Journal of Cultural Economics
https://doi.org/10.1007/s10824-024-09508-y

ORIGINAL ARTICLE

The never-ending book: the role of new material and peer feedback in user-generated content production

Maria Marchenko[1] · Hendrik Sonnabend[2]

Received: 28 March 2023 / Accepted: 27 February 2024

Abstract

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: writing times increase for the average writer and even more for the elite of prolific writers. We explain this finding with quality concerns. In addition, we find supportive evidence that community feedback encourages first-time contributors to continue publishing. However, for more established writers, community feedback has a rather dampening effect on text lengths and writing times. These effects are more pronounced for more informative community feedback ('reviews') than less informative community feedback ('following', 'favoriting').

Keywords Fanfiction · User-generated content · Online public goods · Voluntary contribution

JEL Classification H41 · C31 · D01 · Z11

1 Introduction

Digital technology, i.e. the representation of information in bits, has reduced the costs of storage, computation, and transmission of data (Goldfarb & 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). UGC refers to online content created not by paid professionals but by the general public.

✉ Hendrik Sonnabend
hendrik.sonnabend@fernuni-hagen.de

Maria Marchenko
maria.marchenko999@gmail.com

1 Independent researcher, Vienna, Austria

2 Department of Economics and Business Administration, University of Hagen, Universitätsstraße 11, 58097 Hagen, Germany

Published online: 30 March 2024

Examples include video sharing, social media platforms, blog posts, and question-and-answer websites. UGC has an apparent free-rider problem; therefore, a growing body of literature tries to understand why people change their role from consumers to creators of public goods and what motivates users to contribute content for free.

Previous research has highlighted several different drivers behind user-generated content production. Among them is social motivation like group size (Zhang & Zhu, 2011), social ties/'friends' (Shriver et al., 2013; Goes et al., 2014), social norms (Chen et al., 2010; Burtch et al., 2018), past contributions (Aaltonen & Seiler, 2016), financial incentives (Sun & Zhu, 2013; Cabral & Li, 2015; Khern-am nuai et al., 2018; Burtch et al., 2018; Wang et al., 2022), non-monetary awards (Gallus, 2017; DeVaro et al., 2018; Burtch et al., 2022) and performance feedback (Huang et al., 2019), and image-related or 'glory-based' utility derived from contributing (e.g. Toubia & Stephen 2013; Goes et al.,2016; Xu et al., 2020).

However, the question of how the *subject itself* motivates the production of UGC has received less attention so far. For instance, fans may respond to new music from their favourite artists, and users could engage in writing unofficial patches after the release of a new video game (we will call all of this 'subject' from now on). This lack of research might be because most studies mentioned above use data from popular websites related to informational content, like the open online encyclopaedia *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 addresses this issue by using data from the world's largest fanfiction website *Fanfiction.net*. Fanfiction is literature written by mostly amateurs 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 to work with the new material, fan writers must study it. This is not possible before the release. Additionally, each shock is tightly connected to only one cluster of fanfiction literature, with just a couple of exceptions. For example, a new book or movie about 'Harry Potter' will affect the creation of 'Harry Potter' fanfiction and is not likely to affect the writing of, for example, 'Lord of the Ring' fanfiction.

Results from an event study approach and Cox regressions indicate that no clear effect on the number of contributions can be found. However, when exposed to new material, the average fan writer needs more time to finish a text. 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.

Besides stimuli from the subject itself, community feedback is another source of motivation for the production of UGC. Many online platforms use feedback to stimulate content contribution (Huang et al., 2019). Examples are the badges at *Stack Overflow*, motivating statistics at *Academia.edu* and *ResearchGate*, and 'karma points' at *Reddit*. In a *Reddit*-based field experiment, Burtch et al. (2022) 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) demonstrate that individual-level contributions increase with audience size. However, even a 'silent' audience can affect UGC production when the other users' attention is revealed, for example, by the 'following' feature on social media sites (Goes et al., 2014).

In our setting, the fanfiction website allows for three kinds of community feedback: 'favoriting', 'followers', and reviews. While 'favoriting' and 'following' the work (or even the writer) are the two ways of showing interest in a fan writer's work, the reviews suggest an even higher level of involvement. We find that with more feedback, first-time contributors are more likely to continue publishing their works on the platform, which aligns with the findings of Burtch et al. (2022). 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 more informative feedback and text lengths support the idea of an optimal text length expressed by user reviews.

The rest of the paper is organised as follows. Section 2 introduces the fanfiction platform and provides 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 release of new original content. Section 6 summarises the results and discusses their implications for online communities with UGC production.

2 Institutional background

Fanfiction (also spelt as fan fiction, abbreviated as fan fic, or fanfic) is literature by mostly amateur writers. It is based on sources like books, movies, TV shows, or even the life of a particular celebrity. Fanfiction narratives usually develop an alternative fable or new take on the original story or discuss secondary characters. Most fanfiction writers are 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 marginalised 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 online platforms. The early examples of fandoms go back to Sherlock Holmes stories and Jane Austen's books (Jamison, 2013). Fans have been fantasising about alternative universes or the continuation of their favourite 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 not 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 often feel a connection not only to the narrative of the fanfic but also to the authors. The distribution of the readers' attention to the authors is quite uneven, with a relatively small share of superstar writers. These writers get more feedback on their work in the form of reviews, favouriting or following the updates of the text or the authors.

On online fanfiction platforms, it is implied that the feedback to the new pieces is primarily 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 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. Naturally, the writers would want to respond to the new material; however, the response may differ from the top and average writers. Top writers, often needing to meet the audience's expectations, will likely 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 writers, which should motivate the focused writer to create the content faster, and the feedback from the audience. In the case of top writers, the latter may put pressure on 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. However, the comments from her followers kept appearing on her page even three years later, inquiring for new material.

Like all the other examples of the UGC, fanfiction rarely brings profit to its creators and is published outside the community. As Coppa (2017, p. 14) notes, fanfiction is not written and distributed for financial gain; it is 'made for free, but not 'for nothing'' as the stories are considered gifts to a community 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 *Twilight* fanfiction. *New York Times* bestselling author Cassandra Claire was also a fanfiction writer. In her *Wired* piece, Laurie Penny says that a 'surprising number' of film and TV writers, editors, journalists, and novelists used to write fanfiction (Penny, 2019).[1] It is worth noting that the phenomenon of users who are successful enough to turn their hobby into a profession is not exclusive to fanfiction and can also be found in other UGC domains such as software development, video sharing, and 'blogging'.

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 (mainly written in English) by over ten million registered users. The scope of the platform is comparable with *Wikipedia*, with more registered users—around 43 million, but

[1] Note that there is also a close connection (in terms of content and writers) between fanfiction and 'new adult' literature, which is a sub-genre of the 'romance' genre.

only around 124,000 active editors. The platform allows users to follow a particular author and stories, often published in a series of chapters or smaller bits and pieces.

Users can leave reviews or simplified positive feedback by 'favoriting' or 'following' stories. It is a convention in the community that the readers support the authors, and the reviews bear primarily positive feedback. It is also the strongest signal, as it suggests the highest involvement of the readers because it requires more actions than just hitting the 'like' button. Figure 6 provides a small sample of reviews for work with many reviews. Out of more than 5,000 reviews, only some slightly negative reviews do not agree with the characters or complain about the long waiting time. The reviews are encouraging for the less popular texts, even when there are only a few (see Fig. 7 for a further example). This observation aligns with Evans et al. (2017), who find that only 1% of the reviews in their sample can be coded as 'non-constructive negative'.[2]

Both 'favoriting' and 'following' are sending only positive feedback. Adding a text or an author to favourites signals the 'liking' of their work(s). The 'following' is more substantial feedback as the user signs up to receive updates on the work or all the author's works by following a text or author. So, all three ways of the community response to the text bear positive feedback; however, they differ in the degree of involvement and may affect writers differently. 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 the first step, all available entries in the category 'Books' between 06/1999 and 12/2017 were collected from *Fanfiction.net*. These are 482,838 observations on fictional texts created by 203,233 authors and built on 2342 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'). Note that our data do not include time stamps for the community response variables as they are not part of a text's meta-information. We discuss possible implications in the next section.

Table 1 presents summary statistics. First, on average, the website publishes 3247 contributions per month. Since 1999, we have geometric growth followed by a period of consolidation at a high level, see Fig. 1. The drop in uploads around 2014 might be explained by the emergence of *Amazon*'s fanfiction platform *Kindle Worlds* (Lipton, 2014). Even though the initial number of fandoms joining the platform was relatively small (only 24, see Contrera, 2014), Lipton (2014) suggests the emergence

[2] The platform does not have moderators for reviews, but writers can moderate (guest-)reviews from unregistered users and report inappropriate reviews from registered users to the administrators.

Table 1 Summary statistics for fanfiction

Variable	Obs	Mean	SD	Min	Max
Contributions per month[a]	452,546	3246.764	1230.060	1	5477
Contributions per original	2342	206.165	2401.889	1	80,827
Contributions per author	203,233	2.376	5.273	1	461
Words per text	482,838	8736.842	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

[a] For 30,292 observations, the date of publication was missing

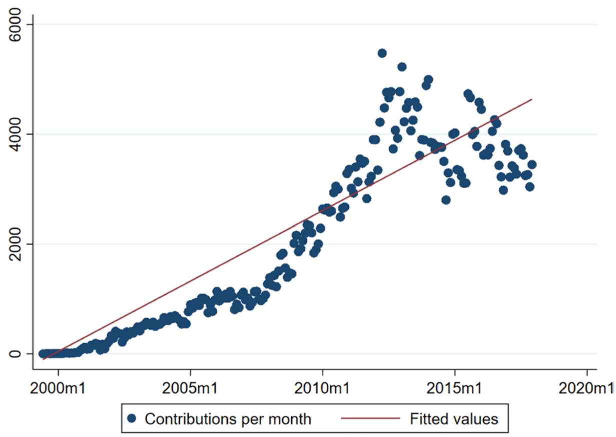

Fig. 1 Contributions per month

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 for the selected fandom. Additionally, there were ongoing community concerns about the introduction of the Stop Online Piracy Act (SOPA), which could have caused some decrease in the activity.[3] However, like Yin et al. (2017), we do not know any legal changes at the time. Yin et al. (2017) also find hints of seasonality with more writing activities in the (northern hemisphere) summer months and smaller peaks in December.

Second, contributions concentrate intensely 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,

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

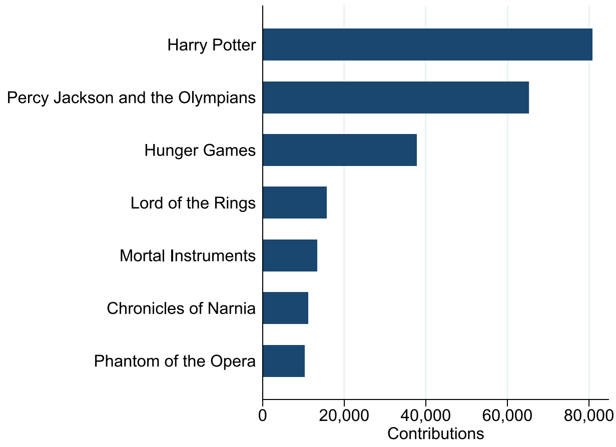

Fig. 2 Contributions per fandom (top 7 titles)

and 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 and the output per author vary widely (see Figs. 8 and 9 in "Appendix A"). While most users have only a limited number of uploads, there are also prolific writers with hundreds of publications. This pattern is common in social network structures where users can contribute openly.[5] Accordingly, most writers with more than one publication target only one source text (60.12%), and 29.01% create texts on two topics. On average, the time between publications is 127 days (median: 27 days).

In the second step, data on the original works were collected from the *Wikipedia* encyclopaedia and the social cataloguing 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 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 analysing 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

[4] We only have the data on the 'favoriting' and 'following' of the particular texts. The 'favoriting' and 'following' of the authors is not available.

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

feedback affects users differently. For instance, Burtch et al. (2022) show that feedback is more important for new members since this group is less connected to the community and feels more uncertain. In the same way, the findings documented by Goes et al. (2014) suggest that the general effect of ‘silent’ peer attention (through ‘following’) on user activities is positive but that the marginal effect decreases in the popularity of the focused user.

As mentioned in Sect. 2, in our setting, three types of feedback are available to users: reviews, (text) ‘following’, and (text) ‘favoriting’. We expect reviews to exceed the others in terms of information and hence expect this type to have the highest impact on UGC production. Moreover, since there is a strong correlation between the different types of feedback (see Sect. 3), we refrain from using them simultaneously as independent variables in most of the cases in our regression models. Finally, as noted in Sect. 3, we do not have information on the timing of feedback. Thus, parameter estimates would only capture the actual effect for cases where all the feedback on the previous work was available to the writer before the writing process related to the focused work ended. While we do not expect a systematic bias coming from ex post (and hence for the writing process irrelevant) feedback, our estimates should be interpreted as upper bounds of the actual effects.

4.1 Community feedback and first-time writers

We first examine whether feedback affects the propensity to stay in the community and contribute more than one text. Therefore, the sample was restricted to the first publications (debuts) (193,163 observations, one per writer). Formally, we estimate a model defined by

$$Stay_{i,j,t} = \alpha_0 + \alpha_1 lnFeedback_{i,j,t} + \xi_i + \theta_t + \varepsilon_{i,j,t} \, , \tag{1}$$

where $Stay_{i,j,t}$ is a binary dependent variable which equals 1 if text i published in month t is not 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 text, 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 *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 Sect. 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 an increase in the probability of continuing writing after the first text by 7.4 percentage points (column (1)). This translates into a 19% increase when

[6] The status of a text (finished/unfinished) is considered not a quality indicator. However, it indicates that for the writer, the continuity of the story is more important than the complete closure of the work.

Table 2 Community feedback on debuts and the probability of leaving

	The first publication is the only text			
	(1)	(2)	(3)	(4)
Log(Reviews+1)	0.074***			0.084***
	(0.001)			(0.002)
Log(Favoriting+1)		0.060***		0.029***
		(0.001)		(0.002)
Log(Followers+1)			0.053***	– 0.041***
			(0.001)	(0.002)
Additional controls	*Yes*	*Yes*	*Yes*	*Yes*
Original work FE	*Yes*	*Yes*	*Yes*	*Yes*
Genre dummies	*Yes*	*Yes*	*Yes*	*Yes*
Month and year FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	193,160	193,160	193,160	193,160
R^2	0.119	0.107	0.103	0.120

Dependent variable: Stay = 1 if the debut in not the only publication (zero otherwise). The sample mean of Stay is 0.389

Coefficients are estimated in an OLS regression framework

Additional controls: rating dummies, status dummy, language dummies

Robust standard errors in parentheses, $*p < 0.1$, $**p < 0.05$, $***p < 0.01$

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 aligns with expectations as reviews can be categorised as the most informative type of feedback. Note that all three feedback types have similar distributions with a mean between 22 and 26 (see Table 1).

Column (4) presents estimates from a model with all three types of feedback. This means that the coefficients measure the contribution from one kind of feedback on explaining variations in the probability of producing more than one text, holding the other two fixed. We find the estimated coefficient of reviews to be almost three times as high as that for 'favoriting'. Furthermore, in case of followers, the estimated coefficient is negative. This counter-intuitive result might be best explained by the notification feature of 'favoriting': The first publication of writers who stopped publishing thereafter may still attract users who hope to be alerted in the event of an update (which never happens).[8]

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

[7] As 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 staying in the community for first-time writers is 0.072 (SE 0.001).

[8] We thank an anonymous referee for pointing this out.

because the unobserved writing talent could work as a confounder. Having said this, the fact that prior research has identified similar effects in randomised field experiments suggests that community feedback coupled with the feeling of inclusion prevents novices from leaving.

4.2 The effect of community feedback on writing times

Next, we investigate the effect of community feedback on the time between publications. Since feedback now relates to the previous 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

$$\begin{aligned}\text{WritingTime}_{i,j,t} = \beta_0 + \beta_1 \ln\text{Feedback}_{i-1,j,t} + \beta_2 \text{ContrCount}_{i,j,t} \\ + \beta_3 \text{ContrCount}_{i,j,t} \cdot \ln\text{Feedback}_{i-1,j,t} + \xi_i + \phi_j + \theta_t + \varepsilon_{i,j,t}.\end{aligned} \tag{2}$$

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 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 substantial effect on writing times: a 1% increase in reviews for the previous text increases the time until the following publication by 22–27 days (columns (1) and (2)). The mean time between publications is 127 days, which means an increase of around 18–21%. Compared to reviews, the effect for the two types of low-information 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 put more effort into 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, who then doubt their creative output more and take more time for publication. However, this effect diminishes when experience increases as $\hat{\beta}_3$ is negative and significantly different from zero.[9] More experienced writers, on average, need less time to produce output. In line with Goes et al. (2014), the habituation effect might best explain this finding.

[9] Note 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.

Table 3 Effect of feedback on writing times

	Time between publications, days					
	Reviews		Favoriting		Followers	
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Reviews$_{i-1}$+1)	22.334***	26.694***				
	(0.585)	(0.859)				
Log(Reviews$_{i-1}$+1)		− 0.418***				
· ContrCount		(0.075)				
Log(Favoriting$_{i-1}$+1)			10.016***	12.723***		
			(0.605)	(0.975)		
Log(Favoriting$_{i-1}$+1)				− 0.226***		
· ContrCount				(0.069)		
Log(Followers$_{i-1}$+1)					10.739***	13.695***
					(0.534)	(1.087)
Log(Followers$_{i-1}$+1)						− 0.275***
· ContrCount						(0.093)
ContrCount	− 2.267***	− 1.574***	− 2.263***	− 1.830***	− 2.264***	− 1.964***
	(0.371)	(0.278)	(0.375)	(0.346)	(0.374)	(0.336)
Author FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Additional controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Original work FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Genre dummies	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Month and year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Constant	84.155***	75.856***	113.171***	107.697***	117.600***	113.648***
	(4.800)	(4.104)	(4.712)	(4.625)	(4.708)	(4.414)
Observations	224,427	224,427	224,427	224,427	224,427	224,427
R^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 an 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)

4.3 The effect of community feedback on text lengths

Finally, to estimate the effect of community feedback on the extent of UGC, we regress the text length on the (average) number of past reviews, followers, and 'favoriting'. Again, debuts were discarded. The model is very similar to the preceding model defined by Eq. (2):

Table 4 Effect of feedback on text lengths

	Text length of the publications, log words				Change of the text length	
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln\text{Reviews}_{i-1}$	0.315***	0.274***	− 0.060***	− 0.066***	− 0.095***	− 0.099***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)
ContrCount	− 0.003***	− 0.003***	0.001***	0.000	− 0.000	− 0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\ln\text{Reviews}_{i-1}$				0.001***		0.000***
· ContrCount				(0.000)		(0.000)
Author FE	*No*	*No*	*Yes*	*Yes*	*Yes*	*Yes*
Additional controls	*No*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Original work FE	*No*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Genre FE	*No*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Month and year FE	*No*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
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, $\text{Words}_i/\text{Words}_{i-1}$ is the relation of text i's length to the length of the previous text. ContrCount: count of contributions

Coefficients are estimated in an 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, and writer's 'fanfiction age' (i.e. the time elapsed since registration)

The number of observations varies due to missing values

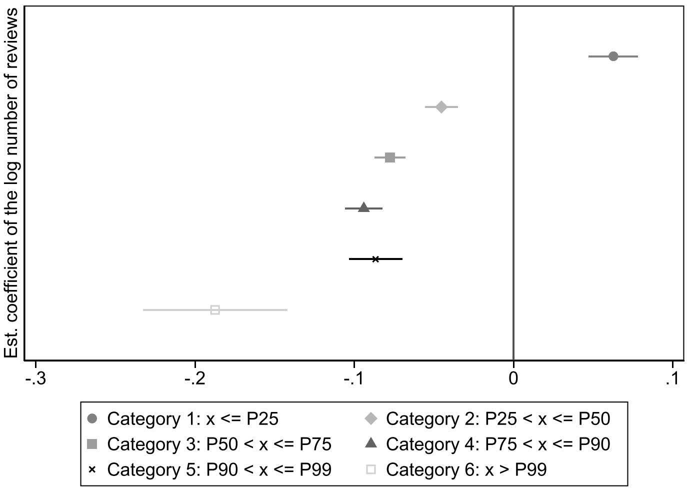

Fig. 3 Effect of reviews by readers on text lengths for different categories of writers. *Notes*: x = average number of words per text of writer j until text i

$$\begin{aligned}\ln\text{Words}_{i,j,t} = \gamma_0 + \gamma_1 \ln\text{Feedback}_{i-1,j,t} + \gamma_2 \text{ContrCount}_{i,j,t} \\ + \gamma_3 \text{ContrCount}_{i,j,t} \cdot \ln\text{Feedback}_{i-1,j,t} + \xi_i + \phi_j + \theta_t + \varepsilon_{i,j,t}.\end{aligned} \tag{3}$$

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

Table 4 shows the results. Surprisingly, after the inclusion of author fixed effects, the estimated γ_1 changes sign (columns (3) and (4)). That is, the estimated coefficient of the log number of reviews switches from positive to negative. As an explanation, we refer to the unbalanced structure of content provision by users on the extensive and 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 varies in sign and magnitude: the producers of short pieces respond to a rise in reviews with increased text lengths, whereas the opposite is true for the producers of long works.

In other words, once we control for within-author variations, community feedback, on balance, decreases the extent 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 6 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 preceding text as the dependent variable (columns (5) and (6)). Results for this relative measure of text length do not differ qualitatively. Including the time between text i and *i-1* or the log of the average writing time for writer j up to text i hardly change any of the estimates.[10]

In light of the considerations raised in Sect. 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 Influence of new material coming from the source

The preceding analyses have helped to understand how peer feedback affects user-generated content production in the online community under consideration. In this part of the analysis, we assess how the release of new original material 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 7 in Appendix for details).

[10] Results can be made available upon request.

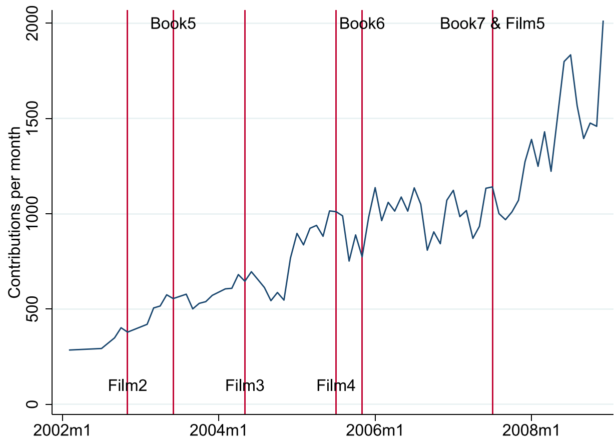

Fig. 4 Contributions to *Harry Potter*

5.1 Do new books or movies increase the number of contributions?

In the first approach, we examine whether new material impacts the extent of fan text releases. Precisely, 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 between two publications is 27 days and fans need time to read first, we use seven binary variables indicating whether a new original book (or movie) was released 3 months before or after the publication month or whether it was published the same month. This procedure also ensures that further inputs from the original work will not bias the estimates as we define a minimum interval of 6 months between the actual new release and subsequent releases of new content. Differences in popularity that may cause a selection bias are captured by the original work fixed effects. 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 8 in Appendix shows that the model does a poor job of explaining variations in fan text publications according to the release of new material. We conclude that 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*, Fig. 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's

Table 5 New material and writing times—Cox regressions

	Time between publications, days				
	All writers			Top writers	
	(1)	(2)	(3)	(4)	(5)
New book	– 0.787***	– 0.697***	– 0.661***	– 0.879***	– 0.832***
	(0.009)	(0.014)	(0.013)	(0.050)	(0.053)
New movie			– 0.695***	– 1.078***	– 0.963***
			(0.014)	(0.041)	(0.045)
Additional controls	*No*	*Yes*	*Yes*	*No*	*Yes*
Original work FE	*No*	*Yes*	*Yes*	*No*	*Yes*
Genre FE	*No*	*Yes*	*Yes*	*No*	*Yes*
Rating FE	*No*	*Yes*	*Yes*	*No*	*Yes*
English dummy	*No*	*Yes*	*Yes*	*No*	*Yes*
Year FE	*No*	*Yes*	*Yes*	*No*	*Yes*
N	136,635	136,635	136,635	11,941	11,941

Dependent variable: WritingTime is the time elapsed since the last contribution

Additional controls: Log(Reviews$_{i-1}$+1), *ContrCount*, 'FanFiction' age, text status (finished/unfinished), author-specific average writing time before text i

Top writers: Writers in the 75th percentile of prior publications (14) and 'Followers' (308)

Robust standard errors in parentheses (clustered on the writer level), $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

heterogeneity and the complexity of the writing process that a monthly measure cannot pin down. We address these issues and focus on the writing time next, which is the period between two releases.

Specifically, we estimate a Cox proportional hazards model defined by

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

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

Since our investigations in the previous section suggest that writing activities are affected by prior community feedback and a writer's experience, we include the log of reviews related to the previous text, a writer's 'platform age' and the number of prior publications as control variables. Further dummy variables control for text characteristics (genre, category, finished/unfinished status, language) and general heterogeneity across the original works. Month and year dummies capture seasonal and general time trends. Finally, we use the

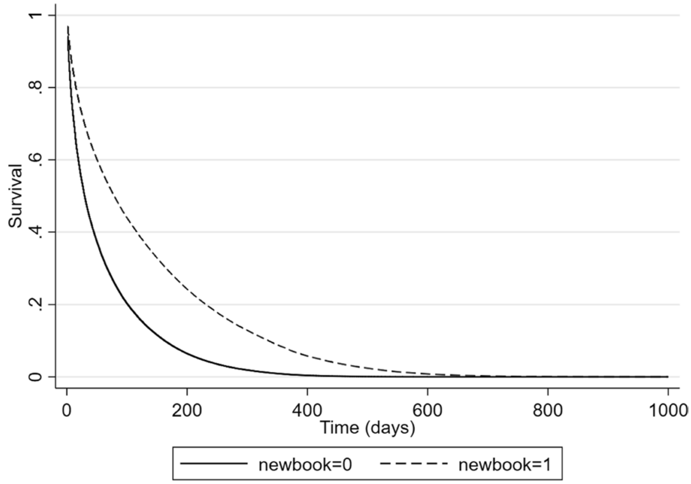

Fig. 5 Cox proportional hazards regression

author-specific average time before i to account for individual differences in text production.[11] The full sample is used to estimate the model.

The estimated coefficients presented in Table 5 show that new original material increases the production time for new content. For instance, the point estimate of −0.70 for new books (column (2)) suggests that they decrease the likelihood of a new release by 50.34% (which is $100 * (\exp(-0.70) - 1)$). 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 ($p_{75} = 14$) *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 a 56.40% reduction in the likelihood of publication (column (5)). This finding suggests that instead of quickly responding to new content to get ahead of the herd, gain the privilege of interpretation, and 'stay in the game', top writers take their time to adequately 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 there is no statistical difference in the number of 'favoriting' per text for the elite of top writers after the release of new material (means of 40.95 vs. 39.31, p-value = 0.842), a test on the equality of means just barely misses statistical significance at the 10% level for the average writer (means of 31.81 and 28.53, p-value = 0.130).

[11] Controlling for text length would help to make reactions in writing times to new material more comparable. Our model does not include the text length due to endogeneity issues. However, estimates of the full model with and without text length are virtually identical.

6 Concluding remarks

In this paper, we explore possible stimuli for 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 original series's new material and expansion may serve as a booster of content provision. In other UGC settings, separating the shock from the subject and other stimuli is much more challenging. At the same time, community feedback, which is mostly positive and encouraging, can serve as a driver for further content provision and improve the quality of the content. Furthermore, the audience's attention, especially towards the fanfiction for popular originals, can be viewed as a source of competition, improving productivity and creative process (as suggested in Gross, 2020; Wu & Zhu, 2022). It is also widely accepted that the fanfiction community is not just the marginalised group of teenagers but is among the drivers of modern popular culture.

Our estimation results suggest that community feedback stimulates the writers to produce more texts and continue creating fan writings after the debut. Moreover, other things equal, the authors receiving more feedback take more time to produce the following text and shorter ones. Short text producers, however, increase the text length after receiving more feedback. This is likely explained by the internal quality control, trying to deliver well-shaped characters in a more concise text in response to the audience's attention. These findings align with other literature on peer feedback in the UGC provision, for example, Burtch et al. (2022).

The novel finding concerns the stimuli coming from the original text. We find that conditional on text length, quasi-exogenous shocks from the release of new material increases the production time for new content. Moreover, the result is even more pronounced for the most popular writers. We interpret this to mean that top writers prefer to develop their ideas properly, 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 does not necessarily give 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, our findings can be transferred to competitive creative work settings beyond UGC. For instance, sharing the first results with peers to receive a 'friendly review' works similarly in science. So, 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 practised) rules for feedback and general conduct.

Appendix A

See the Figs. 6, 7, 8 and 9.

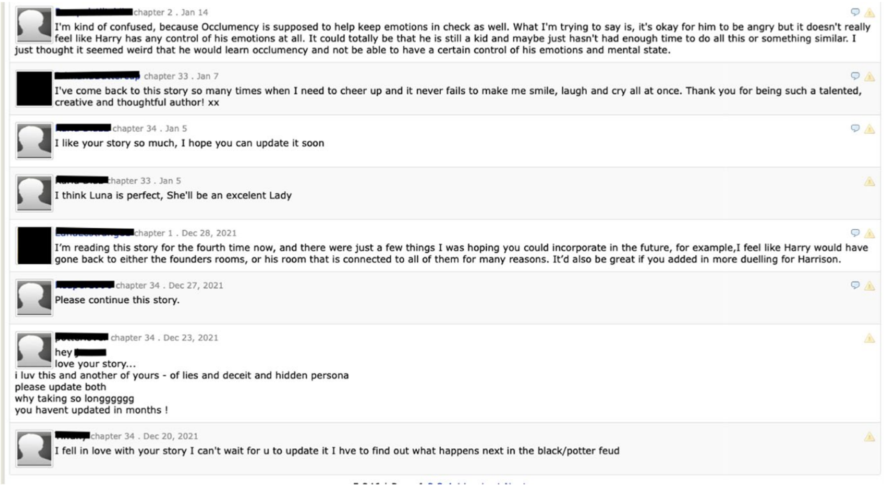
[redacted] chapter 2 . Jan 14
I'm kind of confused, because Occlumency is supposed to help keep emotions in check as well. What I'm trying to say is, it's okay for him to be angry but it doesn't really feel like Harry has any control of his emotions at all. It could totally be that he is still a kid and maybe just hasn't had enough time to do all this or something similar. I just thought it seemed weird that he would learn occlumency and not be able to have a certain control of his emotions and mental state.

[redacted] chapter 33 . Jan 7
I've come back to this story so many times when I need to cheer up and it never fails to make me smile, laugh and cry all at once. Thank you for being such a talented, creative and thoughtful author! xx

[redacted] chapter 34 . Jan 5
I like your story so much, I hope you can update it soon

[redacted] chapter 33 . Jan 5
I think Luna is perfect, She'll be an excelent Lady

[redacted] chapter 1 . Dec 28, 2021
I'm reading this story for the fourth time now, and there were just a few things I was hoping you could incorporate in the future, for example,I feel like Harry would have gone back to either the founders rooms, or his room that is connected to all of them for many reasons. It'd also be great if you added in more duelling for Harrison.

[redacted] chapter 34 . Dec 27, 2021
Please continue this story.

[redacted] chapter 34 . Dec 23, 2021
hey [redacted]
love your story...
i luv this and another of yours - of lies and deceit and hidden persona
please update both
why taking so longggggg
you havent updated in months !

[redacted] chapter 34 . Dec 20, 2021
I fell in love with your story I can't wait for u to update it I hve to find out what happens next in the black/potter feud

Fig. 6 Reviews example 1

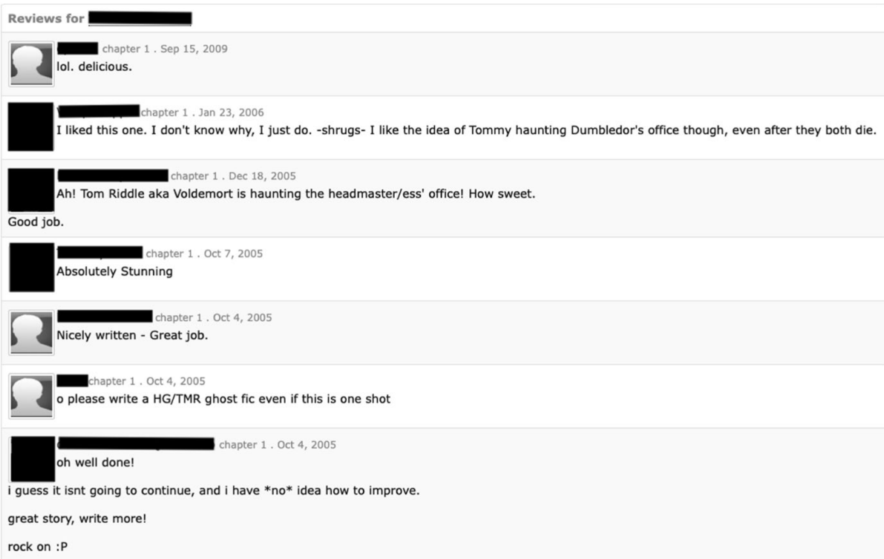
Reviews for [redacted]

[redacted] chapter 1 . Sep 15, 2009
lol. delicious.

[redacted] chapter 1 . Jan 23, 2006
I liked this one. I don't know why, I just do. -shrugs- I like the idea of Tommy haunting Dumbledor's office though, even after they both die.

[redacted] chapter 1 . Dec 18, 2005
Ah! Tom Riddle aka Voldemort is haunting the headmaster/ess' office! How sweet.
Good job.

[redacted] chapter 1 . Oct 7, 2005
Absolutely Stunning

[redacted] chapter 1 . Oct 4, 2005
Nicely written - Great job.

[redacted] chapter 1 . Oct 4, 2005
o please write a HG/TMR ghost fic even if this is one shot

[redacted] chapter 1 . Oct 4, 2005
oh well done!
i guess it isnt going to continue, and i have *no* idea how to improve.
great story, write more!
rock on :P

Fig. 7 Reviews example 2

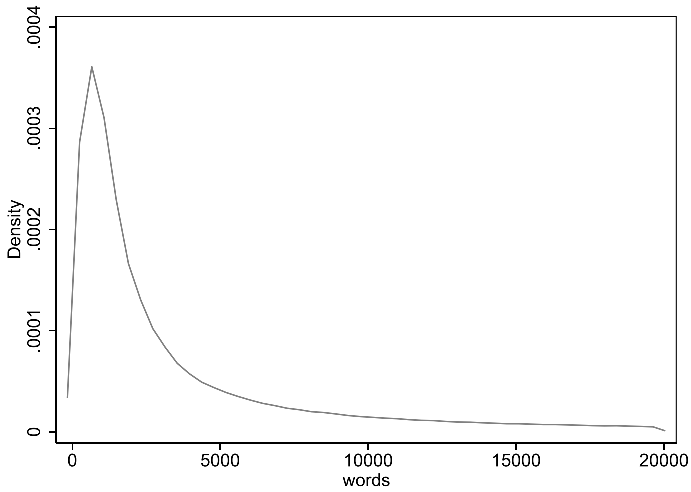

Fig. 8 Number of words (excluding strong outliers)—kernel density estimates. *Notes:* Sample restricted to observations below the 90th percentile. $N = 434,553$

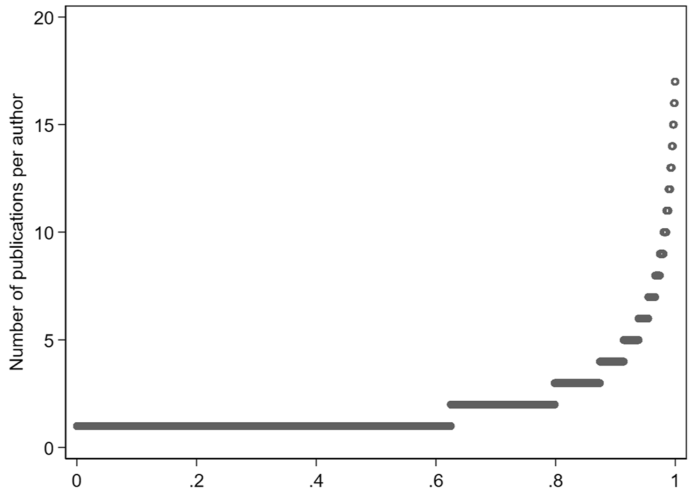

Fig. 9 Number of fan texts per author(excluding strong outliers)—quantile plot. *Notes*: Sample restricted to observations below the 99th percentile. $N = 203,217$

Appendix B

See the Tables 6, 7 and 8.

Table 6 Effect of feedback on text lengths: Favoriting and Followers

	Favoriting			Followers		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln\text{Favoriting}_{i-1}$	0.312***	− 0.044***	− 0.052***			
	(0.004)	(0.003)	(0.003)			
ContrCount	− 0.004***	0.001***	− 0.001	− 0.003***	0.001***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\ln\text{Favoriting}_{i-1}$			0.001***			
· ContrCount			(0.000)			
$\ln\text{Followers}_{i-1}$				0.313***	− 0.053***	− 0.060***
				(0.004)	(0.003)	(0.003)
$\ln\text{Followers}_{i-1}$						0.001***
· ContrCount						(0.000)
Author FE	*No*	*Yes*	*Yes*	*No*	*Yes*	*Yes*
Additional controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Original work FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Genre FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Month–year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
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, $\text{Words}_i/\text{Words}_{i-1}$ is the relation of text i's length to the length of the preceding text. ContrCount: count of contributions

Coefficients are estimated in an 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 7 List of book series used in Sect. 5

Rank	Title	No. fan texts
1	Harry Potter	80,827
2	Percy Jackson and the Olympians	65,287
3	Hunger Games	37,744
4	Mortal Instruments	13,373
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

Books are ranked according to the number of fan texts

Table 8 Amount of fan text publications in response to new material

	(1)	(2)	(3)	(4)
New book	− 11.562	− 11.568	11.051	13.611
	(10.594)	(10.629)	(14.135)	(20.678)
New film		14.106	13.255	13.083
		(45.498)	(41.148)	(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	*No*	*No*	*Yes*	*Yes*
Month dummies	*No*	*No*	*Yes*	*Yes*
Original work FE	*Yes*	*Yes*	*Yes*	*Yes*
N	4400	4400	4053	4053
R^2	0.000	0.000	0.048	0.050

Dependent variable: Fan contributions per month and topic (original work)

Coefficients are estimated in an OLS regression framework (robust standard errors in parentheses)

Acknowledgements We are grateful for the valuable comments provided by Sophie Einwächter, Karol Borowiecki, Samuel Cameron, Anne-Sophie Radermecker, Nicholas Martin Ford, Francesco Angelini, an anonymous referee, 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.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest All authors have no Conflict of interest.

References

Aaltonen, A., & Seiler, S. (2016). Cumulative growth in user-generated content production: Evidence from Wikipedia. *Management Science, 62*(7), 2054–2069.

Burtch, G., He, Q., Hong, Y., & Lee, D. (2022). How do peer awards motivate creative content? Experimental evidence from Reddit. *Management Science, 68*(5), 3175–3973.

Burtch, G., Hong, Y., Bapna, R., & Griskevicius, V. (2018). Stimulating online reviews by combining financial incentives and social norms. *Management Science, 64*(5), 2065–2082.

Cabral, L., & 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., & 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.

DeVaro, J., Kim, J.-H., Wagman, L., & Wolff, R. (2018). Motivation and performance of user-contributors: Evidence from a CQA forum. *Information Economics and Policy, 42*, 56–65.

Evans, S., Davis, K., Evans, A., Campbell, J. A., Randall, D. P., Yin, K., & Aragon, C. (2017). More than peer production: Fanfiction communities as sites of distributed mentoring. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing* (pp. 259–272).

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., & Lin, M. (2016). Do incentive hierarchies induce user effort? Evidence from an online knowledge exchange. *Information Systems Research, 27*(3), 497–516.

Goes, P. B., Lin, M., & Au Yeung, C.-M. (2014). "Popularity effect'' in user-generated content: Evidence from online product reviews. *Information Systems Research, 25*(2), 222–238.

Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature, 57*(1), 3–43.

Graddy, K., & 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., & 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.

Jenkins, S. Van Kerm, P. (2015). PARETOFIT: stata module to fit a Type 1 Pareto distribution. http://EconPapers.repec.org/RePEc:boc:bocode:s456832. Accessed 25 Mar 2024

Khern-am-nuai, W., Kannan, K., & 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. *Houston Law Review, 52*, 425.

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., & Hofstetter, R. (2013). Social ties and user-generated content: Evidence from an online social network. *Management Science, 59*(6), 1425–1443.

Sun, M., & 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. & 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.

Wang, J., Li, G., & Hui, K.-L. (2022). Monetary incentives and knowledge spillover: Evidence from a natural experiment. *Management Science, 68*(5), 3549–3572.

Wu, Y., & Zhu, F. (2022). Competition, contracts, and creativity: Evidence from novel writing in a platform market. *Management Science, 68*(12), 8613–8634.

Xu, L., Nian, T., & Cabral, L. (2020). What makes geeks tick? A study of stack overflow careers. *Management Science, 66*(2), 587–604.

Yin, K., Aragon, C., Evans, S., & 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* (pp. 6106–6110). ACM.

Zhang, X. M., & Zhu, F. (2011). Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review, 101*(4), 1601–15.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.