

The Never Ending Book

Marchenko, Maria; Sonnabend, Hendrik

DOI:

[10.57938/8cc8e44c-c21f-4789-9354-fe6321c504e4](https://doi.org/10.57938/8cc8e44c-c21f-4789-9354-fe6321c504e4)

Published: 01/05/2022

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Marchenko, M., & Sonnabend, H. (2022). *The Never Ending Book: The role of external stimuli and peer feedback in user-generated content production*. WU Vienna University of Economics and Business. Department of Economics Working Paper Series No. 320 <https://doi.org/10.57938/8cc8e44c-c21f-4789-9354-fe6321c504e4>

Department of Economics
Working Paper No. 320

The Never Ending Book: The role of external stimuli and peer feedback in user-generated content production

Maria Marchenko
Hendrik Sonnabend

May 2022



The Never Ending Book: The role of external stimuli and peer feedback in user-generated content production*

Maria Marchenko[†]

Hendrik Sonnabend[‡]

27th April 2022

This paper studies the determinants of the voluntary provision of user-generated (online) content. Using data from the largest fanfiction website, we find that writers respond differently to new original material: conditional on text length, writing times increase for the average writer and even more for the elite of prolific writer. We explain this finding with quality concerns. In addition, we find supportive evidence that community feedback encourages first-time contributors to continue publishing. For more established writers, we find that community feedback has a rather dampening effect on text lengths and writing times. Overall, these effects are more pronounced for high-quality community feedback (‘reviews’) compared to low-quality community feedback (‘following’, ‘favoriting’).

Keywords: fanfiction; user-generated content; online public goods; voluntary contribution

JEL classification: H41, C31, D01, Z11

**Acknowledgments:* We are grateful for the valuable comments provided by Sophie Einwächter, Karol Borowiecki, Samuel Cameron, Anne-Sophie Radermecker, Nicholas Martin Ford, Francesco Angelini, and participants at the ACEI meeting in 2021 as well as an online-seminar organised by Christiane Hellmanzik. In addition, we thank Mathias Tirtasana, Julian Herwig, and Houda Ben Said for their assistance in preparing the data.

[†] maria.marchenko@wu.ac.at, Vienna University of Economics and Business, Department of Economics, Welthandelsplatz 1, 1020 Vienna, Austria.

[‡] hendrik.sonnabend@fernuni-hagen.de, University of Hagen, Department of Economics and Business Administration, Universitätsstraße 11, 58097 Hagen, Germany.

1. Introduction

Digital technology, i.e., the representation of information in bits, has reduced the costs of storage, computation, and transmission of data (Goldfarb and Tucker, 2019). It has given rise to new (economic) phenomena like the sharing economy, pricing in the face of zero replication costs, and user-generated content (UGC). The latter has the obvious free rider problem; therefore, a growing body of literature tries to understand why people change their role from a consumer to a creator of public good and what motivates users to contribute content for free.

Previous research has highlighted several different drivers behind user-generated content production. Among them are social motivation like group size (Zhang and Zhu, 2011), social ties / ‘friends’ (Shriver et al., 2013), social norms (Chen et al., 2010; Burtch et al., 2018); past contributions (Aaltonen and Seiler, 2016); financial incentives (Sun and Zhu, 2013; Cabral and Li, 2015; Khern-am nuai et al., 2018; Burtch et al., 2018; Wang et al., 2021); non-monetary awards (Gallus, 2017; DeVaro et al., 2018; Burtch et al., 2021) and performance feedback (Huang et al., 2019); image-related or ‘glory-based’ utility derived from contributing (e.g., Toubia and Stephen, 2013; Goes et al., 2016; Xu et al., 2019).

Yet, stimuli coming from *the subject itself* have received less attention so far. Probably, the reason is that most of the aforementioned studies use data from popular websites related to informational content like the open online encyclopedia *Wikipedia*, platforms that follow a question-and-answer format, and rating sites. Hence, it is difficult to identify clear exogenous shocks from the subject itself that would allow for a causal interpretation of results.

This paper aims to address this issue by using the data from the world’s largest fanfiction website *Fanfiction.net*. Fanfiction is a literature by mostly amateur writers who use the narrative of movies, games, and books to develop their own stories. In the case of books, most of the fanfiction published on this platform is based on book series. This allows us to explore reactions to new content as quasi-exogenous shocks: new material can be anticipated, and there may be speculations (which itself can spur content production), but in order to work with the new material, fan writers must study it.

Like in other social network structures where users are allowed to contribute openly, a small group of individuals produces considerably more output than others. Those prolific writers on the fanfiction platform under consideration receive special attention in terms of text ‘followers’ and ‘favoring,’ among others. We argue that strong user expectations towards these writers shape a competitive environment. Prior studies that try to understand how competition affects creative work suggest that a moderate level of competition can foster innovation, whereas a high level of competition reduces it (Gross, 2020). Using data from a commercial Chinese writing platform, (Wu and Zhu, 2022) find that writers respond to more competitive pressure by increasing their output. Regarding innovation (proxied by user reviews), increased competition can have a positive impact or no impact at all depending on the compensation scheme. In our setting, when exposed to new material, top writers may face a trade-off between getting more attention by being ahead of the herd and meeting quality expectations. Results from an event study approach and Cox regressions indicate that no clear effect on the number of contributions can be found, but that the average fan writer needs more time to finish a text when exposed to new material. Indeed, this effect is more pronounced for popular and heavy writers suggesting that user attention (expressed in expectations, loyalty, and reputation) can work as quality control for UGC.

We also study the effects of different kinds of community feedback (‘favoring,’ ‘followers,’ and reviews) on writers. Many online platforms use feedback as a way to stimulate content contribution (Huang et al., 2019). For example, the badges at *Stack Overflow*, motivating statistics at *Academia.edu* or *ResearchGate*, ‘karma points’ at *Reddit*. The latter can be used by the users for the community rather than platform feedback, specifically by sending the awards to fellow users for useful information or interesting content. In the *Reddit*-based field experiment, Burtch et al. (2021) find that positive peer feedback in the form of awards positively influences both the frequency and the persistence of contributions. Furthermore, Zhang and Zhu (2011) show the importance of the audience for the individual-level contributions due to the social benefits that the bigger audience can create. ‘Favoriting’ and ‘following’ the work (or even the writer) are the two ways

of showing interest in the work of a fan writer, and the reviews suggest an even higher level of involvement. All types of feedback on the fanfiction platforms are predominantly positive, as will be discussed in more detail in the next section. Implemented by the platform’s community guidelines, the constructive way consumers and producers of creative work interact with each other makes *FanFiction.Net* different from other UGC platforms. We find that with more feedback, the first-time contributors are more likely to continue publishing their works on the platform, which is in line with findings in Burtch et al. (2021). Additionally, we detect longer writing times between two texts and mostly shorter texts for the authors with ‘better’ feedback. This result may suggest more efforts and internal quality control in writing new texts. The effects vary with authors’ characteristics though; for example, the effect of increasing writing time diminishes with the authors’ experience. Moreover, our findings on the nexus between high-quality feedback and text lengths support the idea of an optimal text length expressed by user reviews.

The rest of the paper is organized as follows. In Section 2, we introduce the fanfiction platform and provide background information on the topic. Section 3 presents the data set and descriptive statistics. Sections 4 and 5 present the empirical strategies and estimation results for community effects on the writing probabilities, times and lengths, and effects from the original work and new content. In Section 6, we summarize the results and discuss their implications for online communities with UGC production.

2. Institutional background

Fanfiction (also spelled as fan fiction, or abbreviated as fan fic, or fanfic) is a literature by mostly amateur writers. It is based on some original source like books, movies, TV shows, or even a life of a particular celebrity. Fanfiction narratives usually develop an alternative fable or new take on the original story or talk about secondary characters. Most fanfiction writers are themselves fans of the original story or a character; therefore, they have a special attachment to the story they develop. Fanfiction writing used to be considered a very marginalized activity for a narrow group of people (Thomas, 2011). However, it has become more visible and known to a broader public in recent decades,

especially with the growth of the online-platforms. The early examples of fandoms go back to Sherlock Holmes stories and Jane Austen books (Jamison, 2013). Fans have been fantasizing about alternative universes or the continuation of their favorite stories for quite some time already.

The topic of the original narrative may create a close isolated community within the fanfiction community, and fans are no longer just the consumers of the storyline, but also the creators and motivators (Thomas, 2011). Even though the publications are anonymous, and the authors do not always reveal information about themselves, the users quite often feel a connection not only to the narrative of the fanfic, but to the authors as well. The distribution of the readers' attention to the authors is quite uneven, with a relatively small share of the superstar writers who get more feedback on their work in the form of reviews, favoriting or following the updates of the text or the authors.

On the online fanfiction platforms, it is implied that the feedback to the new pieces is mostly positive (Yin et al., 2017), unlike on other social media. Of course, there is still some criticism; however, most of it is either constructive feedback on particular character development or an impatience about the update or new chapters. Hence, the authors receive a lot of encouragement and community support to keep on generating free content, and some readers might even transform into writers (Thomas, 2011). At the same time, the top writers might feel more pressure from the audience than the average writers to deliver an interesting story. In the cases of the serial original (series of books, movies, TV shows), the new content may call for new fanfiction writing. The writers would want to respond to the new material; however, the response may be different from the top writers and the average writers. Top writers, often needing to meet the expectation of the audience, are likely to face the trade-off between writing something fast and creating a more thought-through story. This trade-off can also be viewed as the two drivers of the content: competition from the other authors, which should motivate to create the content faster, and the feedback from the audience. The latter, in the case of top authors, may put pressure in terms of the quality of the writing, but at the same time, allow the authors more time since the audience will be willing to wait longer for their works. Thomas (2011)

provides an example of the popular writer Carol, who stopped writing in 2007, but the comments from her followers kept appearing on her page even three years later inquiring for new material.

As all the other examples of the UGC, fanfiction is rarely bringing profit to its creators and being published outside the community. Fanfiction, as Coppa (2017, p.8) notes, is “networked creative work produced within and for a community of fans.” There exists a number of papers highlighting the importance of networks in the creative process. For example, Borowiecki (2013) looks at the relationship between productivity and potential connections inside geographical clusters in classical music, and Vedres (2017) analyzes the joint participation in jam sessions on jazz musicians’ creative success. An article by Fraiberger et al. (2018) shows the importance of connections in the art using co-exhibition networks that affect the artists’ careers’ success. The length and trajectory of the career depend on the artist’s location in the network. And Marchenko (2020) shows the presence of positive peer effects on the secondary art market, with one of the possible mechanisms being the co-development of the artistic excellence of the artists of certain groups.

Fanfiction is not written and distributed for financial gain; it is “made for free, but not ‘for nothing’” (p.14), as the stories are considered gifts to a community (p.16) of like-minded peers. However, there are several successful examples of fanfiction writers, the most prominent being E.L. James. Her *Fifty Shades* series was first published as the *Twilight* fanfiction. *New York Times* bestselling author Cassandra Claire used to be a fanfiction writer as well. Laurie Penny, in her *Wired* piece, tells that ‘surprising number’ of film and TV writers, editors, journalists, and novelists used to write fanfiction (Penny, 2019). Some ideas in TV series (for example, *Torchwood* or *Smallville*) are thought to appear from the ideas in fanfiction, making the fanfiction community more visible and important for the broader culture than it is often thought.

Fanfiction is a unique example of User-Generated Content. *First of all*, it is very much connected and built upon the original narrative. New content in the series might affect the creation process and change the writing behavior of the creators.

Secondly, the fanfiction community provides strong positive feedback and a feeling of

belonging (Penny, 2019), which can be a strong motivator for creating new products (see an example of *Reddit* awards in Burtch et al., 2021).

Such feedback may even motivate readers to become writers. An increasing number of writers can create additional competition, and as a result, the behavior of the creators on the fanfiction platforms will combine both their reaction to the positive feedback and to competition. These reactions might be different for top writers and average writers, resulting in different writing strategies.

Lastly, the fanfiction community can be seen as a space for creative people to open up and find their own voices. In this sense, fanfiction can be seen as a starting point for creative production.

Fanfiction.net Established in 1998, *Fanfiction.net* is a multi-fandom online archive for fanfiction. The largest of its kind, it hosts millions of stories (mostly drafted in the English language) by over ten million registered users. The scope of the platform is comparable with *Wikipedia*, with more registered users - around 43 million, but only around 124000 active editors. The platform allows users to follow a certain author and stories, often published in a series of chapters or smaller bits and pieces. Registered users can leave reviews or simplified positive feedback by ‘favoriting’ or ‘following’ stories. It is a convention that reviews also bear mostly positive feedback, and it is the strongest signal, as it suggests the highest involvement of the readers, because it requires more actions than just hitting the ‘like’ button. Figure A.1 provides a small sample of reviews for a work with a high number of reviews. Out of more than 5,000 reviews, there are only some slightly negative reviews that do not agree with the characters or complain about the long waiting time. For the less popular texts, the reviews are quite encouraging, even when it is a few of them (see Figure A.2 for a further example).

Hence, the authors should get more motivation with receiving more reviews. Both ‘favoriting’ and ‘following’ are sending only positive feedback. Adding a text or an author to favorites signals the ‘liking’ of their work(s). The ‘following’ is stronger feedback as by following a text or an author, the user signs up to receive updates on the work or all the author’s works. Authors and readers can also use an implemented private messaging

function and search the community for a ‘beta-reader’ who will supervise and/or comment on their work regularly.

3. Data and descriptive statistics

Our data come from three different sources. In a first step, all available entries in the category ‘Books’ between 06/1999 and 12/2017 were collected from the website *Fanfiction.net*. In sum, these are 482,838 observations on fictional texts created by 203,233 different authors and build on 2,342 source texts.

The data include formal information about the writer’s ‘fanfiction age’ (i.e., the time elapsed since registration), the text (like language, category (i.e., age recommendation), genre, status (finished/unfinished), number of words, date of upload), and the online community response (number of reviews, number of followers, number of users who marked the text as a ‘favorite’)¹.

Table 1: Summary statistics for fanfiction

Variable	Obs.	Mean	Std. Dev.	Min	Max
contributions per month ^a	452,546	3,246.764	1230.060	1	5,477
contributions per original	2,342	206.165	2401.889	1	80,827
contributions per author	203,233	2.376	5.273	1	461
words per text	482,838	8,736.842	25,544.760	0	225,0144
reviews per text	482,838	26.022	120.493	0	16,113
favoriting per text	482,838	24.627	135.000	0	23,847
followers per text	482,838	22.027	141.886	0	23,424

Notes: ^a For 30,292 observations, the date of publication was missing.

Table 1 presents summary statistics. First, on average, the website publishes 3,247 contributions per month. Since 1999, we can observe a geometric growth followed by a period of consolidation at a high level, see Figure 1. To a certain extent, the drop in uploads around the year 2014 can be explained by the emergence of *Amazon’s* fanfiction platform *Kindle Worlds* (Lipton, 2014). Even though the initial number of fandoms joining the

¹The numbers for reviews, followers and favorites are only available as the final output rather than the dynamic process. Unfortunately, we cannot observe how the feedback density was changing over time: whether it was mostly obtained immediately after publication, or constantly in time, or after the new information from the original material, etc.

platform was rather small (only 24, see Contrera, 2014), Lipton (2014) suggests the emergence of the *Amazon* platform will “change the landscape for fanfiction writers”. One of the possible reasons is the change of the copyright agreement, with *Amazon* getting the copyright permission to the selected fandom. Additionally, there were some community concerns about the introduction of the SOPA² act in 2014, after it was postponed in 2012, which could have caused some decrease in the activity³. However, like Yin et al. (2017), we do not know about any legal changes at the time. Note that Yin et al. (2017) also find hints for seasonality with more writing activities in the (northern hemisphere) summer months and smaller peaks in December.

Second, contributions are highly concentrated on a few original works. Figure 2 shows that the top 3 titles are *Harry Potter*, *Percy Jackson and the Olympians*, and *The Hunger Games*. Apart from the *Phantom of the Opera*, all titles were released as a sequence of books. Third, the three types of feedback (reviews, favoriting, following) have a similar (but not identical) distribution.⁴ Yet, the correlation between the two types of low-effort feedback (favoriting and following) is much higher ($\rho_{FAV,FOL} = 0.944$, $\rho_{REV,FOL} = 0.690$, and $\rho_{REV,FAV} = 0.733$). Finally, the number of words per contribution as well as the output per author vary widely (see Figures A.3 and A.4 in Appendix A). While the majority of authors restrict to a limited number of uploads, there are also prolific writers with hundreds of publications. This pattern is common in social network structures where users are allowed to contribute openly.⁵ Accordingly, the majority of authors with more than one publication target only one source text (60.12 %), and 29.01% create text on two topics. On average, the time span between publications is 127 days (median: 27 days).

In a second step, data on the originals works was collected from the *Wikipedia* Encyclopedia and the social cataloging website *Goodreads*. As we are also interested in the fan writer’s response to new material, we identified 33 book series within the top 300 most

²Stop Online Piracy Act, the bill first introduced in 2011 to strengthen the power of US law enforcement to fight online copyright infringement.

³See, for example, a forum discussion here: <https://www.fanfiction.net/topic/2872/108508804/1/>

⁴We only have the data on the ‘favoriting’ and ‘following’ of the particular texts. The ‘favoriting’ and ‘following’ of the authors is not available.

⁵Fitting a Pareto distribution to the number of uploads per author (in ascending order) by using the Stata routine *paretofit* developed by Jenkins and van Kearn (2015) gives us an value of $\alpha = 2.172$.

Figure 1: Contributions per month.





Figure 2: Contributions per fandom (top 7 titles).



popular topics in our sample that had at least one new release (book or movie) in the observation period. This reduced sample includes 282,010 observations.

4. Community effects

We begin with analyzing stimuli coming from the community. Prior literature such as Toubia and Stephen (2013) and Huang et al. (2019) has shown that peer recognition can encourage UGC production. Furthermore, there is evidence that feedback affects users differently. Burtch et al. (2021), for instance, show that feedback is more important for new members since this group is less connected to the community and feels more uncertainty.

As mentioned in Section 2, there are three types of feedback which are available to users: reviews, (text) ‘following’, and (text) ‘favoriting’. We expect reviews to exceed the others in terms of quality and hence expect this type to have the highest impact on UGC production.

4.1. Just the debut?

We first examine whether feedback affects the propensity to stay in the community and to contribute more than one text. Therefore, the sample was restricted to debuts (193,163 observations altogether). To control for variations between source texts and categories, we use a set of text-related fixed effects. We also control for the general time trend in the popularity of fanfiction and the platform by including year fixed effects. Formally, we use the linear fixed effect model

$$single_{i,j,t} = \alpha_0 + \alpha_1 \ln Feedback_{i,j,t} + \xi_i + \theta_t + \varepsilon_{i,j,t} , \quad (1)$$

where $single_{i,j,t}$ is a binary dependent variable which equals 1 if text i published in month t is the only publication of author j , and zero otherwise. $Feedback$ is a placeholder for the log number of reviews, followers, and ‘favoriting’ attributed to text i . ξ_i is a set of control variables related to the article including the genre, the language, the text status (finished/unfinished), and the rating (i.e., the age of the target group).⁶ We also add original work fixed effects as groups within the community (e.g., fans of *Harry Potter* or

⁶The status of a text (finished/unfinished) is considered not to be a quality indicator but indicates that for the writer the continuity of the story is more important than the full closure of the work.

Mortal Instruments) may differ in terms of their interaction cultures and activity levels, among others. θ_t are month and year fixed effects to control for seasonality in writing activities and general time trends (see Section 3). $\varepsilon_{i,j,t}$ denotes the error term.

Table 2 indicates that first-time writers are indeed more inclined to keep publishing the more community feedback they get: a 1% increase in the number of reviews is associated with a decrease in the probability to publish only one text by 7.4 percentage points (column (1)). This translates into a 12% decrease when evaluated at the sample mean. The estimated β_1 is slightly smaller for followers and substantially smaller for ‘favoriting’ (columns (2) and (3)).⁷ This is in line with expectations as reviews can be categorized as feedback of the highest quality. Note that all three feedback types have similar distributions with a mean between 22 and 26 (see Table 1).

Our estimates, however, should be interpreted with caution and may not reflect a causal effect of feedback on the willingness to stay in the community. This is because the unobserved writing talent could work as a confounder. Having said this, the fact that prior research has identified similar effects in randomized field experiments suggests that community feedback coupled with the feeling of inclusion prevents novices to leave.

4.2. Writing times

Next, we investigate the effect of community feedback on the time span between publications. Since feedback now relates to the foregoing text ($i - 1$), debuts were discarded. We estimate a model similar to (1) supplemented by author fixed effects (ϕ_j) and the count of contributions authored by writer j (*ContrCount*). That is, we estimate

$$\begin{aligned} WritingTime_{i,j,t} = & \beta_0 + \beta_1 \ln Feedback_{i-1,j,t} + \beta_2 ContrCount_{i,j,t} \\ & + \beta_3 ContrCount_{i,j,t} \cdot \ln Feedback_{i-1,j,t} + \xi_i + \phi_j + \theta_t + \varepsilon_{i,j,t}. \end{aligned} \quad (2)$$

While author fixed effects account for unobserved heterogeneity between writers, we take the contribution count as a proxy for experience. The interaction term then refers to the

⁷As a robustness check, we replicate the analysis using a Logit estimator. Results are very close to the LPM estimates. For instance, the average marginal effect of reviews on the probability of leaving the community for first-time writers is -0.072 (SE 0.001).

Table 2: Community feedback on debuts and the probability to leave.

	(1)	(2)	(3)
Log(Reviews+1)	-0.074*** (0.001)		
Log(Favoriting+1)		-0.060*** (0.001)	
Log(Followers+1)			-0.053*** (0.001)
additional controls	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
original work FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
genre dummies	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
month and year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	193,163	193,163	193,163
R^2	0.119	0.107	0.103

- Dependent variable: *single* = 1 if the debut remains the only publication (zero otherwise). Mean of *single* is 0.611.
- Coefficients are estimated in an OLS regression framework.
- Additional controls: rating dummies, status dummy, language dummies.
- Robust standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

idea that the effect of previous feedback might vary with experience. As further controls, we added the text status and the writer’s ‘fanfiction age’ (i.e., the time elapsed since registration).

Table 3 presents the results. It shows that *reviews* have a strong effect on writing times: a 1% increase in reviews for the foregoing text increases the time until the next publication by 22 to 27 days (columns (1) and (2)). Given that the mean time span between publications is 127 days, this means an increase by around 18 to 21%. Compared to reviews, the effect for the two types of low-quality feedback is only half the size (columns (3) to (6)).

Note that these are the estimates from a fixed effects model and hence represent within-individual responses to feedback. In this sense, we interpret our results to mean that community feedback may work as ‘positive pressure’ which makes fan writers to put more effort in their works, working longer to improve the quality of the text. Alternatively, more community attention may put some ‘pressure of high expectations’ on the authors, causing them to have more doubts about their creative output, and to take more time before publishing it. However, this effect diminishes when experience increases as $\hat{\beta}_3$ is

negative and significantly different from zero.⁸ More experienced writers, on average, need less time to produce output.

Table 3: The effect of feedback on writing times.

	<i>Reviews</i>		<i>Favoriting</i>		<i>Followers</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Reviews _{<i>i</i>-1} +1)	22.334*** (0.585)	26.694*** (0.859)				
Log(Reviews _{<i>i</i>-1} +1) · ContrCount		-0.418*** (0.075)				
Log(Favoriting _{<i>i</i>-1} +1)			10.016*** (0.605)	12.723*** (0.975)		
Log(Favoriting _{<i>i</i>-1} +1) · ContrCount				-0.226*** (0.069)		
Log(Followers _{<i>i</i>-1} +1)					10.739*** (0.534)	13.695*** (1.087)
Log(Followers _{<i>i</i>-1} +1) · ContrCount						-0.275*** (0.093)
ContrCount	-2.267*** (0.371)	-1.574*** (0.278)	-2.263*** (0.375)	-1.830*** (0.346)	-2.264*** (0.374)	-1.964*** (0.336)
author FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
additional controls	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
original work FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
genre dummies	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
month and year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Constant	84.155*** (4.800)	75.856*** (4.104)	113.171*** (4.712)	107.697*** (4.625)	117.600*** (4.708)	113.648*** (4.414)
Observations	224,427	224,427	224,427	224,427	224,427	224,427
<i>R</i> ²	0.578	0.579	0.573	0.573	0.573	0.574

- Dependent variable: *WritingTime* is the time elapsed since the last contribution. *ContrCount*: count of contributions.
- Coefficients are estimated in a OLS regression framework.
- Robust standard errors (clustered on the author level) in parentheses, * p<0.1, ** p<0.05, *** p<0.01.
- Additional controls: rating dummies, text status dummy (finished/unfinished), language dummies, writer's 'fanfiction age' (i.e., the time elapsed since registration).

⁸Note that results are qualitatively the same when *Feedback* is defined as the average number of reviews before text *i*. In the same way, including the number of words of text *i* or the log of the average number of words of writer *j* up to text *i* hardly changes any of the estimates and the *R*². We therefore conclude that text length preferences are already captured by the author fixed effects.

4.3. Text lengths

Finally, to estimate the effect of community feedback on the extend of UGC, we regress the text length on the (average) number of past reviews, followers, and ‘favoriting’. Again, debuts were discarded. The model, despite being log-linear, is very similar to the foregoing model defined by Equation (2):

$$\begin{aligned} \ln Words_{i,j,t} = & \gamma_0 + \gamma_1 \ln Feedback_{i-1,j,t} + \gamma_2 ContrCount_{i,j,t} \\ & + \gamma_3 ContrCount_{i,j,t} \cdot \ln Feedback_{i-1,j,t} + \xi_i + \phi_j + \theta_t + \varepsilon_{i,j,t}. \end{aligned} \quad (3)$$

The dependent variable $Words_{i,j,t}$ is the length of text i written by author j in month t .

Table 4: The effect of feedback on text lengths.

	$\ln Words_i$				$\ln Words_i / \ln Words_{i-1}$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln Reviews_{i-1}$	0.315*** (0.004)	0.274*** (0.003)	-0.060*** (0.003)	-0.066*** (0.003)	-0.095*** (0.001)	-0.099*** (0.001)
$ContrCount$	-0.003*** (0.000)	-0.003*** (0.000)	0.001*** (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.001*** (0.000)
$\ln Reviews_{i-1}$ $\cdot ContrCount$				0.001*** (0.000)		0.000*** (0.000)
author FE	<i>No</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
additional controls	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
original work FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
genre FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
month and year FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	279,597	258,393	224,422	224,422	221,796	221,796
R^2	0.088	0.299	0.646	0.616	0.342	0.343

- Dependent variable: $Words_i$ length of text i , $Words_i / Words_{i-1}$ is the relation of text i 's length to the length of the foregoing text. $ContrCount$: count of contributions.
- Coefficients are estimated in a OLS regression framework.
- Robust standard errors (clustered on the author level) in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
- Additional controls: rating dummies, text status dummy (finished/unfinished), language dummies, writer's ‘fanfiction age’ (i.e., the time elapsed since registration).
- The number of observations varies due to missing values.

Table 4 shows the results. Surprisingly, after the inclusion of author fixed effects, the estimated γ_1 changes sign (columns (3) and (4)). As an explanation, we refer to the unbalanced structure of content provision by users on the extensive and the intensive margin. To illustrate the heterogeneity among authors, we estimated the specification

used in column (2) of Table 4 for different categories of writers. Specifically, we determined the average text lengths per author prior to text i and then grouped individuals into six categories according to five percentiles of the overall distribution (P25, P50, P75, P90, P99). Figure 3 shows that the estimated γ_1 's varies in sign and magnitude: the producers of short pieces respond to a rise in reviews with an increase of text lengths, whereas the very opposite is the case for the producers of long works.

In other words, once we control for within-author variations, we find that community feedback, on balance, decreases the extend of user-generated content. Specifically, if the number of reviews for text $i-1$ increases by 1%, the length of text i decreases by around 6% (4.4% for 'favoriting' and 5.3 % for followers, see Table B.1 in Appendix B). In addition, $\hat{\gamma}_3 < 0$ indicates that this effect weakens when experience increases. Given that the mean of the dependent variable is 8,737 (median: 1,934), this means a reduction of 524 (116) words. To provide further evidence, we use the relation of text i 's length to the length of the foregoing text as dependent variable (columns (5) and (6)). Results for this relative measure of text length do not differ qualitatively. Including the time span between text i and $i-1$ or the log of the average writing time for writer j up to text i hardly changes any of the estimates.⁹


Overall, in light of the considerations raised in Section 2, we may take this as evidence that feedback (reviews) works as a correction towards an optimal text length: authors of short texts are encouraged to write more and vice versa.

5. Stimuli from the original work

The forgoing analyses has helped to understand how peer feedback affects user-generated content production in the online community under consideration. In this part of the analysis, we aim to assess in which ways the release of new content affects fan output. We therefore focus on book series which allow us to estimate the effect of new content on fan work: 33 within the top 300 titles in our sample are book series with at least one new release (book or movie) in the observation period (see Table B.2 in the Appendix for

⁹Results can be made available upon request.

Figure 3: Estimated γ_1 for different categories of writers (reviews).




details).

5.1. More contributions?

In a first approach, we examine whether new material has an impact on the extent of fan text releases. Specifically, we follow Graddy and Lieberman (2018) and create binary variables which refer to the time before and after the month the new material was published in order to estimate a fixed effects model for the original work with the sum of fan texts per month on that topic being our dependent variable. Since the median time span between two publications is 27 days and fans need time to read first, we use seven binary variables indicating whether there was the release of a new original book (or movie) three months before or after the publication month or whether it was published the same month. This procedure also ensures that the estimates will not be biased by further inputs from the original work as we define a minimum interval of six months between the actual new release and subsequent releases of new content. Taken together, estimates from that model would indicate variations in output as a direct response by fans exposed to new content or visual material.

Table B.3 in the Appendix shows that the model does a poor job in explaining variations in fan text publications according to the release of new material. We conclude that

Figure 4: Contributions for *Harry Potter*



aggregate data on a monthly basis may not be suitable to capture the complexity of the writing process or that countervailing effects (some are challenged to go deeply into the new content while others wish to contribute as soon as possible) cancel each other out, at least in the short run. For the case of *Harry Potter*, Figure 4 suggests that it is hard to identify a clear pattern of fan responses as we have a drop after the release for some events and a boost for others.

5.2. Does new material affect writing times?

The prior analysis suggests that new content, on average, does not directly affect the number of fan publications. This might be because of the writer heterogeneity and the complexity of the writing process that cannot be pinned down by a monthly measure. We try to address this issues and focus on the writing time next, that is the period between two releases.

Specifically, we estimate a Cox proportional hazards model defined by

$$h_i(t) = h_0(t) \exp(\beta'z), \tag{4}$$

where $h_0(t)$ is the baseline hazard function, z a set of covariates similar to those used

in Section 4, and β a vector of regression coefficients. In our setting, the hazard rate is the likelihood that a new text is published by writer i at time t . We define those writers who were exposed to new original content (books or movies) as ‘treated’, meaning that the original content was released within the interval between two publications. In other words, we aim to estimate the causal effect of the fresh material on writing times.

We include the text length, proxies for prior feedback (the log of reviews related to the previous text) and experience (‘platform age’ and number of prior publications), and a set of dummy variables that account for original work, genre, category, text status (finished/unfinished), language, month and year of release as control variables. The full sample is used to estimate the model.

Table 5: New material and writing times – Cox regressions


	All writers			Top writers	
	(1)	(2)	(3)	(4)	(5)
New book	-0.788*** (0.009)	-0.767*** (0.012)	-0.719*** (0.012)	-0.881*** (0.050)	-0.845*** (0.054)
New movie			-0.756*** (0.012)	-1.078*** (0.041)	-0.962*** (0.045)
additional controls	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
original work FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
genre FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
rating FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
english dummy	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
year FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
<i>N</i>	136,633	136,633	136,633	11,932	11,932

- Dependent variable: *WritingTime* is the time elapsed since the last contribution.
- Additional controls: *lnReviews*, *ContrCount*, ‘FanFiction’ age, number of words, text status (finished/unfinished).
- *Top writers*: Writers in the 75th percentile of prior publications (14) and ‘Followers’ (308).
- Robust standard errors in parentheses (clustered on the author level), * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The estimated coefficients presented in Table 5 indicate that new original material significantly reduces the ‘hazard of publishing’: conditional on text length, it takes more time for writers to produce content. For instance, the point estimate of -0.77 for new books (column (2)) suggests that they decrease the likelihood of a new release by 54.62% (which is $100 * (\exp(-0.77) - 1)$). Figure 5 illustrates this result.

To account for writer heterogeneity, we run additional regression with a sample reduced to the ‘stars of the scene’, meaning writers in the 75th percentile of publications ($p_{75} = 14$)

Figure 5: Cox proportional hazards regression.



and ‘Followers’ ($p_{75} = 308$). The estimates in columns (4) and (5) show that the effect is even more pronounced in the top writer segment. For instance, the estimated coefficient of *New book* translates into an 57.90% reduction in the likelihood of publication (column (5)). This finding suggests that instead of a quick response to new content in order to get ahead of the herd, gain the privilege of interpretation, and to ‘stay in the game’, top writers take their time to properly address the new ideas, characters, and twists. As an explanation, we refer to user attention coupled with expectations that work as a kind of quality control. Note that while the number of ‘favoriting’ per text does slightly decrease for the average writer after the release of new material (means of 31.81 and 28.53, p-value = 0.130), this is not true for the elite of top writers (40.95 vs. 39.31, p-value = 0.842).

6. Concluding remarks

In this paper, we explore possible stimuli of generating new free content. Online fanfiction communities provide a unique example of User-Generated Content (UGC). It is built upon the original narrative, so the new material and expansion of the original series may itself serve as a booster of a content provision. In other UGC settings, it is much harder to separate the shock coming from the subject and other stimuli. At the same time, community feedback, which is mostly positive and encouraging, can serve as a driver for the further content provision, as well as the improvement of the quality of the content. The audience attention, especially towards the fanfiction for the popular originals, can be viewed as a source of competition improving productivity and creative process (as suggested in Gross, 2020; Wu and Zhu, 2022). It is also widely accepted that the fanfiction community is no longer just the marginalized group of teenagers, but is among the drivers of the modern popular culture.

Our estimation results suggest that the community feedback indeed stimulates the writers to produce more texts and to continue creating the fan writings after the debut. Moreover, other things equal, the authors receiving more feedback take more time to produce the next text, and produce shorter texts. Short text producers, however, increase the text length after receiving more feedback. This is likely to be explained by the internal quality control, trying to deliver well-shaped characters in a more concise text in the response to the audience attention. These findings are in line with other literature on the peer feedback in the UGC provision, for example Burtch et al. (2021).

The novel finding concerns the stimuli coming from the original text. We find that symmetric shocks from the release of new material significantly reduce the ‘hazard of publishing’: conditional on text length, it takes more time for writers to produce content. Moreover, the result is even more pronounced for the most popular writers. We interpret this to mean that top writers prefer to properly develop the ideas, knowing that the vast and loyal audience will be willing to wait for the text by them specifically.

Hence, our findings support the idea that user attention (expressed in expectations, loyalty, and reputation) can work as quality control for UGC. Related to the findings of Wu

and Zhu (2022) and Gross (2020), it appears that a competitive environment not necessarily gives incentives to content producers to increase output at the expense of quality. We attribute this result to the review tool coupled with the positive, constructive atmosphere required by the *FanFiction.Net* community guidelines. In this regard, the website might provide a positive example for other UGC platforms. Moreover, we think that our findings can be transferred to competitive creative work settings beyond UGC. For instance, in science, the sharing of first results with peers in order to receive a ‘friendly review’ works in a similar way. So, it appears likely that self-upload platforms in other areas of creative work such as (amateur and early stage) music and film production would also benefit from (enforced, respected, and practiced) rules for feedback and general conduct.

References

- Aaltonen, A. and Seiler, S. (2016). Cumulative growth in user-generated content production: evidence from wikipedia. *Management Science*, 62(7):2054–2069.
- Borowiecki, K. J. (2013). Geographic clustering and productivity: An instrumental variable approach for classical composers. *Journal of Urban Economics*, 73(1):94–110.
- Burtch, G., He, Q., Hong, Y., and Lee, D. (2021). How do peer awards motivate creative content? experimental evidence from reddit. *Management Science*.
- Burtch, G., Hong, Y., Bapna, R., and Griskevicius, V. (2018). Stimulating online reviews by combining financial incentives and social norms. *Management Science*, 64(5):2065–2082.
- Cabral, L. and Li, L. (2015). A dollar for your thoughts: Feedback-conditional rebates on eBay. *Management Science*, 61(9):2052–2063.
- Chen, Y., Harper, F. M., Konstan, J., and Li, S. X. (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review*, 100(4):1358–98.
- Contrera, J. (2014). From ‘fifty shades’ to ‘after’: Why publishers want fan fiction to go mainstream. *The Washington Post*. https://www.washingtonpost.com/lifestyle/style/from-fifty-shades-to-after-why-publishers-want-fan-fiction-to-go-mainstream/2014/10/24/825d6a94-5a04-11e4-b812-38518ae74c67_story.html, accessed March 2022.
- Coppa, F. (2017). *The Fanfiction Reader: Folk tales for the digital age*. University of Michigan Press (USA).
- DeVaro, J., Kim, J.-H., Wagman, L., and Wolff, R. (2018). Motivation and performance of user-contributors: Evidence from a CQA forum. *Information Economics and Policy*, 42:56–65.
- Fraiberger, S. P., Sinatra, R., Resch, M., Riedl, C., and Barabási, A.-L. (2018). Quantifying reputation and success in art. *Science*, 362(6416):825–829.
- Gallus, J. (2017). Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia. *Management Science*, 63(12):3999–4015.
- Goes, P. B., Guo, C., and Lin, M. (2016). Do incentive hierarchies induce user effort? evidence from an online knowledge exchange. *Information Systems Research*, 27(3):497–516.
- Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.
- Graddy, K. and Lieberman, C. (2018). Death, bereavement, and creativity. *Management Science*, 64(10):4505–4514.

- Gross, D. P. (2020). Creativity under fire: The effects of competition on creative production. *Review of Economics and Statistics*, 102(3):583–599.
- Huang, N., Burtch, G., Gu, B., Hong, Y., Liang, C., Wang, K., Fu, D., and Yang, B. (2019). Motivating user-generated content with performance feedback: Evidence from randomized field experiments. *Management Science*, 65(1):327–345.
- Jamison, A. (2013). *Fic: Why fanfiction is taking over the world*. BenBella Books (USA).
- Khern-am nuai, W., Kannan, K., and Ghasemkhani, H. (2018). Extrinsic versus intrinsic rewards for contributing reviews in an online platform. *Information Systems Research*, 29(4):871–892.
- Lipton, J. D. (2014). Copyright and the commercialization of fanfiction. *Hous. L. Rev.*, 52:425.
- Marchenko, M. (2020). Peer effects in art prices. *WU Working paper*.
- Penny, L. (2019). We can be heroes: how the nerds are reinventing pop culture. *Wired*. <https://www.wired.com/story/culture-fan-tastic-planet-fanfic/>, accessed February 2022.
- Shriver, S. K., Nair, H. S., and Hofstetter, R. (2013). Social ties and user-generated content: Evidence from an online social network. *Management Science*, 59(6):1425–1443.
- Sun, M. and Zhu, F. (2013). Ad revenue and content commercialization: Evidence from blogs. *Management Science*, 59(10):2314–2331.
- Thomas, B. (2011). What is fanfiction and why are people saying such nice things about it?? *Storyworlds: A Journal of Narrative Studies*, 3:1–24.
- Toubia, O. and Stephen, A. T. (2013). Intrinsic vs. image-related utility in social media: Why do people contribute content to twitter? *Marketing Science*, 32(3):368–392.
- Vedres, B. (2017). Forbidden triads and creative success in jazz: the Miles Davis factor. *Applied Network Science*, 2(1).
- Wang, J., Li, G., and Hui, K.-L. (2021). Monetary incentives and knowledge spillover: Evidence from a natural experiment. *Management Science*. Forthcoming.
- Wu, Y. and Zhu, F. (2022). Competition, contracts, and creativity: Evidence from novel writing in a platform market. *Management Science*. Forthcoming.
- Xu, L., Nian, T., and Cabral, L. (2019). What makes geeks tick? A study of stack overflow careers. *Management Science*.
- Yin, K., Aragon, C., Evans, S., and Davis, K. (2017). Where no one has gone before: A meta-dataset of the world’s largest fanfiction repository. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 6106–6110. ACM.
- Zhang, X. M. and Zhu, F. (2011). Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review*, 101(4):1601–15.

A. Additional Figures

Figure A.1: Reviews example 1

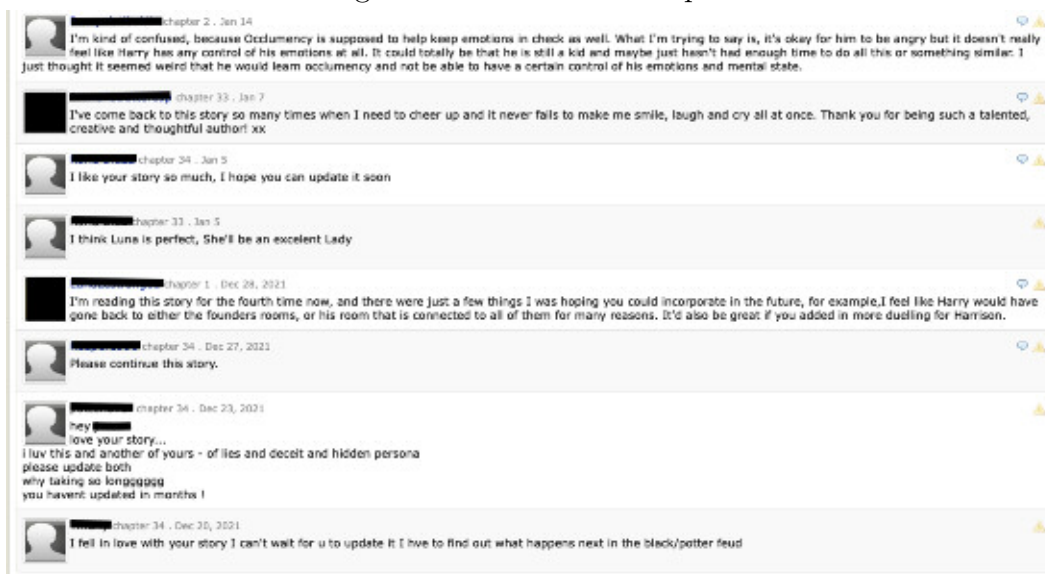


Figure A.2: Reviews example 2



Figure A.3: Number of words (excluding strong outliers) – kernel density estimates


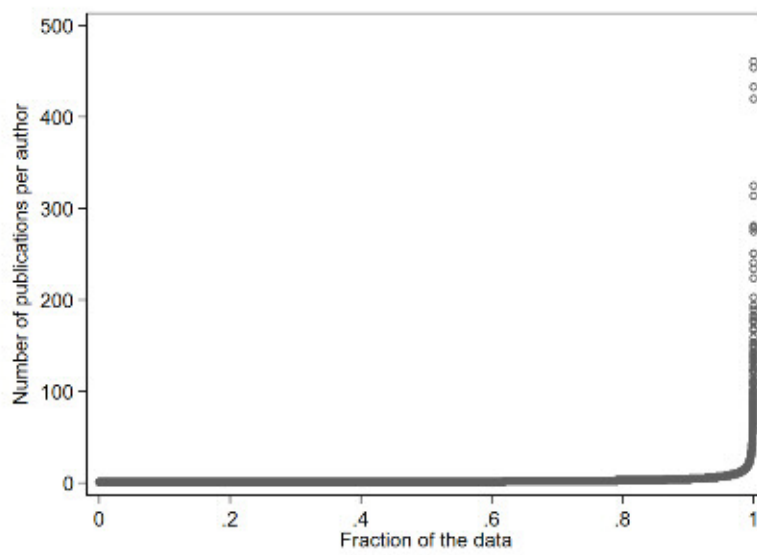


Figure A.4: Number of fan texts per author – Quantile plot



B. Additional Tables

Table B.1: The effect of feedback on text lengths: favoriting and Followers.

	<i>favoriting</i>			<i>Followers</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln\text{favoriting}_{i-1}$	0.312*** (0.004)	-0.044*** (0.003)	-0.052*** (0.003)			
ContrCount	-0.004*** (0.000)	0.001*** (0.000)	-0.001 (0.000)	-0.003*** (0.000)	0.001*** (0.000)	0.000 (0.000)
$\ln\text{favoriting}_{i-1}$ · ContrCount			0.001*** (0.000)			
$\ln\text{Followers}_{i-1}$				0.313*** (0.004)	-0.053*** (0.003)	-0.060*** (0.003)
$\ln\text{Followers}_{i-1}$ · ContrCount						0.001*** (0.000)
author FE	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>
additional controls	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
original work FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
genre FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
month-year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	279,597	224,422	224,422	279,597	224,422	224,422
R^2	0.093	0.645	0.645	0.100	0.646	0.646

- Dependent variable: $Words_i$ length of text i , $Words_i/Words_{i-1}$ is the relation of text i 's length to the length of the foregoing text. ContrCount : count of contributions.

- Coefficients are estimated in a OLS regression framework.

- Robust standard errors (clustered on the author level) in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

- Additional controls: rating dummies, status dummy, language dummies, writer's 'fanfiction age' (i.e., the time elapsed since registration).

- The number of observations varies due to missing values.

Table B.2: List of Book series used in Chapter 5

Rank	Title	No. fan texts
1	Harry Potter	80827
2	Percy Jackson and the Olympians	65287
3	Hunger Games	37744
4	Mortal Instruments	13373
5	Gossip Girl	8466
6	Divergent Trilogy	6686
7	A song of Ice and Fire	5870
8	Maximum Ride	5850
9	Inheritance Cycle	5242
10	Twilight	5100
11	Artemis Fowl	4832
12	Gallagher Girls	4332
13	Clique	3684
14	Alex Rider	3376
15	39 Clues	2806
16	Maze Runner Trilogy	2540
17	Vampire Academy	2200
18	Warriors	2103
19	Sisters Grimm	1915
20	Skulduggery Pleasant series	1874
21	Infernal Devices, Cassandra Clare	1800
22	Ranger's Apprentice	1775
23	Series Of Unfortunate Events	1665
24	Morganville Vampires	1649
25	Darren Shan Saga / Cirque Du Freak	1593
26	Vampire Diaries	1495
27	House of Night	1392
28	Darkest Powers	1356
29	Fifty Shades Trilogy	1133
30	Wheel of Time	1090
31	Gone	1047
32	Sookie Stackhouse / The Southern Vampire Mysteries	1044
33	Kane Chronicles	864

Notes: Books are ranked according to the number of fan texts.

Table B.3: The amount of fan text publications in response to new material.

	(1)	(2)	(3)	(4)
New book	-11.562 (10.594)	-11.568 (10.629)	11.051 (14.135)	13.611 (20.678)
New film		14.106 (45.498)	13.255 (41.148)	13.083 (52.761)
New book (t+1)				14.748 (22.658)
New book (t+2)				8.205 (20.753)
New book (t+3)				6.184 (19.202)
New book (t-1)				3.869 (20.246)
New book (t-2)				4.746 (19.407)
New book (t-3)				5.124 (17.648)
New film (t+1)				41.839 (66.794)
New film (t+2)				15.674 (51.961)
New film (t+3)				6.208 (51.368)
New film (t-1)				-20.500 (43.019)
New film (t-2)				-22.933 (42.026)
New film (t-3)				-20.329 (38.897)
Year dummies	<i>No</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Month dummies	<i>No</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Original work FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
<i>N</i>	4,400	4,400	4,053	4,053
<i>R</i> ²	0.000	0.000	0.048	0.050

^{*} Dependent variable: Fan contributions per month and topic (original work).

^{*} Coefficients are estimated in a OLS regression framework (robust standard errors in parentheses)