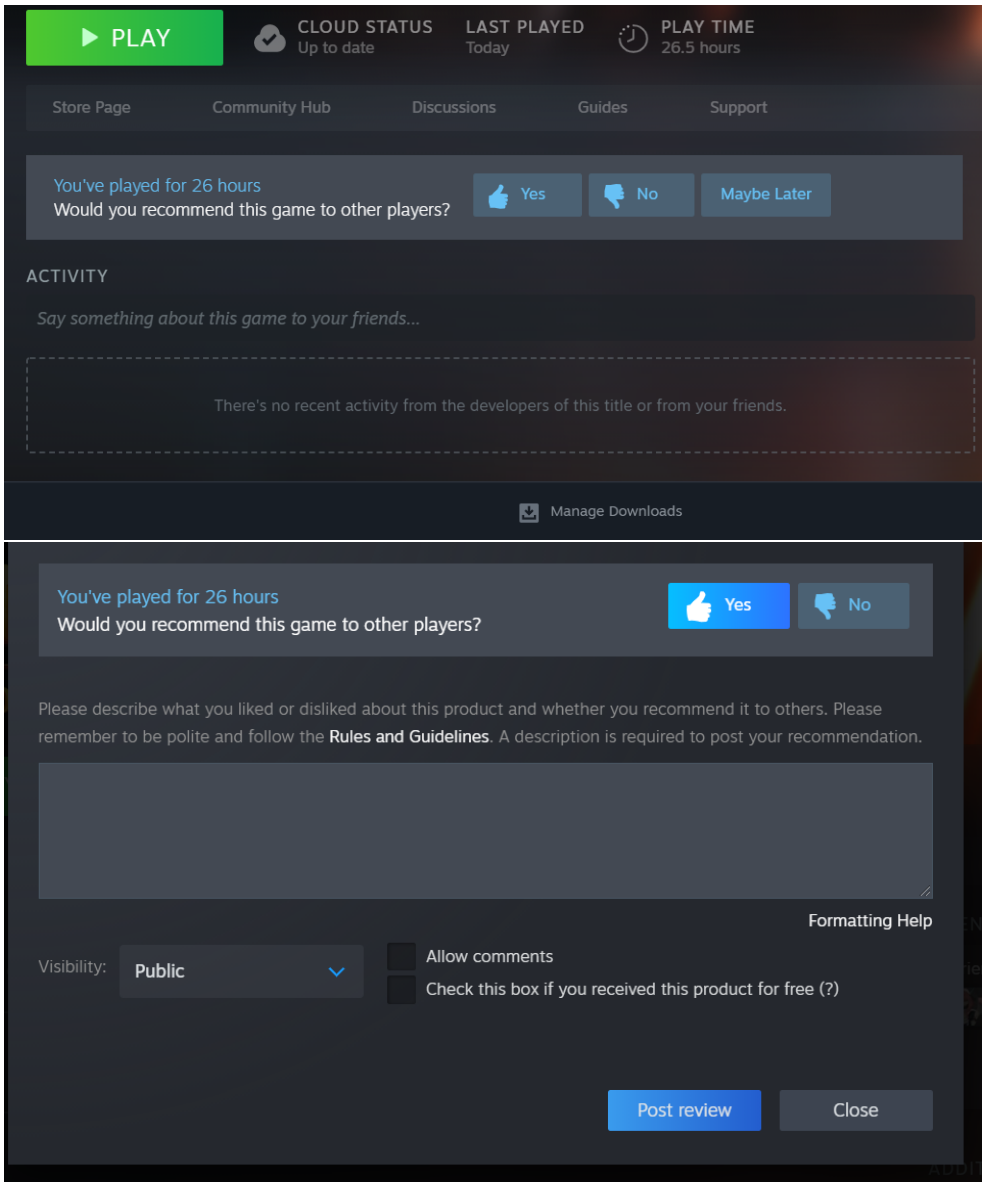
An example of the initial review posting system. The rating of others is immediately obvious when going to leave a review, and the exact rating is clearly visible above the review box. This channel was the default way of reviewing before the Library Update in 2019. It was possible to review via the bottom of the library page but this is not mentioned in any of Steam's blog posts and all instructions to review prior to the policy direct users to the Store page. This is still a usable channel for reviewers, but directly more obtuse than the new method. Even in the worst case where large groups of users still use the old method which seems unlikely, this would depress the coefficients in magnitude.

do compete with the outside option of illegally downloading games, or "pirating" games. Gabe Newell, the founder of Valve has gone on record detailing this:

'The easiest way to stop piracy is not by putting antipiracy technology to work. It's by giving those people a service that's better than what they're receiving from the pirates."

Broadly speaking, Steam as a platform has a goal that aligns with consumers in building a review system that increases the probability a game purchased is actually enjoyed by the consumer. This ensures users trust the review system and continue to purchase games on Steam. In addition fake reviews are unlikely to be an issue here. This paper focuses on the top 250 reviewed games at the time of the policy, and since an account must have purchased and played the game to be able to leave a review, it is not particularly feasible when the average game has 100,000s of reviews.

Figure A.3: Review System After the Policy Change



An example of the review posting system after the policy. The user is prompted upon going to play the game or exiting the game to leave a review directly. When they click on the yes or no buttons, the pop up box below appears and allows them to directly leave a review within the window. Note that the box is very similar to before, the only difference being the average rating not shown to the reviewer.

Appendix B: Data

Table B.1: Descriptive Statistics for Main Data

	Mean	Standard Deviation	Min	Max
Positive	0.88	0.32	0	1
Number of Games Owned	95.70	301.00	0	31,071
Number of Reviews Left	14.67	53.53	1	10,606
Total Playtime	31,601.10	72,630.80	0	6,049,533
Playtime in Last Two Weeks	107.16	695.32	0	127,474
Playtime at Review (Minutes)	15,740.75	42,609.80	1	4,776,595
Review Length (In Characters)	169.05	2,400.33	1	9,638,999
Policy (At Created)	0.55	0.50	0	1
Policy (At Updated)	0.58	0.49	0	1
Rating (At Created)	0.88	0.10	0	1
Rating Last 30 Days (At Created)	0.88	0.12	0	1
Rating (At Updated)	0.89	0.10	0	1
Rating Last 30 Days (At Updated)	0.89	0.12	0	1
Price (USD)	18.13	15.52	0	59.99
Percent of Total Owners Playing	12.97	16.46	0	100
Player Numbers	151,603.10	324,669.90	0	3,080,310
Updated	0.06	0.24	0	1
Rating (From SteamSpy)	0.87	0.11	0	1
Game Released after 2015	0.24	0.42	0	1
Log Playtime at Review	7.92	2.01	0	15.38
Log Number of Reviews	1.73	1.29	0	9.27
N	20,485,646			

Table B.2: Simple Correlations

	Positive	Length	Games	Playtime	Price	Play Num	Rating
Positive	1.00						
Length	-0.12	1.00					
Games	-0.04	0.13	1.00				
Playtime	0.00	0.00	-0.04	1.00			
Price	-0.12	0.12	0.05	-0.00	1.00		
Play Num	-0.13	-0.03	-0.05	0.06	0.01	1.00	
Rating	0.32	-0.07	-0.00	-0.04	-0.36	-0.28	1.00

Variables used as stored in the data set with associated descriptive statistics.

Table B.3: Descriptive Statistics for Weekly Market Data

	Mean	Standard Deviation	Min	Max
Price	27.60	18.60	0.00	59.99
Discount	8.03	18.82	0.00	92.00
Rating	0.84	0.13	0.31	0.99
SteamSpy Rating	0.84	0.13	0.29	1.00
Sales	26,459.66	114,338.61	0.00	3,937,275.00
Owners	3,930,472.45	7,064,129.74	1,000.00	85,161,578.75
Action	0.13	0.09	0.00	0.43
Adventure	0.09	0.08	0.00	0.29
Atmospheric	0.08	0.08	0.00	0.27
Classic	0.01	0.02	0.00	0.13
Co-op	0.06	0.08	0.00	0.30
Competitive	0.00	0.02	0.00	0.17
Dystopian	0.00	0.01	0.00	0.07
Fps	0.04	0.06	0.00	0.32
First Person	0.05	0.07	0.00	0.29
Free to Play	0.03	0.07	0.00	0.41
MMORPG	0.00	0.02	0.00	0.15
MOBA	0.00	0.02	0.00	0.16
Massively Multiplayer	0.02	0.05	0.00	0.27
Moddable	0.02	0.04	0.00	0.16
Multiplayer	0.14	0.11	0.00	0.41
Puzzle	0.01	0.04	0.00	0.22
Shooter	0.04	0.06	0.00	0.23
Singleplayer	0.14	0.10	0.00	0.40
Story Rich	0.06	0.09	0.00	0.34
Strategy	0.09	0.15	0.00	0.56
Team Based	0.01	0.03	0.00	0.11
E-Sports	0.00	0.01	0.00	0.07
N	16,098			

Descriptive Statistics for the data used in Section 6. Variables are a weekly average, and all taken from *SteamSpy* except for Rating which is calculated as described in the main text by aggregating reviews, then further aggregated by week. Using *SteamSpy* rating instead has virtually no difference on the regression. The horizontal line break shows the tags used in generating the vector of characteristics for each product j. Their sums are used to generate BLP instruments. The results are robust to excluding these tags altogether and using product fixed effects.

Review Data

Direct characteristics of the review are displayed both to users and when scraping, such as whether it was voted up, the number of games owned, number of reviews and playtime at review. Steam users can also rate reviews as funny or helpful, and this data is also collected though has little bearing on the specific channel of interest. The number of games owned and the number of reviews left is a measure that is not fixed at time of review, so it reflects the number of games owned or reviews left at the time of scraping. This is oddly how the data is displayed for Steam users too: when reading a review from someone with 100 reviews left, it could be that this was their first review and they have left 99 others since then. Some users can private the number of games they own if their profile is set to private. They are displayed as 0s in the data-set, but everything else is forced to be shown if a review is left.

SteamSpy

SteamSpy is a website that has been scraping the *Steam* Store and user profiles daily since mid 2015. The website as a result has the largest archive of price data, player data and additional data available for analysis. Price and player number data is fairly straightforward as described in the main text: *SteamSpy* is needed to save the daily history but price and player numbers are directly shown on *Steam* and so these are completely accurate for the given day.

Sales Data

Unfortunately, *Steam* does not reveal sales data for games. However, *SteamSpy* uses an algorithm that trawls through user profiles, sees which games users own, and then estimates the total number of owners based off this information. It is very accurate anecdotally according to the site creator and some developers who can compare estimates to their actual sales.

There are however issues. If games have free weekends, where a game is free to download for a weekend so users can try it out, then estimates for owners will be skewed. After the weekend users must then buy the game in order to keep playing. However, *Steam* incorrectly considers them owners for much longer than the weekend lasts. This can result in a reduction in owners spread out over oddly long periods of time, and therefore daily sales calculated as the difference in owners can be negative. If this occurs, the negative number is replaced with a missing value [32] This is far from ideal, but this is the most accurate data available to anyone outside of *Steam*. In addition, rather than look at daily markets, the structural model in Section 6 aggregates weekly to smooth over potential measurement issues.

Another issue discussed on *SteamSpy* is that in 2018 user profiles became private rather than public by default. The algorithm is still anecdotally accurate but the owner numbers jump for all games around that time. All regressions in Section 6 add a unique game dummy at this time to control for the mean shift after this date.

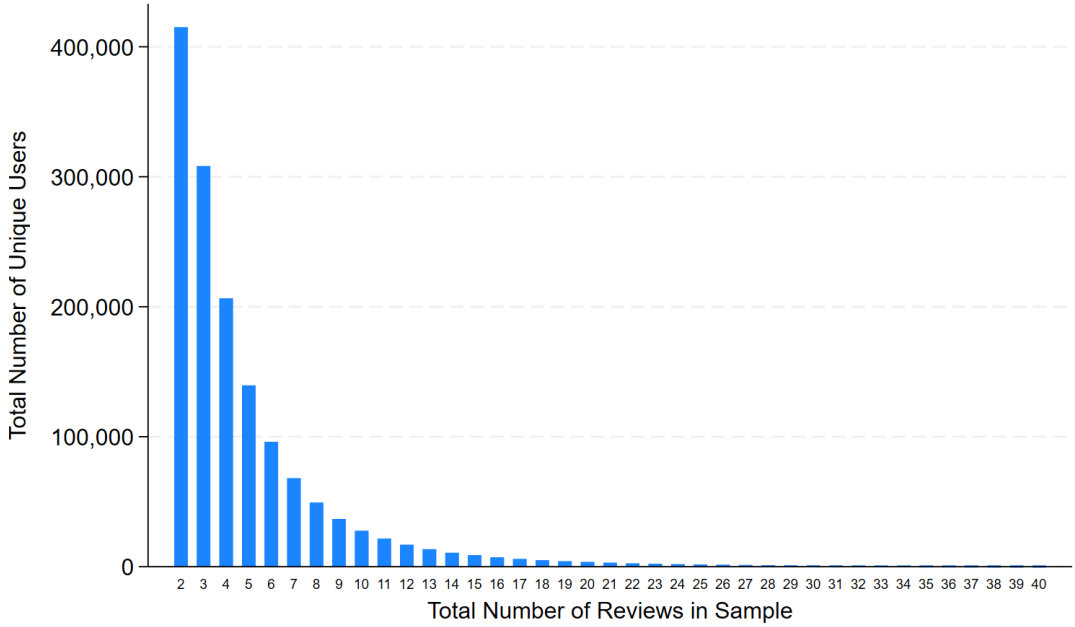
Game Tags

While not necessary for the results, the usual set of BLP instruments and game characteristics are constructed using user submitted tags for games. These include things like a game being single player, multi-player, sci-fi, survival, horror etc. This data is also available on SteamSpy. The tags change over time, but the proportions stay broadly similar and the tags for a game at May 2019 are used as measures of a game's characteristics. To account for bigger games being tagged more, the tags are normalised by the total number of tags left to get a proportion between 0 and 1. For example, game j tagged 50 times as single-player and 50 times as survival would have characteristics of 0.5 each. There are a huge list of possible tags, so the tags are pruned into more main categories and are shown in table B.3.

[32] see Chen and Roth (2023) for a discussion as to why replacing with a small value is not a good idea.

Appendix C: Additional Checks

Figure C.1: Frequency of Lifetime Reviews



This frequency chart displays the number of unique users with a given lifetime number of reviews left. This is exclusively taken from those with at least one review before the policy and one after, hence the minimum value at 2. The graph is cut off at 40 for ease of display but there are users with more than 100 reviews in the sample. While the modal number is 2, there are many in the sample leaving more than 2 reviews. We can see the average user leaves 5 reviews, and there are many 100,000s of users leaving more than 4 reviews. Note that the identification strategy demeans individuals from their fixed effects, and so provided demeaning can occur (i.e there is more than 1 observation) there is no advantage for an individual having 6 reviews rather than 5 except in how well the user can be demeaned. To confirm, the results displayed in Table C.1, are from re-running the baseline regression dropping users with a certain number of reviews and are very similar.

Table C.1: Robustness to Number of Reviews

	(1) Positivity	(2) Positivity	(3) Positivity
Policy	0.251***	0.250***	0.317***
	(0.054)	(0.048)	(0.054)
Rating	0.547***	0.538***	0.539***
	(0.078)	(0.076)	(0.078)
Policy × Rating	-0.252***	-0.239***	-0.331***
	(0.061)	(0.053)	(0.061)
Reviews in Sample	> 2	> 4	< 4
Observations	4,781,368	2,355,491	3,882,213
Games (Clusters)	249	249	249
Controls	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes
Game FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

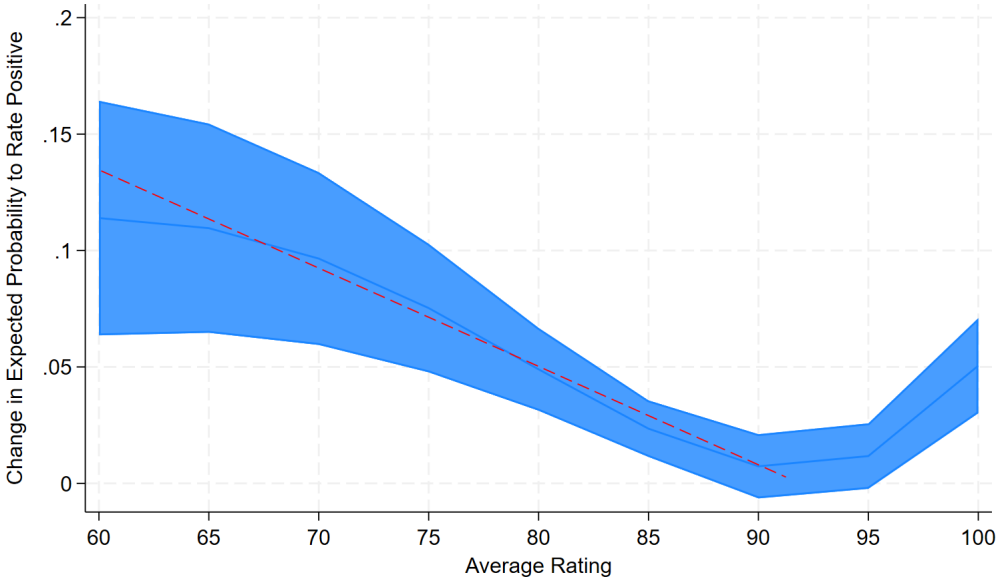
The baseline regression is re-run dropping users based on how many reviews they have left in the sample. This is to ensure results are not being driven by users with few reviews or many reviews. Column 1 drops all users with only 2 reviews. Column 2 drops all users with 4 or fewer reviews. Column 3 drops all users with more than 4 reviews. In all cases, the results stay very similar. The higher coefficients for column 3 compared to column 2 reflects the heterogeneity by experience. Reviews in sample correlate with lifetime number of reviews and hence we would expect the effect to be larger for these users. Standard errors are clustered at the game level and in parentheses.

Table C.2: Alternate Rating Measures

	(1) Positivity	(2) Positivity	(3) Positivity
Policy	0.220***	0.318***	0.141*
	(0.060)	(0.054)	(0.074)
Rating	0.360***	0.550***	0.421***
	(0.071)	(0.074)	(0.101)
Policy × Rating	-0.261***	-0.328***	-0.178**
	(0.067)	(0.062)	(0.084)
Measure	Created	No Updated	SteamSpy
Observations	6,434,055	5,185,856	1,478,341
Games (Clusters)	249	249	76
Controls	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes
Game FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

The baseline regression is re-run with alternate measures for rating. "Created" reflects that the date a review is considered to be left is when it was created rather than updated. "No updated" simply drops all updated reviews from the data set (around 2 million total, and 1 million from those leaving multiple reviews). "SteamSpy" uses the rating calculated from SteamSpy for games released in 2016 and after. This is available for 76 games in my sample, and is only daily. This lower variation results in the greater attenuation. Standard errors are clustered at the game level and in parentheses.

Figure C.2: Non Linear Responses



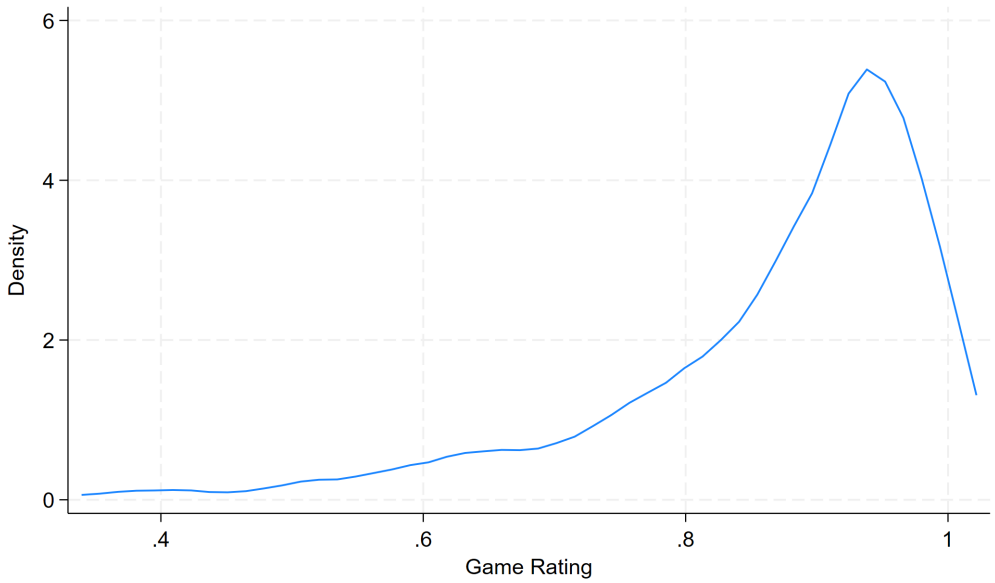
The main specification is re-run with both rating and rating interaction terms up to and including the fourth power. The above graph shows the marginal effect of the policy change at different levels of average rating. For example, for a game with an average rating of 60%, there would be an approximate 12% increase in the probability said game was rated positively. There would be no significant change to games with an average rating of 90%. The effect is calculated for each 5% jump in rating. Even allowing the dependency to take a more general functional form, the change is roughly linear with respect to rating up to about 90%, after which there is no clear effect for the highest rated games and the uptrend is likely an artefact of the functional form. Standard errors are clustered at the game level, and confidence intervals are at 95% coverage and calculated via the *margins* command in *Stata*.

Table C.3: Group Level Results

	(1) Positivity
60% - 70%	-0.124***
	(0.021)
70% - 80%	-0.072***
	(0.012)
80% - 90%	-0.019**
	(0.008)
Constant (90% - 100%)	0.884***
	(0.006)
Policy × 60% - 70%	0.106***
	(0.037)
Policy × 70% - 80%	0.068***
	(0.017)
Policy × 80% - 90%	-0.005
	(0.009)
Policy (90% - 100%)	0.025***
	(0.007)
Observations	6,203,303
Games (Clusters)	249
Controls	Yes
Individual FE	Yes
Game FE	Yes
Time FE	Yes

The companion regression to Figure 3. The main specification is repeated with rating bins rather than a continuous rating. The base category is the 90%-100% group before the policy change. All other aspects are identical to the main specification, other than reviews left at an average rating less than 60% are dropped due to too few observations. To see the total effect as displayed on the graph, add coefficients if the dummy variable is 1 to the constant. For example, the expected probability to rate positive for the 70-80 group after the policy is 0.884 - 0.072 + 0.068 + 0.025. Standard errors are clustered at the game level and in parentheses.

Figure C.3: Distribution of Average Ratings by Game



The kernel density estimate of games in my sample. This is run across games not reviews in order to capture what the average consumer might expect a game's rating could be. Of course, reviewers and users are not necessarily exposed to every game and indeed if they are forgetful of a given games rating it is unlikely they would be aware of the full distribution. However, the extent to which poorly rated games are so rare means that an individual is going to be much more surprised when exposed to an average rating between 60% and 70% when they themselves have a slightly positive experience.

Table C.4: Mechanism Checks

	(1) Abs(Sentiment)	(2) Sentiment if Positive	(3) Daily Total of Reviews
Policy	0.008	-0.004	-1.763
	(0.005)	(0.012)	(3.733)
Rating	0.028***	-0.018	-10.358
	(0.008)	(0.023)	(9.818)
Policy × Rating	-0.010	0.006	11.103**
	(0.006)	(0.013)	(4.444)
Observations	6,403,981	5,331,172	629,643
Games (Clusters)	249	249	-
Controls	Yes	Yes	Yes
Individual FE	Yes	Yes	-
Game FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

Column 1 re-runs the baseline regression with the absolute value of sentiment as the dependent variable. Concern is that if selection is higher on lower rated games then this could be driving the results. The absolute value of sentiment is a proxy for how extreme a user's opinion is and so if it falls for the lower rated games it suggests selection is occurring more on these groups. However, there is no change to the absolute value of sentiment across rating categories and the coefficient is insignificant. While not a perfect measure it helps corroborate the social influence mechanism. Column 2 is the same with unmodified sentiment on only positive reviews. Since the results are driven by an increase in positivity we would expect that if selection occurs we would see a decrease in sentiment as reviewers select for less extreme reviews. All coefficients of interest are insignificant and the coefficient is minuscule compared to the pre-treatment average of around 0.2. Column 3 re-runs the baseline regression with daily total of reviews left by those in the sample. All variables are aggregated by day hence the smaller number of observations and no individual fixed effects. This is done to confirm that selection is not the greatest on the lowest rated games. We see that there are more reviews posted, but increasing with rating suggesting there is not greater selection on lower rated games that drives results. The average total number of reviews pre-treatment is actually slightly higher for games rated 60%-70% compared to 90%-100%, and all are around 9-10 and so these results reflect a proportional increase.

Table C.5: Game × Time Fixed Effects

	(1) Positivity
Policy	0.243***
	(0.058)
Rating	0.095***
	(0.105)
Policy × Rating	-0.240***
	(0.052)
Observations	6,402,502
Games (Clusters)	249
Controls	Yes
Individual FE	Yes
Game × Time FE	Yes

The regression is re-run including a unique fixed effect for each week-game pair. The variation here comes from looking at the rating correlation within a given game at a given week before the policy compared to after the policy. If there is a social influence effect, we would expect that even controlling for the average positivity in a given week we would still see more correlation when users are exposed to the average rating. Results are robust to these fixed effects. Even if games get updated or have good or bad weeks, it seems that this is not driving the results. Standard errors are in parentheses and are clustered at the game level.

Table C.6: Number of Negative Reviews

	(1) Daily Total of Negative Reviews
60% - 70%	2.314***
	(0.109)
70% - 80%	1.759***
	(0.084)
80% - 90%	0.751***
	(0.057)
Constant (90%-100%)	0.945***
	(0.036)
Policy × 60% - 70%	-3.136***
	(0.120)
Policy × 70% - 80%	-1.415***
	(0.097)
Policy × 80% - 90%	0.132**
	(0.059)
Policy (90%-100%)	0.139***
	(0.053)
Observations	692,686
Game FE	Yes
Time FE	Yes

Similar to table C.4, the data is collapsed daily and the regression is run on the number of negative reviews left by rating category. If there is social influence, there should be a fall in the number of negative reviews across the worst rated games. If there is selection, there should be no change to the number of negative reviews or even an increase. The results show that there is indeed a substantive fall in the number of negative reviews after the policy. The pre treatment mean for the 60%-70% group is 4.44 and for the 70%-80% group 2.84. The results reflect a 71% fall and a 51% fall in negative reviews respectively. Standard errors are robust and in parentheses.