Title

Essays on Platform Design Choices for Reward Systems and Tipping in Online Communities

Author

Paridar, Mahsa

Publication Date

2025

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays on Platform Design Choices for Reward Systems and Tipping in Online Communities

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy in Management

by

Mahsa Paridar

2025

ABSTRACT OF THE DISSERTATION

Essays on Platform Design Choices for Reward Systems and Tipping
in Online Communities

by

Mahsa Paridar

Doctor of Philosophy in Management

University of California, Los Angeles, 2025

Professor Elisabeth Honka, Co-Chair

Professor Mina Ameri, Co-Chair

Online platforms depend heavily on sustained user engagement and high-quality content to generate revenue through advertising impressions and by receiving a portion of tips, where users financially reward content creators directly, providing platforms with an additional revenue stream. To achieve these goals, platform owners implement strategic mechanisms that incentivize continuous participation and content production. Two crucial mechanisms in this regard are peer-driven financial support (tipping) and platform-driven incentive systems, both significantly impacting user behavior, content quality, and platform profitability. Understanding these mechanisms is essential for platform owners to optimize engagement strategies, maximize advertising revenue, and effectively leverage direct financial interactions.

In the first essay, I examine tipping behavior and develop a model wherein users determine tip amounts based on their beliefs about an evolving tipping norm, as well as content quality, personal characteristics, and contextual factors. These beliefs are derived from two primary signals: the tips that I receive directly and those I observe others giving to similar content.

A novel aspect of my model allows for the correlation of these signals within a type, across different types, and over time. Through Bayesian updating, users assimilate these signals into their perception of the prevailing tipping norm. My findings reveal that both signals significantly influence user behavior, with tips received playing a more decisive role. I further show that tip amounts are primarily driven by the inferred tipping norm, followed by the quality of the content and individual user characteristics. Prediction exercises suggest that strategic information disclosure on the platform can significantly influence tipping behavior even in later stages.

In the second essay, I quantify the effects of peer and platform rewards on the quantity and quality of user-generated content. My analysis indicates that monetary rewards from peers, such as tips, robustly increase the frequency and length of user posts, while monetary rewards from the platform promote longer submissions at the cost of posting frequency. Moreover, although individual non-monetary rewards (likes) exert a modest influence, their cumulative volume markedly enhances content production. I also find that non-monetary platform rewards (badges) exhibit a nonlinear effect, with content generation declining upon reaching a milestone and subsequently rising as users approach the next target. These findings offer valuable guidance for designing effective reward systems that encmy age desired sustainable user engagement.

The dissertation of Mahsa Paridar is approved.

Randolph E. Bucklin

Brett William Hollenbeck

Elisabeth Honka, Committee Co-Chair

Mina Ameri, Committee Co-Chair

University of California, Los Angeles

2025

To my beloved mother, whose strength, wisdom, and boundless love shaped the core of who I am. I was truly blessed to share 34 precious years with her, each moment filled with warmth, laughter, and deep understanding. The bond I shared defies description, and I would never exchange a single second of it. She enriched my soul in countless ways, and if there is anything in me that appears uniquely kind and strikingly special, it is her beautiful spirit shining through me. Her memory lives vividly within me, forever guiding and inspiring my path.

To Sina, my biggest supporter, my best friend, my love, my husband, and my constant companion through every twist and turn. His quiet strength, deep understanding, and endless patience carried me through the hardest days and made the good ones even brighter. He believed in me when I doubted myself, held me when I needed comfort, and brought lightness into even the heaviest moments. This jmy ney became bearable, even beautiful, because of him.

To my father, whose constant support, understanding, and unconditional love have shaped my life profoundly. His guidance, wisdom, and unwavering belief in my abilities have encmy aged me through every obstacle. His example has taught me resilience, compassion, and integrity, continually motivating me to pursue my dreams and achieve my goals.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This dissertation would not have been possible without the guidance and support of my advisors, Drs. Elisabeth Honka and Mina Ameri. I am also grateful to my committee members, Drs. Randolph E. Bucklin and Brett Hollenbeck, for their valuable feedback, encmy agement, and continued support throughout this process.

April 2025

VITA

2008–2012	B.S. Industrial Engineering, Iran University of Science and Technology
2012–2015	M.B.A. Management, Sharif University of Technology
2016–2018	Co-founder and Marketing Manager, Adshot Co. (Insight)
2019–2025	Anderson Fellowship, UCLA Anderson School of Management
2020–2025	Teaching and Research Assistant, UCLA Anderson School of Management

Chapter 1

Online Tipping under an Evolving Social Norm

1.1 Introduction

Many social platforms have launched tipping features on their websites in recent years. For instance, YouTube announced the "Super Thanks" feature in July 2021, TikTok launched "Tip Jar" in October 2021, and Instagram introduced "Gifts" for Reels in November 2022.[1] The introduction of a tipping feature highlights the growing recognition of the importance of financial support from peers in sustaining the efforts of content creators. An immediate question that a platform faces when launching such a new feature is how to design the tipping environment on the platform. Should the platform publicly display who tipped which amount to a piece of content? Or should it keep such information private? Or should the platform reveal some but not all information, e.g., by showing the average tip amount? In this paper, I examine the effects of different tipping information disclosure strategies on total tip amounts, the number of tippers, and other tipping-related outcomes and discuss their implications for platform design.

[1] Other platforms such as Twitter, Clubhouse, and Twitch have also introduced tipping features, reflecting a broader trend towards enabling direct financial support for content creators within online communities.

Tipping decisions are generally driven by multiple factors (e.g., gratitude, generosity, quality, etc.). One important driver of tipping decisions are social norms (Akerlof 1980; Bernheim 1994; Azar 2004). Social norms are the widespread convergence or the unplanned, unexpected result of individuals' interactions that determine what is/is not acceptable in a group or community (Muldoon et al. 2013). These norms are important as they provide order, predictability, and harmony in any social group by creating an expected idea of how one should behave (Young 1993). In the context of online tipping, where there is relatively little precedence, these norms are emergent properties, arising from individuals' actions and decisions. While the effects of quality and (established) social norms on tipping behavior have been documented (Azar 2007, 2020), there is a lack of empirical research investigating how an *evolving* social norm impacts users' tipping and how individuals' tipping decisions influence the development of a social norm. Examining the new practice of online tipping and the factors impacting it provides insights into how norms form and evolve in digital communities.

Social norms typically develop through repeated interactions and learning (Young 1993). In the context of online tipping, users "interact" by producing and consuming content and "learn" by observing others' tipping decisions. Because social platforms govern which information users can see about others' tips, e.g., who tipped what, when, and how much, they can influence the evolvement of a tipping norm. Using prediction exercises, I first investigate how different information provision strategies affect the development of a tipping norm and users' tipping decisions. I then examine how "sticky" a tipping norm is, i.e., can social platforms still significantly change a tipping norm and users' tipping decisions in later stages? Or is a platform forever "stuck" with the tipping norm that arose based on the platform's initial information disclosure decisions?

I use data from an online board game platform called BoardGameGeek.com (BGG). BGG is a special interest online community where individuals who are interested in board games can learn about them and interact with other board game fans.[2] More importantly, because users provide all the content on this platform, they can act as content creators, generating valuable information and reviews about board games as well as entertaining content.[3] The platform also has its own currency, and starting May 13th, 2005 allowed users to tip content creators using this currency. BGG is an ideal environment to study online tipping because all users' interactions and tipping behaviors, especially after the tipping feature was first introduced, are observable. I study users' decisions for 22 months following the introduction of tipping. During this time, users who gave tips, on average, gave 5.73 tips and users who received tips, on average, received 15.60 tips annually. My data show that the standard deviation of tip amounts decreased over time, suggesting that users tipped more similar amounts as time progressed.

In my model, social norms are incorporated as users' perceptions of the tipping norm, which are continuously updated via Bayesian updating. These perceptions are driven by signals from two smy ces: users' self-experience of tips they received (Young 2015) and observed tipping behavior in the BGG community (Schuster, Kubacki, and Rundle-Thiele 2016). A new feature of the model is that I allow the signals to be correlated when deriving the posterior distribution. More specifically, signals can be correlated within a smy ce and a time period, across smy ces and within a time period, and across time. The perceived norm, along with characteristics of the focal content and a user's personal tendency to tip, govern the user's tip decision.[4] The model is estimated using a Tobit framework.

My results show that users learn about the current tipping norm through both their own self-experience of receiving tips and observed tipping behavior on the platform. On a per-tip

[2]Consumers increasingly prefer special interest online communities over (general) social media, e.g., there are over 2.2 million subreddits and more than 10 million Facebook groups (https://www.amity.co/blog/40-statistics-you-should-know-about-online-communities). The number of members in special interest online communities has increased by 81% since 2019. Examples of other prominent special interest online communities are goodreads.com, cyclechat.net or soundcloud.com.

[3]In terms of the ratio of content consumers create and consume, BGG is similar to other online forums such as reddit.com, stackoverflow.com or stackexchange.com.

[4]While reciprocity has been shown to be another driver of tipping decisions, my data does not suggest that reciprocity plays a role in this empirical context (see Section 2.4 for a detailed discussion).

basis, users find the tips they receive themselves to be more informative than tips observed in the community in shaping their perception of the tipping norm. However, because of the much larger number of tips users observe in the community than receive themselves, the total effect of tips observed in the community on the perceived norm is larger than the total effects of tips received. Furthermore, I separate the portions of the utility that come from content quality, individuals' beliefs about the norm, and individual characteristics (via user fixed effects). I show that users' beliefs about the norm, on average, represents 67% of a tip given on the platform followed by content quality and individual characteristics with 28% and 5%, respectively.

Next, I examine how information disclosure affects users' tipping behavior. I do so by implementing three prediction scenarios: in the first one, users update their perception about the tipping norm only based on personal experience, i.e., they cannot see the tips given in the broader community; in the second one, users update their perceptions about the tipping norm only based on the tips given in the broader community, i.e., they cannot see the tips they receive; and in the third one, users update their perception about the tipping norm based on complete information about tips received personally, but only partial information about tips in the broader community, i.e., they observe the average tip given by the broader community. My results show that information disclosure (visibility) of personal signals has little impact on tipping behavior, but information disclosure of community signals significantly affects users' tipping decisions. When tips in the broader community are *not* visible, users tip smaller amounts but much more often, increasing the total tip amount by 39%. These findings suggest that platforms can strategically manage tip visibility to increase overall tipping activity. However, signal invisibility also leads to larger uncertainty about the tipping norm, highlighting the multifaceted effects of different information disclosure strategies.

And lastly, I study how sticky the perceived tipping norm and tipping behavior are. I do so by comparing outcomes between the same three information disclosure scenarios discussed in the previous paragraph but introduced in the second half of the study period only and the scenario when users observe all signals throughout the whole study period, my main model. My predictions show that the perceived tipping norm and tipping behavior are quite sticky even in the medium-run, i.e., nine months after the change in information disclosure. This is especially the case for aspects of tipping that speak to the breadth of this behavior: the number of unique tippers, the number of unique tippees, and the number of unique tipped content. For example, if a platform removes the visibility of community signals after the first half of the study period, the number of unique tippers, the number of unique tippees, and the number of unique tipped content are smaller by −5.12%, −22.01%, and −14.59%, respectively, even a year after the change compared to a scenario where community signals were never visible.

The contribution of this paper is two-fold. First, I add to managers' and academics' understanding of the impact of different information disclosure strategies on the perceived tipping norm and users' tipping decisions. This is particularly relevant since digital platforms often adopt varied approaches to the visibility of such tip incidences. At one end of the spectrum, platforms such as Twitch or YouTube make tipping visible and salient on users' screens. On the other end of the spectrum, platforms such as Patreon or Cameo keep monetary contributions private between the supporter and the content creator. This study empirically evaluates the impact of these and other information disclosure strategies on tipping decisions.

Second, I show how a tipping norm as a collective of individual decisions evolves over time and how it affects individuals' tipping behavior. By investigating the impact of perceived norms along with content quality while controlling for users' intrinsic motivation, I shed light on how users decide to tip. I further show that, while users learn from self-experience and observing others' actions, these two smy ces of information play different roles in shaping users' beliefs about the norm. By examining how these factors interplay in shaping individual tipping behaviors, I provide more insight into the evolution of tipping norms in online communities, where norms are emerging and evolving.

The remainder of this paper is organized as follows: In the next section, I review the relevant literature. In Section 2.3, I describe my data. I present my model in Section 2.4 and discuss the results in Section 2.5. In the following section, I perform prediction exercises and conclude in section 2.8.

1.2 Relevant Literature

In this section, I review three streams of literature on tipping, social norms, and special interest communities and delineate the positioning of my research in relation to the findings from the extant literature.

1.2.1 Tipping

Previous literature has found three main reasons as to why people tip offline: (i) as an incentive/reward for higher-quality service (Azar 2007; Lynn and Sturman 2010), (ii) because of psychological reasons, e.g., gratitude, social reputation (Conlin, Lynn, and O'Donoghue 2003; Lynn 2014), and (iii) to adhere to social norms (Azar 2010). Furthermore, previous research has also found that default options affect people's tipping decisions (e.g., Haggag and Paci 2014; Everett et al. 2015).

Few papers have investigated digital tipping. Using data from a field experiment on Uber, Chandar et al. (2019) find that tipper characteristics explain much more of the observed variation in tipping than tippee characteristics. Similarly, in the context of an online freelance marketplace, Kim, Amir, and Wilbur (2023) show that tipping decisions are largely driven by tipper characteristics, such as geography and satisfaction. The authors demonstrate that an injunctive norm message significantly increases tipping rates among new buyers, while reciprocity-related messages have no significant impact. Lu et al. (2021) investigate the relationship between audience size and tip revenue of live streamers. They find that a larger audience amplifies social image benefits, thereby increasing both the number of viewers and the revenue from tips for live streamers.

Similar to the before mentioned three papers, I also study digital tipping. However, I develop a micro-founded model that incorporates the main drivers of online tipping decisions and quantifies their influence. Further, my model allows for the development of a social norm related to tipping and for users to be affected by it.

1.2.2 Social Norms

Social norms are the unwritten codes and informal understandings that define what others expect of us and what I expect of others (Young 2015), as well as the unplanned result of individuals' interactions that determine what is/is not acceptable in a group or community (Bicchieri, Muldoon, and Sontuoso 2011). Three aspects are important in the evolution of social norms: (i) they are the result of repeated interactions, (ii) they evolve through learning, and (iii) they underpin social order (Young 1993).

There is a vast amount of literature in different fields, such as marketing, economics, psychology, health, and the environment, investigating the effects of social norms on behavior. In marketing, researchers have examined how social norms influence different types of consumer behavior, e.g., the reuse of hotel towels (Goldstein, Cialdini, and Griskevicius 2008, Chen et al. 2010), loyalty (Lee, Murphy, and Neale 2009), and responses to new products (Homburg, Wieseke, and Kuehnl 2010).[5] One of the few papers studying the effects of social norms online is Burtch et al. (2018). The authors run an experiment to infer the effects of financial incentives and social norms on online reviews. Burtch et al. (2018) find that monetary rewards increase the number of reviews, while social norms increase reviews' length, and combining the two yields the greatest benefit.

While researchers have studied the effects of social norms, few papers have investigated how social norms evolve. The papers that have studied social norm development mostly

[5] See Melnyk, Carrillat, and Melnyk (2022) for a meta-analysis of the effects of social norms on consumer behavior.

use a game-theoretic or computational approach (e.g., Young 1993; Sen and Airiau 2007; Epstein 2001). To the best of my knowledge, there are only two papers that have studied aspects of social norm development empirically. Garrod and Doherty (1993) analyze the effects of interacting with peers as opposed to isolated individuals on the speed of social norm development. Schuster, Kubacki, and Rundle-Thiele (2016) show that increasing the visibility of a target behavior can change the perceived social norm related to the behavior.

My paper belongs to the small group of papers studying social norm development empirically. In contrast to the two previously mentioned papers, I explicitly model the perceived social norm at each point in time, how individuals' actions affect it, and how it affects individuals' actions.

1.2.3 Special Interest Communities

Lastly, my paper is also related to the literature on special interest communities. Previous research has investigated different aspects of online communities. For example, Hendricks and Sorensen (2009) study an online music market and find that releasing a new album causes a substantial and permanent increase in the sales of the artist's old albums. Zhang and Godes (2018a) study Goodreads.com and show that with sufficient experience, having more ties leads to better decision-making. Nevskaya and Albuquerque (2019) use data from a massive online video game platform. They find that improving reward schedules and imposing time limits leads to shorter usage sessions among players and longer product subscriptions. And lastly, Ameri, Honka, and Xie (2023a) study how strangers become friends on an anime platform. To the best of my knowledge, no empirical study has investigated the board game industry.

1.3 Data

My data come from Boardgamegeek.com, an online community revolving around board games. It was established in 2000 and has become the largest online database for board games as well as the largest online community for board game fans with over 3 million users worldwide in 2024. Figure B-1 shows the number of users joining BGG over time.

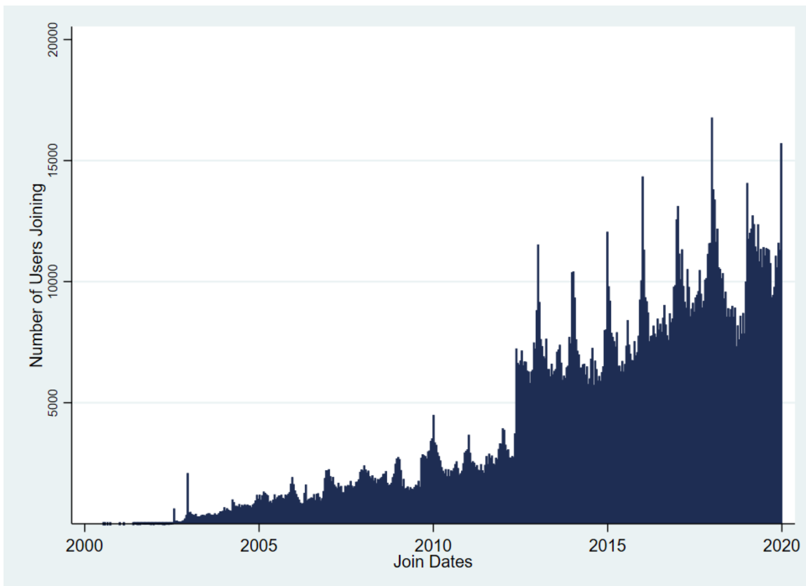

Figure 1.1: Number of Users Joining BGG Over Time

Users create all content on BGG. They provide detailed information about new and existing games via reviews, upload files and images, create their favorite board game lists ("Geeklist"), and also engage in a variety of conversations with other users in the discussion forum.

BGG utilizes a platform-specific virtual currency called GeekGold (GG) for all monetary transactions. GG cannot be directly bought GG from the platform.[6] Users can earn 1 - 5 GG as compensation for writing a review or starting a new discussion thread. Users can

[6]The platform rewards users who donate money to BGG by giving them GG. Some users may also buy GG from other users privately. However, neither donations nor GG purchases are common.

also earn GG in the form of tips from other users for the content they create. Users can tip any amount they want. Aside from tipping, users can use their GG to buy virtual cosmetic items for their profile page or to buy board games from peers. Users can also use their GG to participate in special events, such as lotteries, to win board games.

As is common in most online communities, users can react to the content produced by others not only by tipping but also by giving "likes." Figure B-2 shows a post for which the content creator received both likes and tips from other users. Users can see who tipped and the amount of each tip by clicking on the cent icon and who liked the content by clicking on the thumbs-up icon.

Figure 1.2: Example of a Post for Which the Creator Received Tips and Likes

1.3.1 Data Collection, Cleaning, (Re)Construction

BGG introduced tipping on May 13th, 2005. At that point in time, BGG had about 80,000 users. I study tipping behavior on BGG during the next 22 months ("study period")[7] and focus on users who tipped at least once during the study period. This gives us 1,785 users with 6,672 tipping incidences.[8] I drop 109 tip incidences with tip amounts of more than 20 GG.

[7]On March 27th, 2007, BGG added suggested default tip amounts.

[8]My sample also includes users who joined after March 13th, 2005, as long as they tipped at least once before March 27th, 2007.

For my sample of users, I collected all the content they created, all tips they gave, and information on other spending activities such as purchasing symbolic badges. Furthermore, I tracked all user activities that left a digital footprint on the platform, e.g., liking content, participating in a lottery, adding to board game collections, etc.

Two limitations of my data are that I do not observe user logins and the content users viewed on the platform. Since this information is not available, the following data patterns motivate and support assumptions I make: In 100% of the tipping incidences, users also engaged in at least one other activity on the platform e.g., linking content, buying a badge etc. Therefore, I focus on days on which users engaged in at least one other activity. Additionally, in 95% of the tipping incidences, users had a non-monetary reaction (like, comment, or reply) to the content they tipped. Hence, I focus on content for which users had a non-monetary reaction.

In my data, I observe that users typically tip on the same day or on days following a non-monetary reaction. Therefore, I model users' tipping decisions for content on a daily basis for up to 30 days following the non-monetary reaction (depending on the UGC type).[9]

These last two steps result in a sample of 1,785 users who engaged in 6,672 tipping incidences during the study period. My panel contains 3.9 million user-content-day observations.

1.3.2 Data Description

Figure 1.3 illustrates the number of tip incidences and average tip amount for each UGC category. Replies receive the highest average tip amount with 2.26 GG, and files receive the lowest average tip amount with 1.66 GG.

[9]The following time periods cover 90% of tipping incidences for each type of UGC: 1 day for files,10 days for threads and Geeklists, 30 days for replies, and 14 days for images.

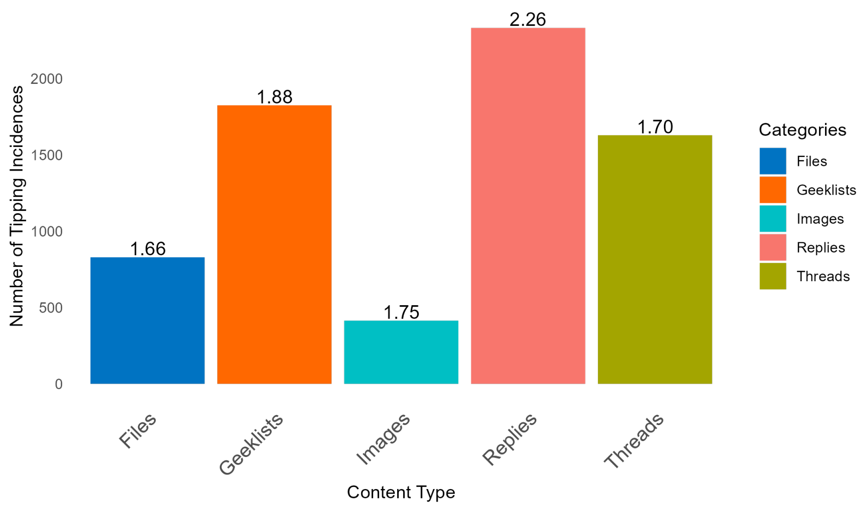

Figure 1.3: Number of Tip Incidences and Average Tip Amount by Content Category

Table 1.1 summarizes key statistics of my data. On average, the per-incidence tip amount a user gives is 1.86 GG, while the per-incidence tip amount a user receives is 2.07 GG, indicating that users tend to receive slightly higher tips than they give. Furthermore, on average, a user gives 5.73 tips and receives 15.60 tips annually, with the maximum number of tips given and received being 260 and 81, respectively. Additionally, a focal user has, on average, 43.67 GG available on any day. On average, 234.08 pieces of content are created on BGG every day.

	Mean	Std. Dev.	Min	Median	Max	N
Tip Amount Per Tip						
Avg. tip amount a focal user gives	1.86	1.92	0.01	1.00	20.00	1,785
Avg. tip amount a focal user receives	2.07	2.02	0.01	1.42	20.00	939
Tip Frequency (Annually)						
Avg. tip frequency a focal user gives	5.73	11.73	1.00	2.00	260.00	1,785
Avg. tip frequency a focal user receives	15.60	23.48	0.17	5.55	81.35	939
Daily available GG	52.51	157.82	0.00^{+}	15	3,425.75	1,785
Daily UGC production	234.08	131.89	16.00	218.00	1,330.00	683

Table 1.1: Descriptive Statistics

Figure 1.4 depicts the monthly standard deviation of tipping amounts over time. The decrease in the standard deviation of tip amounts suggests that users tip more similar amounts as time progresses.

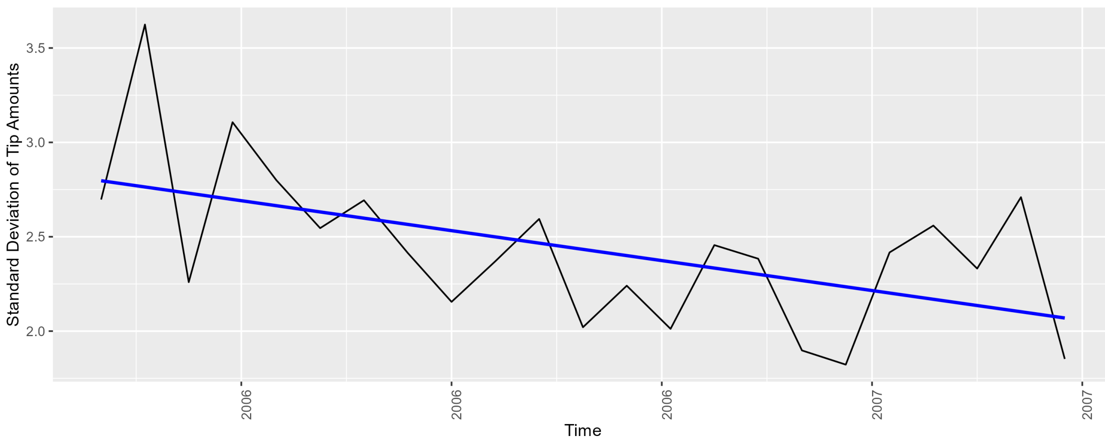

Figure 1.4: Standard Deviation of Tips Given in a Month Over Time

Figure 1.5 shows the within-user standard deviation of tip amounts over time with 95% confidence intervals for users who tipped at least twice.[10] The standard deviation of tip amounts decreases over each 6-month interval, suggesting that users tip more similar amounts over time at the individual level, similar to the pattern observed on the aggregate level.

[10] The pattern is similar for users who tipped at least three or at least fmy times.

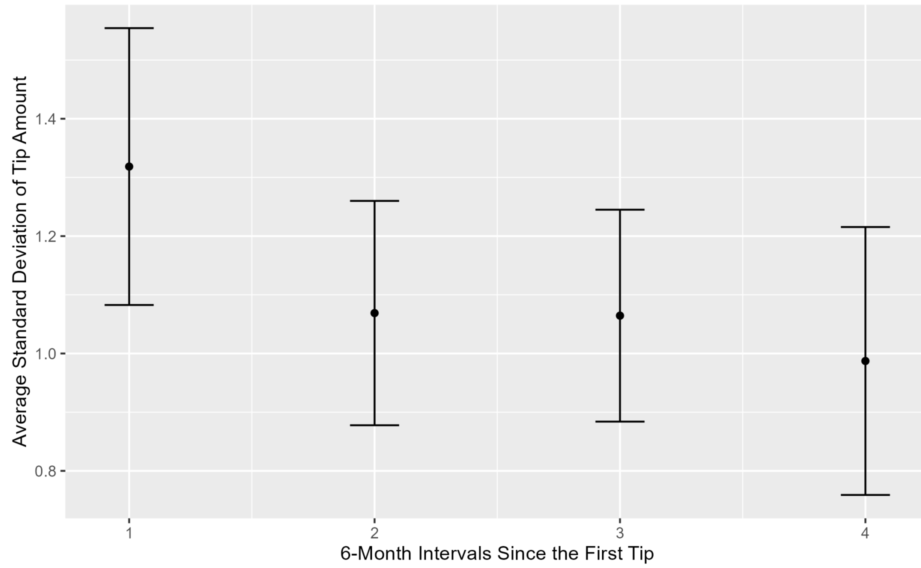

Figure 1.5: Within-User Standard Deviation of Tip Amounts Over Time

Users might have different tendencies to tip because of their nationality or culture. Table 1.2 shows the number of tip incidences, average tip amount per tip incidence, and the percentage of users in my data coming from each country.

Table 1.2: Tip Statistics By Country

	Country	Number of Tip Incidences	Average Tip Amount Per Tip Incidence	% of Users
1	United States	5,607	2.00	70.85
2	Canada	696	1.89	8.40
3	Australia	367	1.83	4.43
4	United Kingdom	270	1.99	3.26
5	Germany	242	2.22	2.92
6	Other	841	1.92	10.15

1.4 Model and Estimation

1.4.1 Model

1.4.1.1 Assumptions

I make several assumptions regarding users' tipping decisions. Users are assumed to be myopic, basing their tipping decisions on today's utility without considering future implications (Azar 2004; Lynn 2016, 2018; Azar 2020). In other words, users are not strategic in their decision of whom and how much to tip and make each tipping decision independently. In the context of tipping, being strategic might arise for two reasons: budget limitation and reciprocity. With respect to budget limitations, users may need to strategically decide which content to tip and how much to tip if they feel constrained by their available budget relative to the amount of content they consume. In other words, strategic behavior may occur when the tipping budget is limited compared to the number of consumed content items, necessitating a careful allocation of tips. However, if the budget is sufficiently large, i.e., the tip amounts are small relative to the available budget, users do not need to be strategic about their tipping decisions. Given that, in my empirical context, the average ratio of budget to tip for users is 126, I assume that the budget is not limiting for users when deciding to tip.

Reciprocity, as discussed by Fehr and Gächter (2000), suggests that individuals may tip others in return for having been tipped. In my data, only 2% of tips are given by a pair of users to each other, suggesting that reciprocity is not a major driver in my setting. Therefore, I do not include it in my model and assume that users make tipping decisions independently of each other.

Each time a user visits the platform, she can tip content she sees. As discussed in Section 3.1, because I neither observe logins nor which content users see, I make the following assumptions based on data patterns. First, my data indicate that on the days on which users tipped a piece of content, they also *always* engaged in some other form of activity. Hence, I assume that user i can only make tipping decisions on $t \in T_i$, where T_i contains the days user i engaged in any activity other than tipping on the platform, i.e., t does *not* represent calendar days. Second, my data show that in 95% of tipping incidences, users also had a non-tipping reaction to the piece of content. Therefore, I consider a user as having *seen* a piece of content if the user had a non-monetary reaction to the piece of content. And lastly, depending on the type of content, tipping happens within 1 - 30 days following the non-monetary reaction. Thus, I assume that user i makes tipping decisions for content $j \in J_{it}$, where J_{it} contains content that the user has shown a non-monetary reaction to within a certain number of days prior to t. The number of days is one for files, ten for threads and Geeklists, 14 for images, and 30 for replies.

1.4.1.2 Utility Function

Formally, user $i = 1, \ldots, M$ decides how much to tip each piece of content $j \in J_{it}$ on day $t \in T_i$. User i's utility U_{ijt} from tipping content j on day t is given by:[11]

$$U_{ijt} = \alpha_i + \beta\mu_{it} + \gamma' Q_{ijt} + \eta' C_{ijt} + \epsilon_{ijt}, \tag{1.1}$$

where α_i represents the user's intrinsic tip tendency and captures internal factors such as generosity, status-seeking, and cultural background, which have been shown to impact tipping behavior (Akerlof 1980; Bernheim 1994; Azar 2007). μ_{it} is user i's posterior belief about the tipping norm on day t (discussed in detail in the next subsection), Q_{ijt} captures content quality, C_{ijt} contains control variables, and ϵ_{ijt} is a normally distributed error term.

Naturally, users are more likely to give (higher) tips to higher-quality content to show their gratitude and to encmy age more (high-quality) content creation in the future (Azar 2007; Paridar, Ameri, and Honka 2023). Q_{ijt} contains variables capturing content quality. For textual content, I use the length of the text, operationalized as number of sentences, that

[11] This utility function is equivalent to an indirect utility function with choice of the amount of a product with price of 1; the available amount of money acts as the budget constraint bounding the solution space.

has been shown to be a good proxy for content quality (Blumenstock 2008; Demberg and Keller 2008; Hasan Dalip et al. 2009; Anderka, Stein, and Lipka 2012). For images, quality is assessed by multiplying the dimensions (width x height). I do not have quality measures for files.

C_{ijt} contains the control variables. The number of likes given to a piece of content captures its popularity. I include dummy variables for each type of content, control for user i's membership length, and user i's available GG. Since newer content might be more engaging and thus more likely to receive tips, I also control for content age in days. A user can only receive tips if she made a post in the past. Relatedly, a user who wrote multiple posts in the past is more likely to receive tips than a user who wrote one post. To account for this, I include dummy variables which indicates whether a user has ever received any tips before and control for the number of content pieces a user created in the past 7 days. To control for the overall activity level on the platform, I also include the number of content pieces created by all users on the platform on day t.

1.4.1.3 Perceived Tipping Norm and Signals

μ_{it} captures user i's belief about the mean of the tipping norm on day t. Initially, user i holds an uncertain prior belief about the tipping norm that follows a normal distribution denoted by $\mu \sim \mathcal{N}(\mu_0, \sigma_0^2)$, where μ_0 and σ_0 are initial beliefs about the mean and variance of the norm at time 0, respectively. In each time period t, user i updates her belief about the current tipping norm using a set of received signals, Ψ_{it}, from two smy ces in a Bayesian fashion.

Users use two types of signals to update their belief about the norm. The first type of signal is the *Personal Signal*, which captures user i's personal experience (Sen and Airiau 2007; Parrett 2011). A user may receive several tips in a day with each tip providing additional information and acting as a separate signal. In addition, since many users do not visit the platform every day, they would observe all the tips received since their last visit once they visit the platform on day t. As result, I model the personal signals user i received on day t as the tips she received in the past seven days. Formally, on each day t, user i receives N_{it}^{p} personal signals, where each personal signal s_{itn}^{p} for $n = 1, \ldots, N_{it}^{p}$ is normally distributed with mean μ and variance σ_p^2. μ represents the mean of the tipping norm and σ_p^2 is the noise associated with the personal signals, i.e.,

$$p\{s_{itn}^{p} \mid \mu, \sigma_p^2\} \sim \mathcal{N}(\mu, \sigma_p^2). \tag{1.2}$$

The second type of signal is the *Community Signal*, which captures the tips user i observes other users to give to others (Schuster, Kubacki, and Rundle-Thiele 2016). I model the community signal as the tips other users have given to all the content that user i is looking at on day t. On each day t, user i receives N_{it}^{c} community signals, where each community signal s_{itn}^{c} for $n = 1, \ldots, N_{it}^{c}$ is normally distributed with mean μ and variance σ_c^2, i.e.,

$$p\{s_{itn}^{c} \mid \mu, \sigma_c^2\} \sim \mathcal{N}(\mu, \sigma_c^2). \tag{1.3}$$

Both personal and community signals point to the same mean tipping norm μ. However, the noise or precision of the two signals are not necessarily the same resulting in different variances for the two types of signals. Note that although the signals from the two smy ces and the dependent variable are the same in nature, signals are a sample from the pool of tipping decisions over different time periods. Thus, I do not need to assume that the variances of the dependent variable and the signals are the same.

Furthermore, because I not only observe when an individual receives a signal but also the value of the signal, I can remain agnostic about whether the signals come from a distribution with constant or time-varying mean. Thus, I use the term μ as a generic term, without any indices, to refer to the tipping norm without taking a stand on whether the tipping norm is constant or time-varying.

Since I do not observe the exact time of the day when a user makes her tipping decision for each post j, I do not incorporate the sequence of user i's decisions in a single day, but instead assume that the user makes all her decisions about posts J_{it} simultaneously. In other words, the user updates her belief about the tipping norm once per day, using all the received signals from all posts J_{it}, and then makes her tipping decisions for posts J_{it}.[12] Given the number of content users interact with, I assume users do not remember the signals they received after using them to update their belief. In other words, observing the same tip on two different days results in receiving two signals of the same value.

1.4.1.4 Bayesian Updating with Correlated Signals

In my empirical context, signals are likely not independent draws from their underlying distributions. For example, in the short run, a user may receive several large tips (personal signals) due to having written a high-quality post. These correlated signals are less informative about the tipping norm than independent signals. The same issue applies to community signals. For example, good posts may receive several large tips, leading to correlated community signals. In the long-run, correlation between signals is also likely. Higher quality posts receive higher tips, motivating users to increase the quality of their content (Paridar, Ameri, and Honka 2023), which, in turn, leads to higher future tips and creates correlation between signals over time.

Correlation between signals means that the i.i.d assumption for signals no longer holds, i.e., I can no longer assume that signals come from univariate normal distributions. Instead, I assume that the signals come from a multivariate normal distribution given by

$$p\{s^p_{i1,1}, s^p_{i1,2}, \ldots, s^p_{it,M^p_{it}}, s^c_{i1,1}, s^c_{i2,2}, \ldots, s^c_{it,M^c_{it}} \mid \mu, \Sigma\} \sim \mathcal{MVN}(\mu, \Sigma) \tag{1.4}$$

where Σ is a $(M^p_{it}+M^c_{it})\times(M^p_{it}+M^c_{it})$ covariance matrix capturing the uncertainty or noise

[12] I assume that a day starts at 6 AM instead of 12 AM to account for users' activities in late hmy s as well as potential time differences. Based on the patterns in the data, the majority of activities on the platform occurs around 12 PM. As a result, I model a user's tipping decisions happening at 12 PM.

of the signals, with $M_{it}^p = \sum_{\tau=1}^{t} N_{i\tau}^p$ and $M_{it}^c = \sum_{\tau=1}^{t} N_{i\tau}^c$ representing the number of personal and community signals, respectively, user i has received until day t. Let $\Omega = \Sigma^{-1}$ be the precision matrix with its diagonal elements capturing the reciprocal variances of the signals and the off-diagonal elements correspond to partial correlations between each pair of signals. This structure can account for correlations between signals originating from the same smy ce as well as between signals from different smy ces. The first M_{it}^p rows and columns contain the precision of personal signals, and the next M_{it}^c rows and columns contain the precision of community signals. Given that all signals point to the same mean, the posterior mean and precision are given by[13]

$$\mu_{it} = \frac{1}{\omega_{it}} \left(\omega_0 \mu_0 + \sum_{k,z=1,1}^{M_{it}^p + M_{it}^c} \frac{\omega_{k,z}(s_{itk} + s_{itz})}{2} \right), \tag{1.5a}$$

$$\omega_{it} = \omega_0 + \sum_{k,z=1,1}^{M_{it}^p + M_{it}^c} \omega_{kz}, \tag{1.5b}$$

where $\omega_0 = \frac{1}{\sigma_0^2}$ and $s_{itk}, s_{itz} \in \{s_{it,1}^p, \ldots, s_{it,M_{it}^p}^p, s_{it,1}^c, \ldots, s_{it,M_{it}^c}^c\}$.

Equations 1.5a and 1.5b utilize the most general version of the precision matrix Ω in which all off-diagonal elements can take on different values. In the estimation, I impose more structure on the precision matrix that results in the estimation of three partial correlations and three decay factors. More specifically, I estimate a partial correlation among all personal signals a user receives within a day, a partial correlation among all community signals a user observes within a day, a partial correlation between the personal and community signals a user sees in a day, and three decay factors (one for each correlation) that capture the relationship between correlations on two consecutive days. In the following, I first describe the structure of the precision matrix that results in the desired correlation structure. I then present the formulas for the posterior mean and precision. All derivations are shown in detail

[13] I provide a detailed derivation of the Bayesian updating process used to compute these posterior distributions in Web Appendix A.2.

in Web Appendix A.3.

The precision matrix Ω can be decomposed into several blocks, each representing the interactions between signals of the same type within the same day, different types within the same day, and signals across different days. The diagonal elements represent the precision of personal and community signals and are denoted by ω_p and ω_c, respectively. The off-diagonal elements, that will be used to calculate the partial correlations between signals of the same type within the same day, are given by λ_p for personal signals and by λ_c for community signals. The element that will be used to calculate the correlation across signal types within the same day is denoted as λ_{pc}.[14] The off-diagonal elements across different days decay according to the decay rates δ_p, δ_c, and δ_{sc}. The decay is applied exponentially based on the time difference, i.e., for two personal signals between different days t and t', the element corresponding to the partial correlation between signals is given by $\delta_p^{|t-t'|}\lambda_p$.

More specifically, at each time period, the precision matrix Ω_t can be written as

$$\Omega_t = \begin{bmatrix} \Omega_p & \Omega_{pc}^T \\ \Omega_{pc} & \Omega_c \end{bmatrix} \tag{1.6}$$

where Ω_p and Ω_c correspond to the precision of signals and the within-smy ce correlations among signals and Ω_{pc} corresponds to the correlations between signals from different smy ces. Ω_p and Ω_c have a $t \times t$ structure of smaller blocks with rows and columns corresponding to time periods $1, \ldots, t$:

$$\Omega_{|\omega,\lambda} \in \{\Omega_{p|\omega_p,\lambda_p}, \Omega_{c|\omega_c,\lambda_c}\} = \begin{bmatrix} \Omega_{11} & \delta^1\Omega_{12} & \cdots & \delta^t\Omega_{1t} \\ \delta^1\Omega_{21} & \Omega_{22} & \cdots & \delta^{t-1}\Omega_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ \delta^t\Omega_{t1} & \delta^{t-1}\Omega_{t2} & \cdots & \Omega_{tt} \end{bmatrix} \tag{1.7}$$

[14]The formulas to calculate the partial correlations ρ are given by $\rho_p = \frac{\omega_p}{\lambda_p}$, $\rho_c = \frac{\omega_c}{\lambda_c}$, and $\rho_{pc} = \frac{\sqrt{\omega_p \times \omega_c}}{\lambda_{pc}}$.

The diagonal Ω_{kk} blocks capture the precision and correlation of the $N_{ik} \in \{N_{ik}^p, N_{ik}^c\}$ signals received from a smy ce in each time period. Thus, each Ω_{kk} is a $N_{ik} \times N_{ik}$ matrix with diagonal elements ω and off-diagonal elements λ. The off-diagonal Ω_{zk} matrix blocks are $\lambda J_{N_{iz} \times N_{ik}}$ matrices, J being an all-ones matrix:

$$\Omega_{kk} = \begin{bmatrix} \omega & \lambda & \cdots & \lambda \\ \lambda & \omega & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \omega \end{bmatrix}_{N_{ik} \times N_{ik}} , \quad \Omega_{zk} = \begin{bmatrix} \lambda & \lambda & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \lambda \end{bmatrix}_{N_{iz} \times N_{ik}} \tag{1.8}$$

I now turn to the matrix capturing the correlations between signals from different smy ces, Ω_{pc}. This matrix also consists of $t \times t$ blocks, with blocks representing the partial correlation between signals of different types within and across time periods. The correlation between two signals is proportional to λ_{pc}, decreasing at an exponential rate of δ_{pc} as the time difference between the two signals increases. Formally, the matrix consists of block matrices of size $N_{iz}^p \times N_{ik}^c$ for personal signals of time period z and community signals of time period k with all elements equal to $\delta_{pc}^{|z-k|} \lambda_{pc}$:

$$\Omega_{pc} = \lambda_{pc} \begin{bmatrix} J_{N_{i1}^p \times N_{i1}^c} & \delta_{pc}^1 J_{N_{i1}^p \times N_{i2}^c} & \cdots & \delta_{pc}^t J_{N_{i1}^p \times N_{it}^c} \\ \delta_{pc}^1 J_{N_{i2}^p \times N_{i1}^c} & J_{N_{i2}^p \times N_{i2}^c} & \cdots & \delta_{pc}^{t-1} J_{N_{i2}^p \times N_{it}^c} \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{pc}^t \; J_{N_{it}^p \times N_{i1}^c} & \delta_{pc}^{t-1} J_{N_{it}^p \times N_{i2}^c} & \cdots & J_{N_{it}^p \times N_{it}^c} \end{bmatrix} \tag{1.9}$$

Given this structure for Ω, the posterior precision ω_t can be derived as

$$\begin{aligned}\omega_t = \omega_0 + (\omega_p - \lambda_p) \sum_{k=1}^{t} N_{ik}^p + (\omega_c - \lambda_c) \sum_{k=1}^{t} N_{ik}^c \\ + \lambda_p \sum_{k,z=1}^{t} \delta_p^{|k-z|} N_{iz}^p \times N_{ik}^p \\ + \lambda_c \sum_{k,z=1}^{t} \delta_c^{|k-z|} N_{iz}^c \times N_{ik}^c \\ + 2\lambda_{pc} \sum_{k,z=1,1}^{t} \delta_{pc}^{|k-z|} N_{iz}^p \times N_{ik}^c\end{aligned} \tag{1.10}$$

and the posterior mean is given by

$$\begin{aligned}\mu_{it} = \frac{1}{\omega_t} \Bigg(\omega_0\mu_0 + (\omega_p - \lambda_p) \sum_{k=1}^{t} \mathbb{S}_{ik}^p + (\omega_c - \lambda_c) \sum_{k=1}^{t} \mathbb{S}_{ik}^c + \\ \lambda_p \sum_{k,z=1}^{t} \left(\delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p\right) + \\ \lambda_c \sum_{k,z=1}^{t} \left(\delta_c^{|k-z|} N_{ik}^c \mathbb{S}_{iz}^c\right) + \\ \frac{\lambda_{pc}}{2} \sum_{k,z=1}^{t} \left(\delta_{pc}^{|k-z|} N_{ik}^c \mathbb{S}_{iz}^p\right) + \frac{\lambda_{pc}}{2} \sum_{k,z=1}^{t} \left(\delta_{pc}^{|k-z|} N_{iz}^p \mathbb{S}_{ik}^c\right) \Bigg)\end{aligned} \tag{1.11}$$

where $\mathbb{S}_{ik}^p = \sum_{r=1}^{N_{ik}^p} s_{ikr}^p$, the sum of the personal signal values on day k, and $\mathbb{S}_{ik}^c$ is defined similarly.

1.4.2 Estimation

The log likelihood function for the Tobit model is given by:

$$\log L(\theta|y,\Psi,Q,C) = \sum_{i=1}^{N}\sum_{\tau=1}^{\mathcal{T}}\sum_{j=1}^{J}\Bigg\{ I(y_{ijt}>0)\log\left[\frac{1}{\sigma}\phi\left(\frac{y_{ijt}-(\alpha_i+\beta\mu_{it}+\gamma' Q_{ijt}+\eta' C_{ijt})}{\sigma}\right)\right]$$
$$+ I(y_{ijt}=0)\log\left[\Phi\left(\frac{\alpha_i+\beta\mu_{it}+\gamma' Q_{ijt}+\eta' C_{ijt}}{\sigma}\right)\right]\Bigg\} \tag{1.12}$$

ϕ is the standard normal probability density function, Φ is the standard normal cumulative distribution function, and $I(\cdot)$ is an indicator function. I set σ_0^2 =1 for identification. $\theta = (\alpha_i,\ \beta,\ \gamma,\ \eta,\ \mu_0,\ \omega_p,\ \omega_c,\ \delta_p,\ \delta_c,\ \delta_{pc},\ \lambda_p,\ \lambda_c,\ \lambda_{pc},\ \sigma)$ is the vector of parameters to be estimated.

In Web Appendix A.3, I show how I rewrite the formulas for the posterior mean and posterior precision to avoid calculating permutations and speed up their computations during the estimation. My data contain about 4.2 million observations and I estimate about 1,800 individual fixed effects.[15] Even though I estimate a non-linear model with normally distributed errors, I do not face the incidental parameter problem in my empirical setting because of the large T, i.e., I have a large number of observations per user (average of 2,350 observations per user) (Neyman and Scott 1948; Arellano and Hahn 2007). Because of the size of the data and the estimation of a large number of fixed effects, the model estimation takes about 14 days. To calculate standard errors of the parameter estimates, I use the BHHH estimator, i.e., the outer product of the gradient, instead of the numerical Hessian (Berndt et al. 1974). All standard errors are clustered at the individual level.

1.5 Results

I present the estimation results in Table 1.3. Column (i) presents the results for a model in which users do not use the two signals to continuously update their belief about the current norm, but instead use their current values independently from the past to decide to how much to tip. Columns (ii) and (iii) depict the results for models in which users utilize both

[15] In a Tobit model, the estimation of fixed effects cannot be avoided by de-meaning the data.

signals to update their beliefs about the tipping norm in a Bayesian fashion as a function of all the signals received so far. For the model whose results are shown in column (ii), I assume that signals are independent. Column (iii) depicts the results for my main model in which signals are correlated with the desired structure described in equations (8) - (12).

In the model without learning (column (i)), the coefficients for personal (Tips User Received in Past Week) and community signals (Prior Tips Given to Focal Content) are both positive and significant. The effect of a community signal is about 3.5 times larger than the effect of a personal signal. The loglikelihood, AIC, and BIC all improve when I move to a model in which users learn the tipping norm via Bayesian updating with independent signals (see column (ii)). This improvement underscores the importance of accounting for learning from past and current signals instead of simply controlling for current signals. These three model fit measures again improve considerably when I transition to my main model depicted in column (iii) in which users learn the tipping norm via Bayesian updating with correlated signals. These improvements underline the importance of accounting for correlations between signals. In discussing the results, I focus on my main model shown in column (iii).

Table 1.3: Empirical Results

Variable	(i) Without Learning	(ii) Learning With Independent Signals	(iii) Learning With Correlated Signals
Learning			
Posterior Belief of Tipping Norm		1.3272***	1.8470***
		(0.0000)	(0.0614)
Prior Belief of Mean of Tipping Norm		2.1652***	2.1049***
		(0.0000)	(0.0662)
Precision of Personal Signals		1.3064***	1.5930***
		(0.0001)	(0.2335)
Precision of Community Signals		0.6657***	0.2746***
		(0.0004)	(0.0212)
Off-Diagonal Elements between Personal Signals			7.6370***
			(0.3180)
Off-Diagonal Elements between Community Signals			1.8157***
			(0.0180)
Off-Diagonal Elements between Personal and Community Signals			4.1867***
			(0.0280)
Decay Factor of Partial Correlation between Personal Signals			0.6580***
			(0.0908)
Decay Factor of Partial Correlation between Community Signals			0.5967***
			(0.0324)
Decay Factor of Partial Correlation between Personal and Community Signals			0.7338**
			(0.1877)
Controlling for Tipping			
Tips User Received in Past Week	0.0406***		
	(0.0018)		
Prior Tips Given to Focal Content	0.1540***		
	(0.0099)		
Content Quality			
Image Quality	0.0600	0.053***	0.0600***
	(0.0423)	(0.0001)	(0.0166)
Thread Quality	0.3640***	0.4067***	0.4640***
	(0.0267)	(0.0235)	(0.0227)
Reply Quality	0.5080***	0.5225***	0.5175***
	(0.0350)	(0.0137)	(0.0128)
Geeklist Quality	0.3790***	0.4580***	0.3789***
	(0.0322)	(0.0253)	(0.0156)
Other Variables			
Constant	-5.6700***	-7.6337***	-7.1216***
	(0.5342)	(0.0064)	(0.1802)
Variance of Dependent Variable	5.5690***	7.9946***	7.2595***
	(0.2620)	(0.0506)	(0.1560)
Control Variables	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes
Calendar Month FEs	Yes	Yes	Yes
Number of observations	3,937,582	3,937,582	3,937,582
AIC	98,144.00	97,006.52	94,038.16
BIC	122,144.04	121,085.31	118,010.45
LogLikelihood	-47,030.18	-46,624.32	-45,201.08

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The coefficient capturing the effect of the posterior belief of the tipping norm is positive, significant, and large, highlighting the important role that social norms play in shaping tipping behavior. The estimate for the mean of the prior is 2.1049, translating into an initial belief about the mean tipping norm of about 7.20 GG. The estimate of the precision of personal signals is 1.5930 and statistically significant. This parameter value implies that it takes receiving about 2 tips for a user's uncertainty to reduce by 90%. Similarly, the estimate for the precision of community signals is 0.2592, implying that 4 community signals (within a day) or 3 community signals (one per day on fmy consecutive days) are needed to reduce the uncertainty by 90%. In other words, receiving a tip is more informative to users than observing tips given to others, as it requires fewer tips to reduce the user's uncertainty about the current tipping norm. This means that users place more weight on personal experiences compared to community feedback when updating their beliefs and making tipping decisions.

All three off-diagonal elements of the precision matrix are statistically significant and range from 1.49 to 7.46. Using the formulas from footnote 14, the partial correlation estimates range from 0.08 to 0.11. These partial correlations are of moderate size and support the notion that signals are not independent in my empirical context, but also show that users do not always tip what the previous tipper has given. The three decay factors are also all statistically significant and range from 0.60 to 0.75. These estimates imply that the three correlations between signals decay quite fast: they are about 10% of their original magnitudes after five to eight time periods.

The estimates for the content quality variables are all significant and positive, indicating that higher quality posts generally receive higher tips. This suggests that users recognize and reward the effort put into creating high-quality content.

Figure 1.6 shows the average predicted tip amount over time as well as users' posterior belief about the norm at each point in time. The decreasing standard deviation of the predicted tip amounts in Figure 1.6(a) indicates less variation and more convergence in tipping behavior. Additionally, the average posterior norm change slows down after a sharp decrease in the first six months, with its standard deviation also decreasing, further supporting a convergence of beliefs about the norm (Figure 1.6(b)). I show the average predicted number of tipping incidences in a day and the predicted total tip amount given by all users in a day in 1.B.

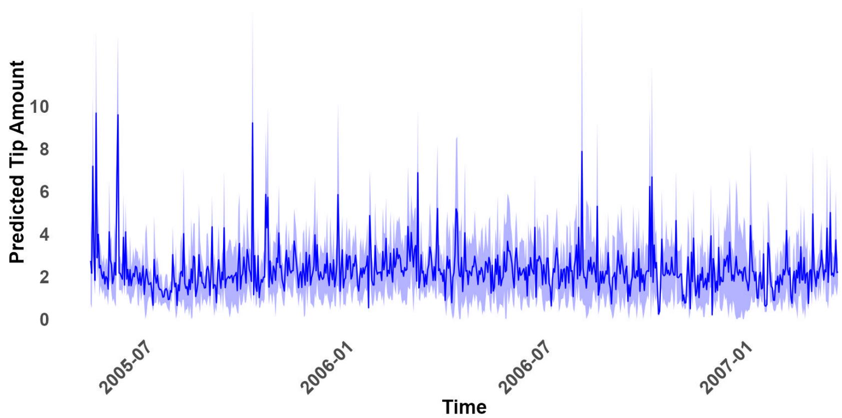

(a) Predicted Tip Amount Per Tip Incidence

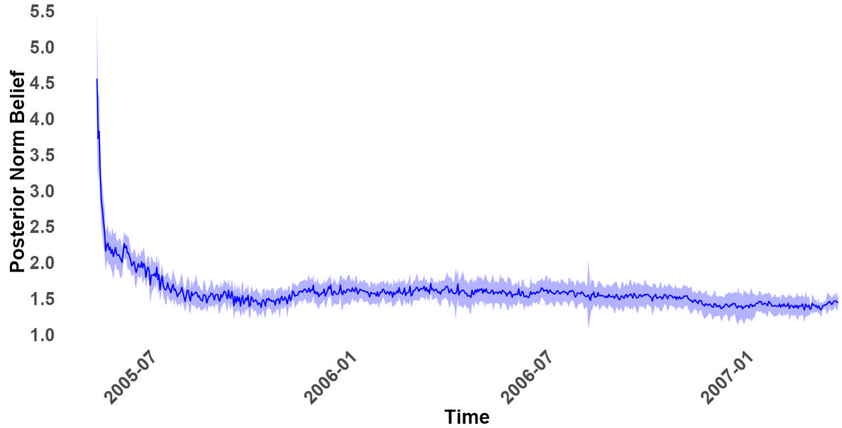

(b) Posterior Belief about Norm

Figure 1.6: Predicted Tips and Posterior Norm Belief with 95% Confidence Intervals

1.5.1 Model Fit

I further examine the predictive performance of the model through a simulation where I predict users' decisions and use the predicted tip amounts as the signals other users receive in following days. My model predicts 1,697 users giving 7,136 tips with an average tip amount of 2.33 GG. In the actual data, I observe 1,785 users giving 6,672 tips with an average tip amount of 1.93 GG. Thus, I conclude that the model fits the data patterns well.

1.6 Prediction Exercises

In this section, I first quantify the relative contributions of drivers of tip decisions. I then investigate how different information disclosure strategies affect the perceived tipping norm development and users' tipping decisions. And lastly, I examine how "sticky" the tipping norm and tipping decisions are and whether social platforms can significantly affect them at later stages.

1.6.1 Tip Decomposition

Here, I examine the contributions of different utility components to the tip amounts users give. I organize utility components into fmy groups: quality, perceived norm, user characteristics (via individual fixed effects) and others, which contains the remaining variables in the utility function. I measure the relative importance of each group by predicting the portion of utility driven by it.

Table 1.4 reports descriptive statistics for the portion of tip amount due to quality, norm, individual characteristics and other factors. On average, the perceived norm represents 67% of the tip amounts given on the platform, while the quality of content contributes 28%. In contrast to previous literature (Chandar et al. 2019; Kim, Amir, and Wilbur 2023), which has found that tipping decisions are largely driven by tipper characteristics, my results indicate that tipper characteristics only play a minor role in individuals' tipping decisions.

Table 1.4: Description of Tip Amount Decomposition

Utility Components	Mean	Median	Min	Max	SD	N
Quality	0.28	0.23	0	1	0.17	7,136
Norm	0.67	0.71	0	1	0.18	7,136
User	0.05	0.00	0	1	0.07	7,136
Other	0.00	0.00	0	1	0.04	7,136

In Figure 1.7, I examine how the contributions of the fmy groups of utility components to tip amounts change over time. Initially, the perceived norm drives nearly 90% of the tip amounts, followed by quality and individual characteristics. Within the first three months, the portion driven by the perceived norm declines rapidly and the portion driven by quality increases rapidly. The portions remain largely stable after the first six months with the exception of individual characteristics whose portion declines somewhat starting in the second half of 2006. To summarize, while the relative contribution of the social norm has decreased compared to the early weeks, it always makes by far the largest contribution.

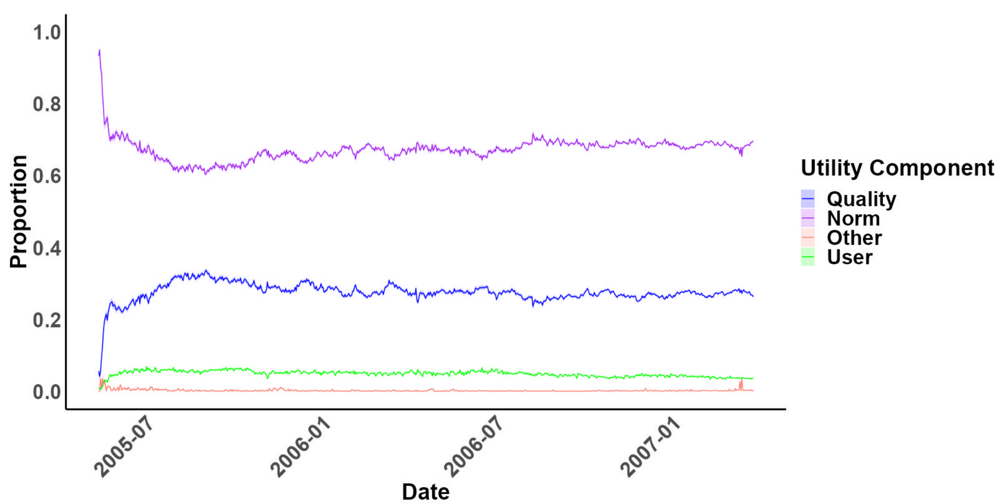

Figure 1.7: Tip Amount Decomposition over Time with 95% Confidence Intervals

1.6.2 Information Disclosure

Users receive two types of signals about the tipping norm on BGG: personal signals (tips received by the focal user) and community signals (tips given to focal content by other users). However, this is not the case on all digital platforms. For example, while platforms such as Twitch or YouTube make tips visible and salient on users' screens, platforms such as Patreon or Cameo keep monetary contributions private between the supporter and the content creator. Here, I empirically evaluate how different information disclosure strategies, i.e., tip visibility, impact users' tipping behavior.

To investigate the effects of information disclosure, I predict tipping behavior when either the personal or community signals are invisible ("No Disclosure") and when users can only see the average community signal ("Partial Disclosure").[16] For each scenario, I predict users' tipping behavior based on their perceptions of the tipping norm derived from the available signals and under the assumption that the quality of the content remains constant. For each scenario, I repeat the prediction calculations 100 times (with different error draws) and then compute the average predictions. I then compare the average predictions from these three scenarios to the average predictions from my main model.

I report the average percentage changes for each scenario compared to my baseline main model in Table 1.5. When community signals are partially or completely invisible (columns (i) and (ii)), users tip more frequently but give smaller amounts when they tip. In both these scenarios, the increases in tip frequency more compensate for the decreases in tip amounts, resulting in increases of total tip amounts by 39%. When personal signals are undisclosed, users also tip more frequently and smaller amounts. However, in this scenario, the increase in tip incidences just offsets the decrease in tip amounts resulting in a slightly larger total tip amount.

[16] I conduct the scenario of personal signals being invisible mainly for comparability reasons. In practice, this scenario could occur when users can only see the tips they received with a time delay, e.g., only see the sum of all tips received at the end of a week or a month.

Table 1.5: Tipping Behavior Under Different Information Disclosure Scenarios

	(i)	(ii)	(iii)
	Community Signal		*Personal Signal*
	Partial Disclosure	No Disclosure	No Disclosure
Number of Tip Incidences	69.44%	100.02%	11.38%
Amount per Tip	-17.99%	-30.82%	-8.81%
Sum of All Tips	38.97%	38.37%	1.58%
Number of Unique Tippers	48.58%	8.24%	-8.98%
Number of Unique Tippees	69.06%	62.49%	1.48%
Number of Unique Tipped Content	61.71%	73.09%	0.66%

Making community signals partially or completely invisible also significantly increases the number of unique tippers, number of unique tippees, and number of unique tipped content when community signals are not or only partially disclosed (columns (i) and (ii) in Table 1.5). The picture is different when personal signals are undisclosed. Then, the number of unique tippers declines by 9%, and the numbers of unique tippees and unique tipped content slightly increase. However, reducing the amount of information users have also results in a delay in norm formation and/or more uncertainty in the norm even in later stages. In Figure 1.8, I plot users' average posterior norm beliefs for all fmy scenarios over time. The variation in average posterior norm beliefs is larger when users do not observe or only partially observe community signals.

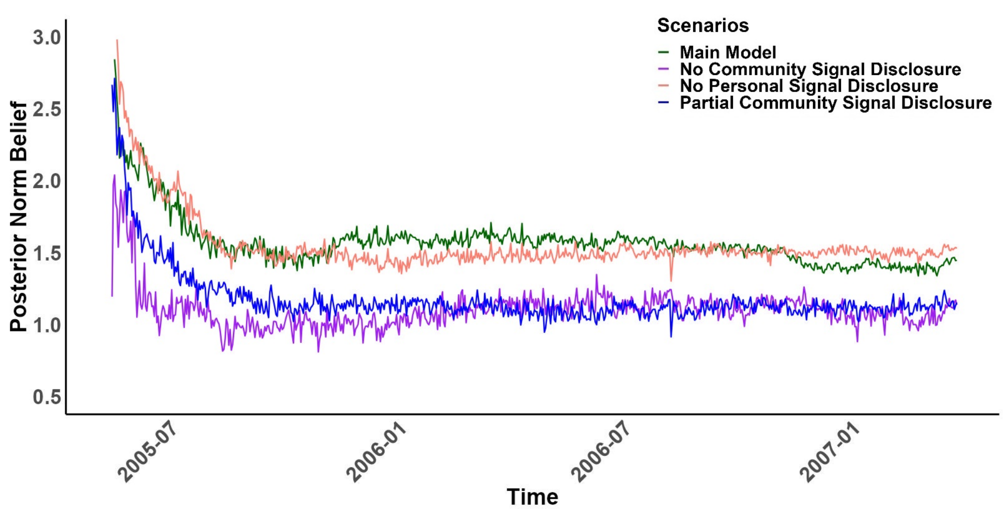

Figure 1.8: Posterior Norm Over Time

Comparing the magnitudes of the changes in outcomes in the three scenarios, the changes are larger when community signals are partially or completely invisible compared to personal signals being undisclosed. These differences in magnitudes arise because the number of signals from each smy ce varies significantly in my data; on average, a user receives 30 community signals for every personal signal received. Thus, while each personal signal is 5 times more informative (see Table 1.3), in practice, community signals result in an effect that is 6 times greater than that of personal signals.[17]

1.6.3 Tipping and Norm Stickiness

In this section, I examine how "sticky" social norms and tipping behavior are. More specifically, I want to understand whether and how quickly platforms can change social norms and tipping behavior after a norm has been established by making changes to the information they disclose. Recall from Section 2.5 that the tipping norm is largely constant after the first six months of the study period. To investigate whether platforms can influence the norm and tipping behavior in later stages, I implement the same three scenarios as in the previous section but only in the second half of the study period, i.e., after the first 11 months of the study period.[18] In these three scenarios, tipping information is fully disclosed (as in my main model) in the first 11 months of the study period. I then compare the predictions from these three scenarios to those from my main model, where the platform discloses tipping information.

Figure 1.9 displays the tipping norm development for the fmy scenarios over time. As expected, it takes longer until tipping norms diverge compared to introducing different information disclosures at the introduction of tipping (see Figure 1.8) because users have already accumulated information about tipping amounts and formed a perceived tipping norm in

[17]The average value of personal signals is 0.94 and the average value of community signals is 0.92. Since these values are very close, it is less likely that the observed differences are due to the values of the signals.

[18]For each scenario, I repeat the prediction calculations 100 times (with different error draws) and then compute the average predictions. I then compare the average predictions from these three scenarios to the average predictions from my main model.

the first 11 months. Compared to the main model, the tipping norms under two of the three information disclosure scenarios, namely, no community signal disclosure and no personal signal disclosure, are higher at the end of the study period. Partial community signal disclosure results in a tipping norm that is very similar to the tipping norm from my main model with full information disclosure. Generally, the differences in tipping norms between the different scenarios and my main model are smaller compared with the end of the study period in Figure 1.8 indicating a degree of stickiness that persists even 11 months after a change in platform information disclosure.

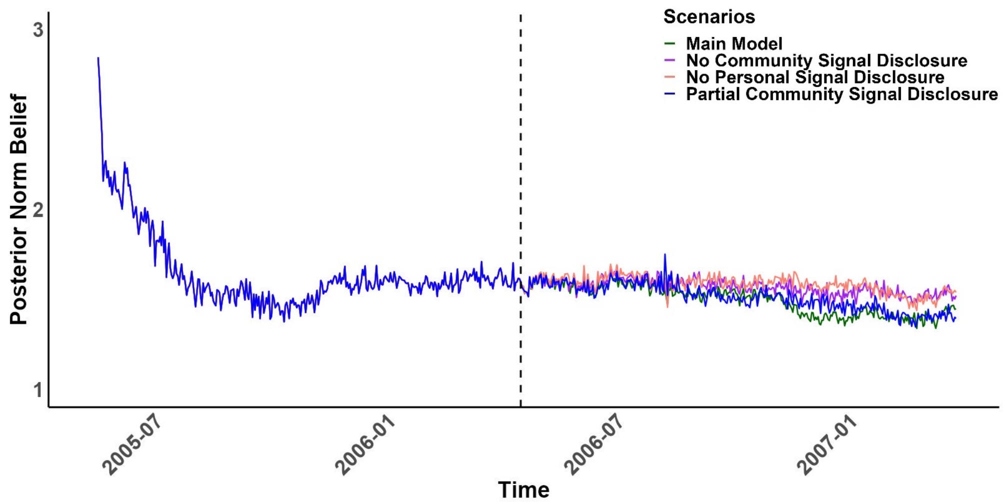

Figure 1.9: Posterior Norm Over Time

The results for tipping decisions are shown in Table 1.6. The percentage changes in Table 1.6 were calculated based on behavior in the last two months of the study period to describe tipping after a transition period. When community signals are only partially or not at all disclosed in the second half of the study period, the resulting changes in tipping behavior are directionally the same as those shown in Table 1.5. However, the magnitudes of the effects are different. Not surprisingly and in line with the results for the tipping norms, the effects are mostly smaller since the changes were only implemented for half of the study period. The results for no disclosure of personal signals are also directionally the same as those shown in Table 1.5 with the exception of a small increase in the number of unique tippers instead of a decrease.

Table 1.6: Tipping Behavior Under Changes in Information Disclosure in Second Half of Study Period (Based on Last 2 Months)

	(i)	(ii)	(iii)
	Community Signal		*Personal Signal*
	Partial Disclosure	No Disclosure	No Disclosure
Number of Tip Incidences	24.47%	51.32%	11.91%
Amount per Tip	-14.61%	-21.53%	-7.45%
Sum of All Tips	6.28%	18.74%	3.57%
Number of Unique Tippers	22.23%	7.90%	1.85%
Number of Unique Tippees	34.32%	31.15%	8.42%
Number of Unique Tipped Content	33.19%	48.32%	6.16%

So far, I have shown that there is some stickiness in users' norm belief and tipping behavior. Next, I want to quantify the amount of stickiness in terms of its medium-run effects on tipping behavior. I do so as follows: for each of the three scenarios, I compare users' tipping behavior for the last two months of the study period for the case when a scenario was introduced in the middle of the study period versus at the beginning of the study period. For example, I compare users' tipping behavior when no community signal was shown in the second half of the study period (but shown in the first half of the study period) to the scenario when no community signal was shown during the whole study period. This comparison measures the (medium-run) stickiness of users having had full information disclosure in the first half of the study period. The results are presented in Table 1.7.

Table 1.7: Measures of Medium-Run Stickiness (Based on Last 2 Months)

	(i)	(ii)	(iii)
	Community Signal		*Personal Signal*
	Partial Disclosure	No Disclosure	No Disclosure
Number of Tip Incidences	21.08%	-3.56%	6.31%
Amount per Tip	-11.31%	2.27%	-3.64%
Sum of All Tips	7.29%	-1.64%	2.56%
Number of Unique Tippers	4.23%	-5.12%	15.52%
Number of Unique Tippees	5.85%	-22.01%	14.65%
Number of Unique Tipped Content	14.49%	-14.59%	7.14%

Let us consider a delayed introduction of no community signal disclosure, i.e., what are the medium-run consequences of having had community signal disclosure in the first half of the study period (see column (ii) in Table 1.7). Here, it is important to keep in mind that the results in Table 1.7 were calculated for the last two months of the study period, i.e., after no community signal disclosure has been in place for 9 months (if it was introduced for 2nd half of study period) or for 20 months (if it was introduced at the beginning of the study period). While the consequences of full community signal disclosure in first half of the study period are of moderate magnitudes for the number of tip incidences, the amount per tip, and the sum of all tips, the picture looks different when I consider the number of unique tippers, the number of unique tippees, and the number of unique tipped content. The stickiness of fewer unique tippers, fewer unique tippees, and fewer unique tipped content, that is characteristic of tipping behavior under full community signal disclosure, persists even 9 months after it was dropped. Similar observations of sticky tipping behavior can also be observed for the other two scenarios shown in columns (i) and (iii) in Table 1.7.

1.7 Discussion and Conclusion

Understanding how to influence users' tipping behavior is crucial for online platforms looking to incentivize content creators and to build an engaged community. It can provide firms with valuable insights into strategies that can motivate their community's stakeholders and boost overall engagement. This paper examines the evolution of tipping norms within an online community, focusing on how users form and update their beliefs about the tipping norm through Bayesian updating with correlated signals. I study how these beliefs about the current tipping norm, combined with content quality and other factors, influence tipping behavior.

My findings reveal that users' tipping decisions are significantly shaped by their current perception of the tipping norm, which is continually updated based on their personal experiences of receiving tips and observing others' tip on the platform. Specifically, I show that that personal experiences with tips received are more informative than observed tipping behavior in the community, impacting users' perceptions of the tipping norm significantly. This dynamic updating process underscores the adaptive nature of user behavior in response to the evolving tipping environment on digital platforms.

I then analyze the impact of different information provision strategies by partially or fully removing the visibility of personal and community signals. My findings indicate that the visibility of these signals significantly affects tipping behavior. Specifically, when community tips are not visible, users tend to tip more frequently but smaller amounts, resulting in a higher total tip amount. These findings suggest that platforms can use tip visibility as a strategic tool to influence tipping behavior. Further, I also examine how sticky the perceived tipping norm and tipping behavior is after a change in the platform's information disclosure. I find evidence for stickiness even in the medium-run, especially as it related to the breadth of tipping.

These findings have practical implications for platform managers and content creators, offering strategies to enhance user engagement and tipping behaviors. For platforms that receive a portion of the tips as revenue, optimizing tipping behavior can directly impact their financial sustainability. Using my findings, platforms can ensure a steady stream of income from tipping activities. Furthermore, my findings suggest that while making tipping signals not visible can increase the number of tips and total tip amount, it might decrease the average amount per tip. This has several implications. Platforms need to balance the visibility of tipping signals to optimize overall tipping behavior while considering the potential impact on individual content creators' income. While increasing the frequency of smaller tips might be financially beneficial for platforms receiving a portion of the tips, it could lead to less desirable outcomes for content creators if their overall income decreases due to lower amounts per tip. Platforms should consider their specific circumstances and design their strategies accordingly.

1.A Derivation of Bayesian Updating Formulas

1.A.1 Bayesian Updating with Independent Signals

In this section of the appendix, I present the derivation of the Bayesian updating process utilized to compute posterior beliefs about the tipping norm, assuming known and deterministic variances for signals. Here, I model the evolution of beliefs as users receive personal and community signals, each assumed to be normally distributed.

The initial prior belief about the tipping norm for each user i is represented as a normal distribution with mean μ_0 and variance σ_0^2, expressed as $\mu \sim \mathcal{N}(\mu_0, \sigma_0^2)$.

The personal signals $s_{it,n}^p$, for $n = 1, \ldots, N_{it}^p$, where each signal $s_{it,n}^p \mid \mu$ is normally distributed with mean μ and variance σ_p^2, are represented by:

$$s_{it,n}^p \mid \mu \sim \mathcal{N}(\mu, \sigma_p^2). \tag{A1}$$

Similarly, community signals $s_{it,n}^c$, for $n = 1, \ldots, N_{it}^c$, where each signal $s_{it,n}^c \mid \mu$ follows a normal distribution with the same mean μ but different variance σ_c^2, are given by:

$$s_{it,n}^c \mid \mu \sim \mathcal{N}(\mu, \sigma_c^2). \tag{A2}$$

The Bayesian updating rule, combining the prior and the likelihoods of all personal and community signals, is formulated as:

$$p(\mu \mid s_1^p, s_2^p, \ldots, s_{n_2}^c) = p(s_1^p \mid \mu)p(s_2^p \mid \mu) \ldots p(s_{n_2}^c \mid \mu)p(\mu). \tag{A3}$$

This equation reflects the multiplication of the likelihoods of observing each signal given the tipping norm μ, with each signal treated as conditionally independent given μ.

Using the normal distribution for $p(\mu \mid \mu_0, \sigma_0^2)$, the prior probability density function of

μ, I express it as:

$$p(\mu \mid \mu_0, \sigma_0^2) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left[-\frac{1}{2\sigma_0^2}(\mu - \mu_0)^2\right] \tag{A4}$$

Substituting this into the equation for Bayesian updating, and considering the normal distributions of s_i^p and s_i^c, the combined density function becomes:

$$\frac{1}{(2\pi)^{\frac{n_1+n_2+1}{2}}} \frac{1}{(\sigma_0^2)^{\frac{1}{2}}(\sigma_1^2)^{\frac{n_1}{2}}(\sigma_2^2)^{\frac{n_2}{2}}} \exp\left[-\tfrac{1}{2\sigma_1^2}\textstyle\sum_{i=1}^{n_1}(s_i^p - \mu)^2\right] \exp\left[-\tfrac{1}{2\sigma_2^2}\textstyle\sum_{i=1}^{n_2}(s_i^c - \mu)^2\right] \exp\left[-\tfrac{1}{2\sigma_0^2}(\mu - \mu_0)^2\right] \tag{A5}$$

$$\propto \exp\left[-\frac{1}{2\sigma_1^2}\sum_{i=1}^{n_1}(s_i^p - \mu)^2 - \frac{1}{2\sigma_2^2}\sum_{i=1}^{n_2}(s_i^c - \mu)^2 - \frac{1}{2\sigma_0^2}(\mu - \mu_0)^2\right] \tag{A6}$$

$$= \exp\left[-\frac{\sum\limits_{i=1}^{n_1} s_i^{p2} + n_1\mu^2 - 2\mu\sum\limits_{i=1}^{n_1} s_i^p}{2\sigma_1^2} - \frac{\sum\limits_{i=1}^{n_2} s_i^{c2} + n_2\mu^2 - 2\mu\sum\limits_{i=1}^{n_2} s_i^c}{2\sigma_2^2} - \frac{\mu^2 + \mu_0^2 - 2\mu\mu_0}{2\sigma_0^2}\right] \tag{A7}$$

Simplifying the exponential and combining terms involving μ, I derive the expression:

$$\exp\left[-\frac{\mu^2}{2}\left(\frac{1}{\sigma_0^2} + \frac{n_1}{\sigma_1^2} + \frac{n_2}{\sigma_2^2}\right) + \mu\left(\frac{\sum\limits_{i=1}^{n_1} s_i^p}{\sigma_1^2} + \frac{\sum\limits_{i=1}^{n_2} s_i^c}{\sigma_2^2} + \frac{\mu_0}{\sigma_0^2}\right) + \left(\frac{\sum\limits_{i=1}^{n_1} s_i^{p2}}{\sigma_1^2} + \frac{\sum\limits_{i=1}^{n_2} s_i^{x2}}{\sigma_2^2} + \frac{\mu_0^2}{\sigma_0^2}\right)\right] \tag{A8}$$

Completing the square allows us to equate the above expression to the standard form of a normal distribution in terms of μ. By matching the coefficients, I can directly derive the corresponding mean and variance:

$$p(\mu \mid s_1^p, s_2^p, \ldots, s_{n_2}^c) \propto \exp\left[-\frac{1}{2\sigma_t^2}(\mu - \mu_t)^2\right] \propto \exp\left[-\frac{\mu^2 - 2\mu\mu_t + \mu_t^2}{2\sigma_t^2}\right] \tag{A9}$$

then:

$$-\frac{\mu^2}{2\sigma_t^2} = -\frac{\mu^2}{2}\left(\frac{1}{\sigma_0^2} + \frac{n_1}{\sigma_1^2} + \frac{n_2}{\sigma_2^2}\right) \tag{A10}$$

$$\implies \sigma_t^2 = \left(\frac{1}{\sigma_0^2} + \frac{n_1}{\sigma_1^2} + \frac{n_2}{\sigma_2^2}\right)^{-1} \tag{A11}$$

and

$$2\mu\mu_t 2\sigma_t^2 = \mu\left(\frac{\sum_{i=1}^{n_1} s_i^p}{\sigma_1^2} + \frac{\sum_{i=1}^{n_2} s_i^c}{\sigma_2^2} + \frac{\mu_0}{\sigma_0^2}\right) \tag{A12}$$

$$\implies \mu_t = \sigma_t^2\left(\frac{\sum_{i=1}^{n_1} s_i^p}{\sigma_1^2} + \frac{\sum_{i=1}^{n_2} s_i^c}{\sigma_2^2} + \frac{\mu_0}{\sigma_0^2}\right) \tag{A13}$$

1.A.2 Bayesian Updating with Correlated Signals

If the signals are not independent from each other, the probability $p(s_1^p, s_2^p, \ldots, s_{n_2}^c \mid \mu)$ does not break into separate probabilities anymore. Instead, I have to use the joint probability of the signals, i.e., the multivariate normal distribution:

$$p(s_1^p, s_2^p, \ldots, s_{n_2}^c \mid \mu) = \left(\frac{1}{(2\pi)^{n/2}}\right)(\det \Sigma)^{-1/2} \exp\left(-\frac{1}{2}(X-\mu)^\top \Sigma^{-1}(X-\mu)\right) \tag{A14}$$

where Σ is the covariance matrix. Its diagonal elements are the variances of the signals and the off-diagonal elements are the covariances between the signals. In the case of a multivariate normally distributed posterior, it is more convenient to write the equations using precision notation. Let $\Omega = \Sigma^{-1}$, with ω_{ij} representing the precision between signals i

and j. This structure can account for the correlations between signals originating from the same type of smy ce as well as between signals from different smy ces. I assume that the first $M_{it}^p = \sum_{k=1}^t N_{ik}^p$ rows and columns contain the precision of personal signals, and the next $M_{it}^c = \sum_{k=1}^t N_{ik}^c$ rows and columns contain the precision of community signals, forming an $(M_{it}^p + M_{it}^c) \times (M_{it}^p + M_{it}^c)$ precision matrix. Note that the diagonal elements corresponding to the first M_{it}^p signals are equal since all personal signals have the same variance, and similarly, the diagonal elements for the next M_{it}^c signals are equal.

I now calculate the posterior distribution. Signals from a multivariate normal distribution are conjugate with a multivariate normal prior, and, in my context, I can simplify the posterior distribution even further because the signals all have the same mean μ:

$$
\begin{aligned}
p\left(\mu \mid s_{i1}^p, s_{i1}^c, \ldots, s_{M_{it}^p+M_{it}^c}\right) \propto\ & \exp\left[-\frac{1}{2}\sum_{k,z=1}^{M_{it}^p+M_{it}^c} (s_{ik}-\mu)\,\omega_{kz}(s_{iz}-\mu) - \frac{\omega_0}{2}(\mu-\mu_0)^2\right] \\
=\ & \exp\left[-\frac{1}{2}\sum_{k,z=1}^{M_{it}^p+M_{it}^c} \omega_{kz}\left(s_{ik}s_{iz} - \mu(s_{ik}+s_{iz}) + \mu^2\right)\right. \\
& \left. -\frac{1}{2}\omega_0(\mu^2 - 2\mu\mu_0 + \mu_0^2)\right] \\
=\ & \exp\left[-\frac{\mu^2}{2}\left(\omega_0 + \sum_{k,z=1}^{M_{it}^p+M_{it}^c}\omega_{kz}\right)\right. \\
& + \frac{\mu}{2}\left(2\omega_0\mu_0 + \sum_{k,z=1}^{M_{it}^p+M_{it}^c}\omega_{kz}(s_{ik}+s_{iz})\right) \\
& \left. -\frac{1}{2}\left(\omega_0\mu_0^2 + \sum_{k,z=1}^{M_{it}^p+M_{it}^c}\omega_{kz}s_{ik}s_{iz}\right)\right]
\end{aligned}
\tag{A15}
$$

After completing the square in the expression for μ, I derive the precision ω_t and the mean μ_t of the posterior distribution as

$$
\omega_{it} = \omega_0 + \sum_{k,z=1,1}^{M_{it}^p+M_{it}^c} \omega_{k,z}\ , \tag{A16}
$$

$$\mu_{it} = \frac{\omega_0 \mu_0 + \sum_{k,z=1,1}^{M_{it}^p + M_{it}^c} \frac{\omega_{k,z}(s_{ik}+s_{iz})}{2}}{\omega_0 + \sum_{k,z=1,1}^{M_{it}^p + M_{it}^c} \omega_{kz}}. \tag{A17}$$

1.A.3 Bayesian Updating with Desired Correlation Structure

To construct Ω in each time period t for each user i, I need to account for all signals from the personal and community smy ces across multiple time periods, i.e.,

$$\begin{aligned}\psi_{it} = [&S_{i1,1}^p, S_{i1,2}^p, \ldots, S_{i1,N_{i1}^p}^p, S_{i2,1}^p, S_{i2,2}^p, \ldots, S_{i2,N_{i2}^p}^p, \ldots, S_{it,1}^p, S_{it,2}^p, \ldots, S_{it,N_{it}^p}^p, \\ &S_{i1,1}^c, S_{i1,2}^c, \ldots, S_{i1,N_{i1}^c}^c, S_{i2,1}^c, S_{i2,2}^c, \ldots, S_{i2,N_{i2}^c}^c, \ldots, S_{it,1}^c, S_{it,2}^c, \ldots, S_{it,N_{it}^c}^c]\end{aligned} \tag{A18}$$

where ψ_{it} is a vector of size $(\sum_{k=1}^t N_{ik}^p + \sum_{k=1}^t N_{ik}^c) \times 1$.

Ω is composed of several blocks, each representing the interactions between signals of the same type within the same day, different types within the same day, and signals across different days.

- The diagonal elements represent the precision of personal and community signals and are denoted by ω_p and ω_c respectively.
- The off-diagonal elements within the same day correspond to the partial correlations between signals of the same type within the same day, given by λ_p for personal signals and by λ_c for community signals. The correlation across signal types within the same day is denoted as λ_{pc}.
- The off-diagonal elements across different days decay according to the decay rates δ_p, δ_c, and δ_{sc}. The decay is applied exponentially based on the time difference, e.g., for two personal signals between different days t and t', the element corresponding to the partial correlation between signals is given by $\delta_p^{|t-t'|}\lambda_p$.

The Ω matrix is of size $(\sum_{k=1}^{t} N_{ik}^{p} + \sum_{k=1}^{t} N_{ik}^{c}) \times (\sum_{k=1}^{t} N_{ik}^{p} + \sum_{k=1}^{t} N_{it}^{c})$ with rows and columns corresponding to signals of ψ_{it}.

$$\Omega_t = \begin{bmatrix} \Omega_p & \Omega_{pc}^T \\ \Omega_{pc} & \Omega_c \end{bmatrix} \tag{A19}$$

Ω_p and Ω_c correspond to the precision of signals and the within-smy ce correlations among signals and Ω_{pc} corresponds to the correlations between signals from different smy ces. Ω_p and Ω_c have a $t \times t$ structure of smaller blocks with rows and columns corresponding to time periods $1, \ldots, t$:

$$\Omega_{|\omega,\lambda} \in \{\Omega_{p|\omega_p,\lambda_p}, \Omega_{c|\omega_c,\lambda_c}\} = \begin{bmatrix} \Omega_{11} & \delta^1\Omega_{12} & \cdots & \delta^t\Omega_{1t} \\ \delta^1\Omega_{21} & \Omega_{22} & \cdots & \delta^{t-1}\Omega_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ \delta^t\Omega_{t1} & \delta^{t-1}\Omega_{t2} & \cdots & \Omega_{tt} \end{bmatrix} \tag{A20}$$

The diagonal Ω_{kk} blocks capture the precision and correlation of the $N_{ik} \in \{N_{ik}^p, N_{ik}^c\}$ signals received from a smy ce in each time period. Thus, each Ω_{kk} is a $N_{ik} \times N_{ik}$ matrix with diagonal elements ω and off-diagonal elements λ. The off-diagonal $\Omega_z k$ matrix blocks are $\lambda \dot{J}_{N_{iz} \times N_{ik}}$ matrices, J being an all-ones matrix:

$$\Omega_{kk} = \begin{bmatrix} \omega & \lambda & \cdots & \lambda \\ \lambda & \omega & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \omega \end{bmatrix}_{N_{ik} \times N_{ik}} , \quad \Omega_{zk} = \begin{bmatrix} \lambda & \lambda & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \lambda \end{bmatrix}_{N_{iz} \times N_{ik}} \tag{A21}$$

Ω_{pc} captures the correlation between signals coming from different smy ces. This matrix also consists of $t \times t$ blocks, with blocks representing the partial correlation between signals of different types within and across time periods. The correlation between two signals is proportional to λ_{pc}, decreasing at an exponential rate of δpc as the time difference between the two signals increases. Formally, the matrix consists of blocks matrices of size $N_{iz}^p \times N_{ik}^c$ for personal signals of time period z and community signals of time period k with all elements equal to $\delta_{pc}^{|z-k|}\lambda_{pc}$:

$$\Omega_{pc} = \lambda_{pc} \begin{bmatrix} J_{N_{i1}^p \times N_{i1}^c} & \delta_{pc}^1 J_{N_{i1}^p \times N_{i2}^c} & \cdots & \delta_{pc}^t J_{N_{i1}^p \times N_{it}^c} \\ \delta_{pc}^1 J_{N_{i2}^p \times N_{i1}^c} & J_{N_{i2}^p \times N_{i2}^c} & \cdots & \delta_{pc}^{t-1} J_{N_{i2}^p \times N_{it}^c} \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{pc}^t \; J_{N_{it}^p \times N_{i1}^c} & \delta_{pc}^{t-1} J_{N_{it}^p \times N_{i2}^c} & \cdots & J_{N_{it}^p \times N_{it}^c} \end{bmatrix} \tag{A22}$$

I now proceed to calculate the posterior mean and precision given the structure for Ω, beginning with ω_t. First, I calculate the sum of elements in Ω_p and Ω_c. I illustrate the calculations for Ω_p; the process for Ω_c is analogous. The sum of the diagonal elements Ω_{kk} in Ω_p can be expressed as:

$$N_{ik}^p \times \omega_p + (N_{ik}^{p\,2} - N_{ik}^p) \times \lambda_p \tag{A23}$$

and the sum of the off-diagonal elements Ω_{zk} as:

$$N_{iz}^p \times N_{ik}^p \times \lambda_p. \tag{A24}$$

Therefore, the total sum for Ω_p is given by

$$\omega_p \sum_{k=1}^{t} N_{ik}^p + \lambda_p \sum_{k=1}^{t} (N_{ik}^{p\,2} - N_{ik}^p) + \lambda_p \sum_{k,z=1,k\neq z}^{t} \delta_p^{|k-z|} N_{iz}^p \times N_{ik}^p \tag{A25}$$

The total sum for Ω_c follows the same structure. For the cross-term Ω_{pc}, the sum of elements is straightforward:

$$\lambda_{pc} \sum_{k,z=1,1}^{t,t} \delta_{pc}^{|k-z|} N_{iz}^p \times N_{ik}^c. \tag{A26}$$

Consequently, the posterior precision ω_t can be calculated as:

$$
\begin{aligned}
\omega_t = \omega_0 &+ \omega_p \sum_{k=1}^{t} N_{ik}^p + \omega_c \sum_{k=1}^{t} N_{ik}^c \\
&+ \lambda_p \sum_{k=1}^{t} (N_{ik}^{p\,2} - N_{ik}^p) + \lambda_p \sum_{k,z=1,k\neq z}^{t} \delta_p^{|k-z|} N_{iz}^p \times N_{ik}^p \\
&+ \lambda_c \sum_{k=1}^{t} (N_{ik}^{c\,2} - N_{ik}^c) + \lambda_c \sum_{k,z=1,k\neq z}^{t} \delta_c^{|k-z|} N_{iz}^c \times N_{ik}^c \\
&+ 2\lambda_{pc} \sum_{k,z=1,1}^{t} \delta_{pc}^{|k-z|} N_{iz}^p \times N_{ik}^c \\
&= \omega_0 + (\omega_p - \lambda_p) \sum_{k=1}^{t} N_{ik}^p + (\omega_c - \lambda_c) \sum_{k=1}^{t} N_{ik}^c \\
&+ \lambda_p \sum_{k,z=1}^{t} \delta_p^{|k-z|} N_{iz}^p \times N_{ik}^p \\
&+ \lambda_c \sum_{k,z=1}^{t} \delta_c^{|k-z|} N_{iz}^c \times N_{ik}^c \\
&+ 2\lambda_{pc} \sum_{k,z=1,1}^{t} \delta_{pc}^{|k-z|} N_{iz}^p \times N_{ik}^c
\end{aligned}
\tag{A27}
$$

Next, I turn to the calculation of the posterior mean μ_t. First, I calculate the second term for Ω_p; the calculations for Ω_c will be similar. Let $\mathbb{S}_{ik}^p = \sum_{r=1}^{N_{ik}^p} s_{ik,r}^p$, the sum of the personal signal values on day k. For the diagonal blocks Ω_{kk}^p at each time k in Ω_p, I have:

$$
\begin{aligned}
\sum_{r,q=1,1}^{N_{ik}^p} \omega_{rq}(s_{ik,r}^p + s_{ik,q}^p) =& 2\omega_p s_{ik,1}^p + \lambda_p(s_{ik,1}^p + s_{ik,2}^p) + \lambda_p(s_{ik,1}^p + s_{ik,3}^p) + \cdots + \lambda_p(s_{ik,1}^p + s_{ik,N_{ik}^p}^p) + \\
& \lambda_p(s_{ik,2}^p + s_{ik,1}^p) + 2\omega^p s_{ik,2}^p + \lambda_p(s_{ik,2}^p + s_{ik,3}^p) + \cdots + \lambda_p(s_{ik,2}^p + s_{ik,N_{ik}^p}^p) + \\
& \quad \vdots \\
& \lambda_p(s_{ik,N_{ik}^p}^p + s_{ik,1}^p) + \lambda_p(s_{ik,N_{ik}^p}^p + s_{ik,2}^p) + \cdots + 2\omega_p s_{ik,N_{ik}^p}^p \\
=& 2\omega_p \mathbb{S}_{ik}^p + (N_{ik}^p - 2)\lambda_p \mathbb{S}_{ik}^p + N_{ik}^p \lambda_p \mathbb{S}_{ik}^p \\
=& 2\omega_p \mathbb{S}_{ik}^p + 2(N_{ik}^p - 1)\lambda_p \mathbb{S}_{ik}^p. \qquad \text{(A28)}
\end{aligned}
$$

For the off-diagonal blocks Ω_{zk}^p in Ω_p, I have:

$$
\begin{aligned}
\sum_{r=1}^{N_{iz}^p} \sum_{q=1}^{N_{ik}^p} \omega_{rq}(s_{iz,r}^p + s_{ik,q}^p) =& \lambda_p(s_{iz,1}^p + s_{ik,1}^p) + \lambda_p(s_{iz,1}^p + s_{ik,2}^p) + \cdots + \lambda_p(s_{iz,1}^p + s_{ik,N_{ik}^p}^p) + \\
& \lambda_p(s_{iz,2}^p + s_{ik,1}^p) + \lambda_p(s_{iz,2}^p + s_{ik,2}^p) + \cdots + \lambda_p(s_{iz,2}^p + s_{ik,N_{ik}^p}^p) + \\
& \quad \vdots \\
& \lambda_p(s_{iz,N_{iz}^p}^p + s_{ik,1}^p) + \lambda_p(s_{iz,N_{iz}^p}^p + s_{ik,2}^p) + \cdots + \lambda_p(s_{iz,N_{iz}^p}^p + s_{ik,N_{ik}^p}^p) \\
=& N_{ik}^p \lambda_p \mathbb{S}_{iz}^p + N_{iz}^p \lambda_p \mathbb{S}_{ik}^p. \qquad \text{(A29)}
\end{aligned}
$$

Thus, the sum $\sum\limits_{zk=1,1} \frac{\omega_{zk}(s_{iz}+s_{ik})}{2}$ for Ω_p is:

$$
\omega_p \sum_{k=1}^{t} \mathbb{S}_{ik}^p + \lambda_p \sum_{k=1}^{t} (N_{ik}^p - 1)\mathbb{S}_{ik}^p + \frac{\lambda_p}{2} \sum_{k,z=1,k\neq z}^{t} \delta_p^{|k-z|} \left(N_{ik}^p \mathbb{S}_{iz}^p + N_{iz}^p \mathbb{S}_{ik}^p\right) \qquad \text{(A30)}
$$

.

Because of symmetry, $\sum\limits_{k,z=1,k\neq z}^{t} \delta_p^{|k-z|} \left(N_{ik}^p \mathbb{S}_{iz}^p + N_{iz}^p \mathbb{S}_{ik}^p\right) = 2 \sum\limits_{k,z=1,k\neq z}^{t} \delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p$. Breaking the second sum and combining the. parts with the first and third sums will give:

$$(\omega_p - \lambda_p)\sum_{k=1}^{t}\mathbb{S}_{ik}^{p} + \lambda_p \sum_{k,z=1}^{t}\delta_p^{|k-z|}\left(N_{ik}^{p}\mathbb{S}_{iz}^{p} + N_{iz}^{p}\mathbb{S}_{ik}^{p}\right) \tag{A31}$$

.

The sum related to Ω_{pc} is given by:

$$\sum_{r=1}^{N_{iz}^{p}}\sum_{q=1}^{N_{ik}^{c}}\omega_{rq}(s_{iz,r}^{p} + s_{ik,q}^{c}) = \lambda_{pc}\sum_{k,z=1}^{t}\delta_{pc}^{|k-z|}\left(N_{ik}^{c}\mathbb{S}_{iz}^{p} + N_{iz}^{p}\mathbb{S}_{ik}^{c}\right). \tag{A32}$$

Finally, I combine these terms to calculate the posterior mean μ_t:

$$\begin{aligned}\mu_{it} = \frac{1}{\omega_t}\Bigg(&\omega_0\mu_0 + (\omega_p - \lambda_p)\sum_{k=1}^{t}\mathbb{S}_{ik}^{p} + (\omega_c - \lambda_c)\sum_{k=1}^{t}\mathbb{S}_{ik}^{c} + \\ &\lambda_p \sum_{k,z=1}^{t}\left(\delta_p^{|k-z|}N_{ik}^{p}\mathbb{S}_{iz}^{p}\right) + \\ &\lambda_c \sum_{k,z=1}^{t}\left(\delta_c^{|k-z|}N_{ik}^{c}\mathbb{S}_{iz}^{c}\right) + \\ &\frac{\lambda_{pc}}{2}\sum_{k,z=1}^{t}\left(\delta_{pc}^{|k-z|}N_{ik}^{c}\mathbb{S}_{iz}^{p}\right) + \frac{\lambda_{pc}}{2}\sum_{k,z=1}^{t}\left(\delta_{pc}^{|k-z|}N_{iz}^{p}\mathbb{S}_{ik}^{c}\right)\Bigg).\end{aligned} \tag{A33}$$

to compute the posterior more efficiently in the optimization process, I construct the terms with the form $\sum_{k,z=1}^{t}\left(\delta_p^{|k-z|}AB\right)$, using $\sum_{k,z=1}^{t}\left(\delta_p^{|k-z|}N_{ik}^{p}\mathbb{S}_{iz}^{p}\right)$ as an example by breaking it as:

$$\sum_{k,z=1}^{t}\left(\delta_p^{|k-z|}N_{ik}^{p}\mathbb{S}_{iz}^{p}\right) = \sum_{k,z=1}^{t-1}\left(\delta_p^{|k-z|}N_{ik}^{p}\mathbb{S}_{iz}^{p}\right) + \sum_{k=1}^{t}\left(\delta_p^{|t-k|}N_{ik}^{p}\mathbb{S}_{it}^{p}\right) + \sum_{k=1}^{t}\left(\delta_p^{|t-k|}N_{it}^{p}\mathbb{S}_{ik}^{p}\right) - N_{it}^{p}\mathbb{S}_{it}^{p} \tag{A34}$$

This can be simplified as:

$$= \sum_{k,z=1}^{t-1} \left(\delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p\right) + \delta_p^t \mathbb{S}_{it}^p \sum_{k=1}^{t} \left(\delta_p^{-k} N_{ik}^p\right) + \delta_p^t N_{it}^p \sum_{k=1}^{t} \left(\delta_p^{-k} \mathbb{S}_{ik}^p\right) - N_{it}^p \mathbb{S}_{it}^p \tag{A35}$$

At each time t, the $\sum\limits_{k,z=1}^{t} \left(\delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p\right)$ can be written as a function of its value at time $t-1$. For $t=1$, it will be $N_{i1}^p \mathbb{S}_{i1}^p$. Thus the sum can be re-written as

$$\sum_{k,z=1}^{t} \left(\delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p\right) = \sum_{z=1}^{t} \left(\delta_p^z \mathbb{S}_{iz}^p \sum_{k=1}^{z} \left(\delta_p^{-k} N_{ik}^p\right) + \delta_p^z N_{iz}^p \sum_{k=1}^{z} \left(\delta_p^{-k} \mathbb{S}_{ik}^p\right) - N_{iz}^p \mathbb{S}_{iz}^p \right) \tag{A36}$$

1.B Supplements for Predictive Exercises

I present the average predicted number of tipping incidences in a day and the predicted total amount of tips given by all users in a day in Figure B-1. In Figure B-1(a), the number of tip incidences is largely stable over time. I observe a similar pattern for the total amount of tips given by all users over time with a moderate increase in the last few months in Figure B-1(b)

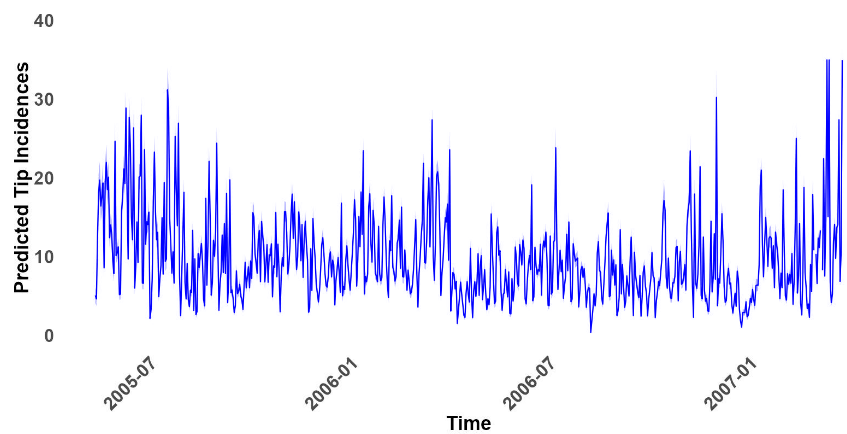

(a) Predicted Number of Tip Incidences

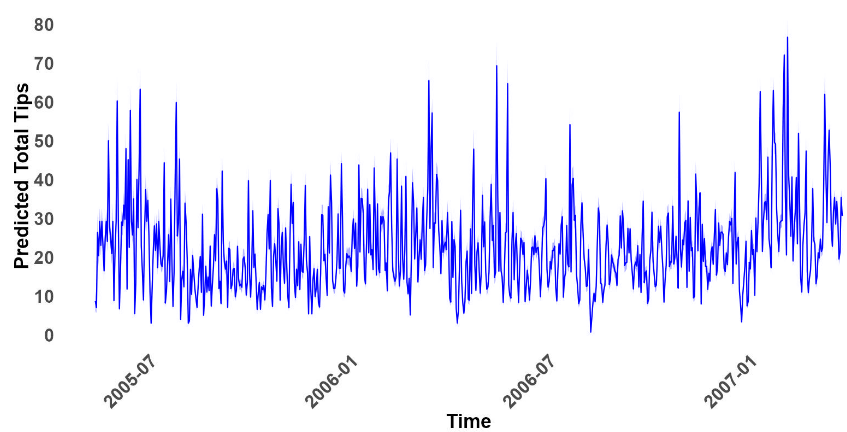

(b) Predicted Total Tips Amount

Figure B-1: Predicted Number of Tip Incidences and Total Tips

Chapter 2

The Impact of Rewards on User-Generated Content

2.1 Introduction

User-generated content (UGC) plays a crucial role for social media platforms: it attracts new users to a platform and keeps existing users engaged. The more and the more engaged users a platform has, the more ads it can show and increase its revenue. Therefore, platforms are keenly interested in increasing high-quality UGC production. To this end, many platforms have implemented a variety of rewards to encmy age users to create content. Table B-1 provides an overview of reward types implemented by a sample of well-known platforms. For example, YouTube utilizes all fmy types of rewards.[1] Several other social media platforms also use multiple types of rewards.

[1]For example, YouTube awards monetary compensation via the YouTube Short Fund to YouTube creators whose Shorts received high engagement and many views (https://blog.youtube/news-and-events/introducing-youtube-shorts-fund/). YouTube awards honorary YouTube Creator Awards plaques to YouTube channels with a large number of subscribers (https://www.youtube.com/creators/how-things-work/get-involved/awards/). On YouTube, users can like content and they can purchase access to the YouTube Super Chat. The proceeds from YouTube Super Chat access purchases go to the content creator (https://support.google.com/youtube/answer/7288782?hl=en).

Table B-1: Examples of Reward Structures (as of March 21st, 2025)

Provider	PLATFORM		PEERS	
Type	Monetary	Non-Monetary	Monetary	Non-Monetary
Examples	Compensation	Badge	Tip	Like
YouTube	X	X	X	X
Instagram	X		X	X
TikTok	X	X	X	X
StackExchange		X		X
StackOverflow		X		X
Twitch	X	X	X	
Wikipedia		X		X
Reddit	X	X	X	X
Goodreads		X		X
CycleChat		X		X
BoardGameGeek	X	X	X	X

This variety of rewards used by platforms raises several important questions. First, rewards can either be given by the platform itself or by other users. Does it matter who provides the reward? On the one hand, economic theory suggests that the giver of a reward should not matter, e.g., a $1 reward from the platform has the same monetary value as a $1 reward from peers. On the other hand, users might value receiving $1 from the platform differently than receiving $1 from peers, e.g., due to the desire for social recognition. The answer to this question is important for platforms which face a trade-off: should they spend monetary resmy ces on directly rewarding content creators or on creating a culture in which users reward each other?

Second, are monetary and non-monetary incentives equally effective in encmy aging UGC? On the one hand, non-monetary rewards enhance motivations related to social benefits of voluntary contributions, such as improved reputation and social status (Hennig-Thurau et al. 2004, Roberts, Hann, and Slaughter 2006, Toubia and Stephen 2013). On the other hand, economic theory suggests that rational individuals aim to maximize their utility and thus monetary incentives can impact content production by triggering financial motivations for content generation (Hennig-Thurau et al. 2004). However, monetary incentives may also weaken the status-enhancing effects of prosocial behaviors and crowd out users' non-monetary motivations (e.g., Qiao et al. 2020, Liu and Feng 2021). While existing evidence supports both types of rewards having significant effects on volume and quality of UGC, most studies have only examined one type of reward, leaving a gap in my understanding of their relative effectiveness.

Third, only a small fraction of users is typically very active on a platform and responsible for creating the majority of content. As these high-volume contributors drive much of the platform's overall engagement, a key challenge for platforms is ensuring that top contributors remain engaged, while also finding ways to encmy age casual users to become more active. It is unclear how different types of rewards influence these two groups of users. Can rewards transform casual users into top contributors or do they primarily reinforce existing engagement patterns? To answer these questions, I jointly measure the impact of multiple rewards, focusing on fmy common types of rewards which differ in their smy ce (awarded by platform and by other users) and prize (monetary and non-monetary).

My data come from an online board game platform called BoardGameGeek.com (BGG). For a random sample of users, I observe different types of UGC they created over a time period of ten years and all rewards they received for their UGC.[2] The platform rewards users with badges and monetary compensation in its virtual currency called GeekGold (GG) for creating UGC.[3] Other users can also reward the focal user for UGC with likes and tips. To summarize, users can receive monetary and non-monetary rewards from both the platform and other users.

One of the challenges of estimating the effects of rewards lies in their endogeneity: users do not randomly receive rewards; they receive rewards for previously produced UGC. Therefore,

[2] I observe five types of UGC: threads, reviews, game session reports, replies/comments, game lists.

[3] GeekGold is a platform-specific virtual currency. Users can buy GG from other users and sell GG to other users in private transactions. Users can also buy board games from other users using GG.

previous literature has mostly relied on experimental variation to measure the effects of different rewards (Burtch et al. 2018, 2022; Huang, Kaul, and Narayanan 2022). However, because of cost and complexity concerns, the experimental approach has mostly been used to study the effects of a *single* type of reward. This makes a comparison of the effects of multiple rewards challenging. I take a different approach in this paper: I address endogeneity concerns taking advantage of my unique data which allow us to include a rich set of controls and fixed effects, and measure the effects of fmy types of rewards in one empirical setting.

I quantify the effects of rewards a user received during the prior three days on the creation of UGC on the focal day using log-log linear regressions. I examine two aspects of the created UGC: its quantity and its quality. UGC quantity refers to the number of posts of a UGC type. I measure quality of the text in terms of its length, which has been shown to have a strong correlation with several dimensions of text quality (Blumenstock, 2008; Mudambi and Schuff, 2010), and its informativeness, which captures the amount of meaningful content in the text relative to its length.

My findings show that monetary peer rewards (tips) are particularly effective in increasing both the volume and quality of posts, while monetary rewards from the platform (compensation) encmy age longer and informative submissions at the expense of overall posting frequency. Furthermore, non-monetary peer rewards (likes) have a positive, albeit small effect on the number of posts. Meanwhile, the effect of non-monetary platform rewards (badges) follows a U-shaped pattern, with users briefly pausing after reaching a milestone and then becoming more active once they get closer to the next goal. In terms of the relative effects of monetary and non-monetary rewards, my results indicate that the effect of receiving 1 GG in tips on UGC quantity is equivalent to the effect of receiving 6 likes.[4] However, taking into account that a post receives an average tip of 0.09 GG and 2.8 likes, the effect of a post's likes becomes 2.5 times as large as of its tips. This is an interesting result because likes, in contrast to tips, are "free," i.e., users do not have to spend GG on them.

[4]During my study period of 2010 to 2020, the exchange rate between GG and US $ varied between 0.01 to 0.07, i.e., 1 GG was worth ¢1 - ¢7.

A second important aspect relates to how different groups of users respond to these rewards. I investigate whether top contributors, who produce a large share of platform content and are of special interest to managers, react differently to rewards than casual participants, and whether a platform can guide casual users toward becoming top contributors through the usage of rewards. Interestingly, top contributors and casual users react similarly to rewards at the beginning of their membership. Over time, when differences emerge, top contributors either become less sensitive to rewards or respond more negatively to them compared to casual users. These findings are consistent with platform-based compensation crowding out intrinsic motivation for highly active individuals, leading to a reduction in the total number of posts they create, even though the quality of each individual post may rise. At the same time, likes also lose much of their impact on top contributors, suggesting that non-monetary peer rewards may need to evolve over a user's lifecycle if they are to remain effective. Based on these results, I conclude that extrinsic rewards alone cannot transform a casual user into a top contributor; more fundamental drivers, such as intrinsic interest, are likely at play.

The contribution of this paper is three-fold. First, I add to managers' and academics' understanding of the relative effectiveness of fmy different types of rewards which vary in their nature (monetary vs. non-monetary) and giving entity (peers vs. platform). Unlike most prior studies, which examine only one or two types of rewards, this paper provides a comprehensive comparison, addressing a gap in the literature. By offering a holistic perspective, my findings provide guidance to managers on how to design reward systems to achieve desired UGC goals.

Second, I investigate the effects of the fmy types of rewards not only on the quantity but also on the quality of UGC. By examining different dimensions of text quality, I contribute to the literature on UGC by highlighting that certain rewards (e.g., compensation) can elevate content depth at the cost of posting volume, while others (e.g., tips) drive frequent activity without necessarily undermining the substantive value of each post. These insights underscore the complexity of UGC creation processes and careful weighing of the trade-offs between maintaining a high volume of user submissions and fostering detailed, in-depth contributions in designing a reward system that is tailored to a platform's short- and long-term objectives.

And third, I show that, while rewards can refine or boost existing engagement, they are not effective tools managers can use to convert casual users into top contributors. Moreover, top contributors tend to be more susceptible to crowding out from platform-based monetary rewards, underscoring the potential pitfalls of one-size-fits-all incentive schemes. Furthermore, non-monetary peer rewards such as likes, that can act as social acknowledgment signals, often lose their impact on top contributors, emphasizing the need for more tailored and effective incentive strategies.

The remainder of this paper is organized as follows: In the next section, I review the relevant literature. In Section 2.3, I introduce and describe my data. I present my model in Section 2.4 and discuss the results in Section 2.5. In the following section, I study the heterogenous effects of rewards. In Section 2.7, I review the robustness checks and conclude in Section 2.8.

2.2 Relevant Literature

In this section, I review the relevant streams of literature on UGC, online rewards, and special interest communities and delineate my research vis-à-vis findings from previous research.

UGC plays an important role in influencing consumer adoption decisions (Li and Hitt, 2008; Zhang and Godes, 2018b; Ameri, Honka, and Xie, 2019) and driving sales (Chevalier and Mayzlin, 2006; Chen, Wang, and Xie, 2011; Moe and Trusov, 2011). It also serves as a key smy ce of entertainment and engagement for online platform users (Chevalier and Mayzlin, 2006; Leung, 2009; Yang, Ren, and Adomavicious, 2019). Given its significant downstream impact, extensive research has explored the external factors that drive UGC creation, including social norms (Burtch et al., 2018), financial incentives (Cabral and Li, 2015; Burtch et al., 2018; Khern-am nuai, Kannan, and Ghasemkhani, 2018), rewards (Gallus, 2017; Burtch et al., 2022), performance feedback (Huang et al., 2019), community commitment (Bateman, Gray, and Butler, 2011), social recognition (Wasko and Faraj, 2005; Chen et al., 2024), and audience size (Zhang and Zhu, 2011).

A growing body of literature examines how platforms use rewards to shape user contributions and engagement. Platforms themselves can directly influence UGC creation through monetary compensation or symbolic recognitions given by the platform (e.g., Goes, Guo, and Lin 2016; Gallus 2017; Khern-am nuai, Kannan, and Ghasemkhani 2018; Liu and Feng 2021; Qiao and Rui 2023). Findings on the effect of monetary compensation are mixed: some studies show financial rewards increase content quantity and quality (Gneezy and Rustichini, 2000; Burtch et al., 2018; Kuang et al., 2019; Yu, Khern-am nuai, and Pinsonneault, 2022), while others suggest that monetary rewards can crowd out intrinsic motivation and status-seeking behaviors and negatively impact users' content generation (Khern-am nuai, Kannan, and Ghasemkhani, 2018; Qiao et al., 2020; Liu and Feng, 2021). For instance, Kuang et al. (2019) find that financial rewards in paid knowledge-sharing platforms encmy age voluntary contributions and social engagement, and Yu, Khern-am nuai, and Pinsonneault (2022) show performance-based rewards lead to more and higher-quality reviews. On the other hand, Khern-am nuai, Kannan, and Ghasemkhani (2018) and Qiao et al. (2020) find that financial rewards result in shorter and lower-quality reviews. Unlike monetary rewards, which provide immediate economic benefits, non-monetary rewards encmy age UGC by serving as status signals, providing role clarity, and motivating of long-term engagement (Gallus, 2017; Hanson, Jiang, and Dahl, 2019; Lu, Xie, and Chen, 2023). For instance, Gallus (2017) and Lu, Xie, and Chen (2023) find that symbolic medals enhance contribution, while Hanson, Jiang, and Dahl (2019) show that labels and badges clarify roles, increasing UGC creation.

Beyond directly providing rewards themselves, platforms can also encmy age peers to reward UGC creation either monetarily, e.g., through tipping (Cabral and Li, 2015; Burtch et al., 2022), or non-monetarily, e.g., through likes and symbolic acknowledgments (Zhang and Zhu, 2011; Restivo and van de Rijt, 2014; Huang, Hong, and Burtch, 2017; Gallus, Jung, and Lakhani, 2020; Mummalaneni, Yoganarasimhan, and Pathak, 2024). Through field experiments, Burtch et al. (2022) find that monetary peer rewards increase content length and frequency on Reddit, particularly among newer users, and Restivo and van de Rijt (2014) find that receiving peer recognition boosts productivity among already active contributors on Wikipedia. Zhang and Zhu (2011) show that contributors reduced their UGC production on Chinese Wikipedia when the number of users decreased due to an exogenous shock.

While prior research has mostly examined reward types individually, few studies compare their effects. Some find that both monetary and non-monetary incentives enhance content generation, with monetary rewards primarily influencing short-term engagement, while non-monetary rewards, such as social approval and status signals, play a greater role in sustaining long-term contributions (Roberts, Hann, and Slaughter, 2006; Burtch et al., 2018; Woolley and Sharif, 2021; Wang, Qiu, and Xu, 2024). However, other studies suggest that monetary rewards can negatively impact contributions, though social rewards may help mitigate these declines (Sun, Dong, and McIntyre, 2017; He et al., 2023). Despite these insights, many questions related to the relative effectiveness of different reward types remain unanswered. I contribute to addressing this gap by systematically comparing the effects of fmy types of rewards – monetary vs. non-monetary and platform-based vs. peer-based – on both the quantity and quality of UGC, offering a more comprehensive understanding of how different incentives shape user contributions.

Lastly, my paper is related to the literature on special interest communities where interactions are based on shared enthusiasm for a specific consumption activity (Kozinets 1999). Special interest communities help people feel more connected, and internet users increasingly prefer special interest online communities over general social media, such as Facebook or Instagram.[5] Recent studies have examined user behavior in different special

[5]https://blog.gwi.com/chart-of-the-week/online-communities/

interest communities. For example, Hendricks and Sorensen (2009) study users' adoption of new music online, and show that new album releases on the platform lead to a substantial and permanent increase in the sales of old albums of the same artist. Zhang and Godes (2018b) study Goodreads, and show that, with sufficient experience, having more ties leads to better decisions. Nevskaya and Albuquerque (2019) study the role of rewards on users' consumption of a game in a massive online video game platform. They find that improving reward schedules and imposing time limits leads to shorter usage sessions and longer game subscriptions. Ameri, Honka, and Xie (2023b) study how strangers become friends within an evolving online social network in an online anime-watching platform, and how this evolving network impacts users' content generation and vice versa. I contribute to this stream of literature by examining how different reward types affect users' content generation in a board game related online community.

2.3 Data

My data come from Boardgamegeek.com (BGG). This website is a consumption-related online community revolving around board games. It was established in 2000 and has become the largest online database for board games as well as the largest online community for board game fans with over 4 million users worldwide. Figure B-1 shows the number of users joining BGG over time.

Figure B-1: Number of Users Joining BGG Over Time

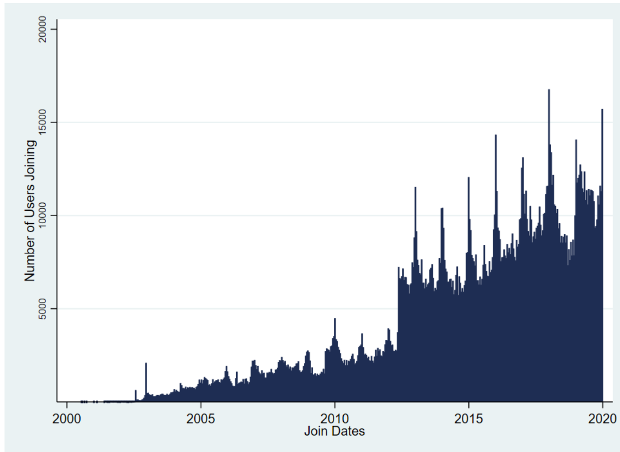

An important aspect of BGG is that all of its content is created by users. Users provide detailed information about new and existing games via reviews and game session reports,[6] curate themed lists of games,[7] and also engage in a variety of conversations with other users in the discussion forum section of the website.[8]

BGG utilizes a platform-specific virtual currency called GeekGold (GG) to reward users for their contributions. Users can receive 1 - 5 GG as compensation for writing a review, session report, or starting a new discussion thread.[9] Users can only earn GG through contributions and cannot directly buy GG from the platform.[10] Users can also earn GG in form of tips from other users for the content they create. Users can tip any amount they want. Aside from tipping, users can spend their GG on virtual cosmetic items for their profile page

[6]Game session reports on BGG are user-submitted logs with detailed gameplay experiences, strategies, and outcomes from a single play session of a board game.

[7]These lists, known as GeekLists on BGG, are user-created collections of games organized around specific themes, personal rankings, or recommendations.

[8]Users can also contribute other forms of UGC such as ratings, uploading game files, and images. These forms of UGC are much less common on BGG and I therefore focus on threads, reviews, session reports, replies (or comments), and game lists.

[9]All reviews/reports go through a process in which other volunteer users vote to approve a review/report and recommend an amount between 1 - 5 GG to award to the content creator. I investigated this process and found that the approval of a review/report is a formality to prevent inappropriate content and takes less than one day. The average amount recommended by other users determines the compensation amount the content creator receives for her contribution.

[10]The platform rewards users who donate money to BGG by giving them GG. Some users may also buy GG from other users privately. However, both donations and GG purchases are not common.

or buy board games from peers.

Users also receive badges for writing a certain number of posts of a type. Reviews, session reports, and lists have each their own badges and milestones, while there are also broader badges that encompass any forum post, which includes threads, reviews, session reports, and replies. The badge system is set up in a way that a user has to produce increasingly more content to reach the next milestone. For example, a user has to write 5 reviews to earn the first badge, 45 additional reviews to earn the second badge, etc. A list of the badges and associated milestones is available in Web Appendix A. Lastly, users can also react to the content produced by others by giving "likes." Figure B-2 shows a thread in which the content creator received likes and tips from other users.

Figure B-2: Example of a Post for Which the Creator Received Tips and Likes

Table B-2 summarizes the available rewards for threads, reviews, game session reports, replies, and lists. Non-monetary peer rewards (likes), monetary peer rewards (tips), and non-monetary rewards (badges) from the platform can be given/earned for all UGC types. Monetary rewards from the platform (compensation) are only awarded for threads, reviews, and session reports.

Table B-2: Available Rewards

	MONETARY		NON-MONETARY	
	Platform	Users	Platform	Users
UGC Types	Compensation	Tip	Badge	Like
Threads	0-5	0.001+	20 Levels	1+
Reviews	1-5	0.001+	6 Levels	1+
Session Reports	1-5	0.001+	6 Levels	1+
Replies	-	0.001+	20 Levels	1+
Lists	-	0.001+	3 Levels	1+

2.3.1 Data Collection and Cleaning

I collected data containing all activities of a random sample of 100,000 users from their join date until August 19th, 2020. The data for each user include details of all the content the user created and the rewards she received for each piece of content.

I took the following steps to construct my final estimation sample. First, to ensure a minimum level of activity, I focus on users with more than 50 contributions during their entire membership. To exclude platform administrators, who create a lot of content, I exclude users with more than 2,000 contributions per year. Excluding very inactive and very active users left us with 47,988 users. Second, I drop users who did not create any UGC of any type after Jan 1st, 2020. I condition on at least one UGC contribution after Jan 1st, 2020, to only keep users who are still active platform members. Otherwise, if a user did not create any UGC, I cannot distinguish between the user having left the website and the user still being an active member but deciding not to create any content.[11] My final sample contains 16,688 users with 42,541,246 daily observations of all their activities and the rewards they received for the content they created from January 2010 to December 2019, my study period of 10 years.

My measure of UGC quantity is straightforward: it is the number of posts of each UGC

[11] I do not observe user log-ins or browsing activity.

type (i.e., initial threads, replies, reviews, session reports, and lists) a user made on a day. For UGC quality, I assess both the length and informativeness of the content. The length of a post, measured by the number of words, serves as a reliable proxy for quality, as prior research has demonstrated the strong relationship between length and perceived quality (Blumenstock, 2008; Mudambi and Schuff, 2010). For each day, I calculate the average number of words a user has written per post of each UGC type. To evaluate informativeness, I calculate the factual density of each post, defined as the number of factual statements, standardized by text length, i.e., the ratio of the number of factual statements to the total number of words in the text (Lex et al. 2012; Horn et al. 2013). I utilize the ReVerb Open Information Extraction framework to extract factual content and to identify informational relationships in the text (Fader, Soderland, and Etzioni 2011).

2.3.2 Data Description

By the end of my study period, on average, users had been BGG members for 8.7 years. Users, on average, write about 74 posts annually (see Table B-3). However, there is considerable variation in activity levels across users.

Table B-3: Annual UGC Creation Activity

	Mean	SD	Min	1st Quart.	Median	3rd Quart.	Max	N
Threads	3.65	8.08	0.00	0.35	1.21	3.65	322	16,688
Reviews	0.22	2.20	0.00	0.00	0.00	0.00	146	16,688
Session Reports	0.10	0.61	0.00	0.00	0.00	0.00	26	16,688
Replies	63.72	157.53	0.00	3.69	12.62	48.53	1,961	16,688
Lists	6.13	28.69	0.00	0.00	0.00	1.50	1,128	16,688
UGC	73.82	173.23	0.00	5.08	16.36	59.77	1,968	16,688

Table B-4 shows summary statistics for the fmy types of rewards. Note that users only earn a badge for a certain number of contributions (and not for each contribution). Therefore, users do not receive badges frequently as opposed to the other types of rewards and the mean number of earned badges is small.[12]

Table B-4: Annual Earned Rewards

	Mean	SD	Min	1^{st} Quart.	Median	3^{rd} Quart.	Max	N
Tips	8.66	45.41	0.00	0.00	0.61	4.49	2,052.58	16,688
Likes	168.50	622.14	0.00	5.27	20.80	87.63	22,110	16,688
Compensation	0.30	2.35	0.00	0.00	0.00	0.10	146.00	16,688
Badge	0.23	0.44	0.00	0.00	0.00	8.00	32.00	16,688

Table B-5 shows descriptive statistics for the UGC quantity and quality measures. The average number of posts of a UGC type a user writes in a day is 0.04. To put it differently, an average user writes 1.5 posts in a week across all five UGC types. Posts are, on average, 82 words long, with a median of 40 words. Finally, an average post has an informativeness score of 0.09, which means that it contains 9 factual pieces of information in every 100 words.

Table B-5: Descriptive Statistics for Quantity and Quality Measures

	Mean	Median	SD	Min	Max	N
Quantity	0.04	0.00	0.58	0.00	804.00	212,706,230
Length	82.49	40.0	219.92	0.01	112.70	4,023,673
Informativeness	0.09	0.10	0.07	0.00	1.00	4,023,673

2.4 Model

My goal is to measure the causal effects of the fmy types of rewards on the quantity and quality, i.e., length and informativeness, of UGC. In interpreting the estimated effects as causal, I will rely on the conditional independence assumption, i.e., the assumption that, conditional on a large number of observed variables and fixed effects, the error term and the focal variables are independent (Angrist and Pischke 2009).

I use the following set-up: For UGC quantity, let Y_{ijt} denote the number of posts of type $j \in \{\text{thread posts, reviews, session reports, replies, lists}\}$ user $i = 1, ..., N$ created on day $t = 1, ..., T$. For UGC quality let Y_{ijt} reflect one of the two quality dimensions of UGC posts

[12] I also report the summary statistics for rewards received per post in Web Appendix A.

of type j user i wrote on day t. I separately estimate the models for the three dependent variables using log-log linear regressions with the following specification:

$$Y_{ijt} = \beta_1 Tips_{ijt} + \beta_2 Likes_{ijt} + \beta_3 Compensation_{ijt} + \beta_4 Badge_{ijt} + \beta_{5j} C_{ijt} + \alpha_{it} + \gamma_{ijt} + \epsilon_{ijt} \quad \text{(B1)}$$

I operationalize the fmy reward types as follows: $Tips_{ijt}$ is the amount of tips (in GG) user i received from other users for UGC type j in the three days prior to day t, i.e., days $t-3$ to $t-1$. I exclude the tips user i received on day t because I cannot determine whether the reward was received before new content was produced that day and, as a result, whether receiving the reward impacted user i's behavior. I include tips received up to three days prior to day t to account for potential lingering effects of receiving rewards as well as for the possibility that user i may not have seen the reward immediately.[13] The variables $Likes_{ijt}$, $Compensation_{ijt}$, and $Badge_{ijt}$ are defined similarly: $Likes_{ijt}$ is the number of likes, $Compensation_{ijt}$ is the compensation amount (in GG), and $Badge_{ijt}$ is the number of badges user i received for UGC type j in the three days prior to day t.

There are several concerns related to the endogeneity of rewards. I address them by including a set of control variables denoted by C_{ijt}, user-day fixed effects α_{it}, and user-UGC-year fixed effects γ_{ijt} in equation (1). In the following, I discuss the endogeneity concerns and how the control variables and fixed effects address them. First, a user does not randomly receive peer rewards, i.e., tips and likes from other users. A user can only receive a peer reward if she made a post in the past. Relatedly, a user who wrote multiple posts in the past is more likely to receive a peer reward than a user who wrote one post. I address these concerns by including five control variables that contain the cumulative number of posts a user has ever written and has written in the past 365/30/7/3 days.

Additionally, some posts might generate many peer rewards, while others do not. To account for this, I include control variables related to the quality and timeliness of posts

[13] I test the robustness of my results regarding the three-day time window by re-estimating my models using one-day and seven-day time windows. The results are robust and presented in Web Appendix D available at `https://minaameri.com/incentives-appendix`.

users have written in the prior three days. Following the literature on content quality, the five quality control variables capture the following aspects of past posts: length, complexity, politeness, linked content, and informativeness (details on the construction of these variables are provided in Web Appendix A). The two control variables related to the timeliness of posts consist of the cumulative number of posts ever written and the number of posts written during the past 30 days by all users on topics the focal individual has posted within the past three days. These control variables capture the existing volume of information and the level of interest in a topic at that time.

A second concern is the non-random timing of peer rewards: on this platform, users commonly receive peer rewards within the first few days after publishing a post. Older posts rarely receive peer rewards. I address this concern by including individual-day fixed effects α_{it}. These fixed effects control for differences in received peer rewards across days for each user. Thus, the identifying variation for the effects of peer rewards is the within-day variation for each user.[14]

Third, a user does also not randomly receive platform rewards, i.e., compensation and badges. When writing a game review or session report, the user knows that she will receive compensation from the platform.[15] However, depending on how the review or report is evaluated by a few random users, the amount of the awarded compensation varies. Although users receive compensation for all their reviews and session reports, receiving compensation for a thread post requires additional steps: The user has to nominate her thread post and the thread post has to pass an evaluation by a few other users. Thus, receiving a reward

[14] I have more than 3.5 million individual-days in which a user wrote a post in one of the five UGC categories, allowing us to identify the effects of rewards despite using granular individual-day fixed effects. I test the robustness of my results by including less granular individual-week instead of individual-day fixed effects. The results are presented in Web Appendix D available at `https://minaameri.com/incentives-appendix`. I also estimated my model with neither the individual-day nor individual-UGC-year fixed effects. The results are also shown in Web Appendix D and are directionally robust.

[15] Formally, reviews have to go through GeekModding, a process in which other users read the posts, approve them, and suggest a compensation reward amount, to receive compensation. However, in practice, all reviews following basic platform guidelines get approved and compensated. The user receives a compensation reward within the allowable range that equals the average compensation amount suggested by users who read her post in GeekModding. GeekModding is fast: reviews get approved and published within a day.

for a thread post is not guaranteed. Furthermore, for all three UGC types, I control for the quality of the post that received the reward, allowing the additional variation in awarded compensation to help identify its effect.

And lastly, a user knows when she has written enough posts to earn the next badge.[16] I follow Goes, Guo, and Lin (2016) in addressing this concern: I include variables that capture the progress towards the next badge in terms of the remaining number of posts needed to reach the next milestone. Since previous literature suggests non-linear effort exertion to reach hierarchical milestones (Lal and Srinivasan, 1993; Goes, Guo, and Lin, 2016), I also include the square of the progress variables. As described in Section 2.3, the badge system is set up in a way that reaching the next badge gets increasingly difficult, i.e., a user has to produce more and more content to earn the next badge. This implies that the number of remaining posts needed to reach the next milestone is not comparable across badges since the same number can imply different completion levels. Therefore, I estimate separate coefficients for each badge.

In addition to the aforementioned variables, C_{ijt} also contains the following controls. To account for unobserved factors that may prevent a user from contributing to a specific type of UGC until day t, e.g., inexperience, I include a dummy variable that indicates if user i has ever produced any content of type j before day t. I also control for the rewards user i received for UGC types other than type j in the three days prior to t. Further, if a user has produced UGC of type j in the past, I control for the number of days since their last post of type j and its quadratic term to account for users engaging in conversations lasting several days.

Recall that I include user-day and user-UGC-year fixed effects, α_{it} and γ_{ijt}. They serve several purposes: they address the endogeneity concern related to the timing of rewards discussed in the previous section, and they capture the inherently heterogeneous tendency of users to create UGC as well as any unobserved, day-specific shocks that may influence user

[16] The number of posts for each UGC type is displayed on each user's personal page.

behavior. Incorporating user-day fixed effects also allows us to control for incidences when a user did not visit the platform and, as a result, did not post anything. User-UGC-year fixed effects control for unobserved tendencies of users for posting in each UGC type of content at different time periods. And lastly, ϵ_{ijt} is the error term and is assumed to follow a normal distribution.

2.5 Results

Next, I present and discuss my results displayed in Table B-6.[17] Recall that I use log-log linear regression models. Thus, the estimated coefficients can be interpreted as elasticities.

Table B-6: Effects of Rewards on UGC Quantity and Quality

	(i)	(ii)	(iii)
	Quantity	Quality	
		Length	Informativeness
Tips	0.0572***	0.0221***	0.0003
	(0.0086)	(0.0050)	(0.0002)
Likes	0.0088***	-0.0020*	0.0000
	(0.0011)	(0.0010)	(0.0000)
Compensation	-0.0841***	0.0751*	-0.0001
	(0.0147)	(0.0326)	(0.0005)
Badge	0.0914	-0.1029	-0.0030
	(0.1234)	(0.1809)	(0.0036)
Controls	Yes	Yes	Yes
Individual-Day Fixed Effects	Yes	Yes	Yes
Individual-UGC-Year Fixed Effects	Yes	Yes	Yes
Number of Observations	212,706,230	212,706,230	212,706,230
R^2	0.29	0.88	0.65

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Starting with peer rewards, both tips and likes have significant positive effects on UGC quantity. However, their effects on UGC length are directionally opposite with a significant

[17] The full set of coefficient estimates are available in Web Appendix B.

positive effect for tips and a negative effect for likes. In other words, monetary peer rewards result in more and longer posts, while non-monetary peer rewards lead to more but shorter posts. Interestingly, neither reward exhibits a significant effect on the informativeness of posts. Recall that the informativeness measure is standardized based on text length to account for varying post lengths, ensuring that longer posts do not artificially appear more informative. As a result, an insignificant effect indicates that, even when the length of a post increases or decreases, the amount of informativeness scales proportionally. Comparing the effect sizes of tips and likes on post quantity, my results indicate that the effect of receiving a 1 GG tip is equivalent to the effect of receiving 6 likes. However, although individual likes appear to have a smaller effect, their sheer volume can have a substantial impact: with a post receiving an average tip of 0.09 GG and 2.82 likes, the impact of likes is effectively 2.5 times the impact of tips.

Next, I examine my findings for platform rewards. I find that receiving monetary rewards from the platform exerts a *negative* effect on UGC quantity. This result is consistent with the literature on monetary incentives and crowding-out theory (Kreps, 1997; Khern-am nuai, Kannan, and Ghasemkhani, 2018; Qiao et al., 2020; Liu and Feng, 2021), Monetary rewards can undermine intrinsic motivation (Deci, Koestner, and Ryan, 1999) and erode the status-enhancing benefits of prosocial behaviors (Bénabou and Tirole, 2006), thereby crowding out non-monetary motivations. This crowding-out effect is especially pronounced when users routinely receive less than the maximum performance-based reward (Deci, Koestner, and Ryan, 1999), or when monetary incentives are not sufficiently large to be salient (Gneezy and Rustichini, 2000; Sun, Dong, and McIntyre, 2017), as is the case in my setting.

An alternative explanation for this finding could potentially be that external events, such as users acquiring a new game, lead to a user writing a review rather than compensation itself. However, by incorporating individual-time fixed effects, my model robustly accounts for such temporal shocks, ensuring that the observed compensation effects are not confounded with coinciding external events. Additionally, one might argue that, incentivized by the compensation, users simply take longer to accumulate or process information, thereby producing longer posts as they recover from the cognitive effort of previous contributions. Yet, my controls for the number of days since the last post and its quadratic form along with measures of cumulative posting activity over varying time windows (365, 30, 7, and 3 days) and prior post quality measures control for these recovery or information-gathering dynamics, ruling them out as alternative explanations.

Furthermore, note that crowding out is often related to a perceived loss of autonomy, i.e., when an activity once freely chosen starts feeling externally mandated (Frey and Jegen, 2001; Deci and Ryan, 2013). As such, formal monetary compensation from the platform may inadvertently make the user feel as though they are now "working for the platform," reducing the intrinsic joy of participating. Tips, on the other hand, are often perceived as gestures of genuine appreciation rather than formal, extrinsic control. Because tips typically directly come from appreciative community members, users tend to view them as voluntary acknowledgments that validate their intrinsic desires to share content (Burtch et al., 2022) rather than as impersonal payments from the platform. Therefore, tips rarely disrupt a contributor's sense of autonomy and do not undermine the enjoyment they derive from posting, but instead reinforce users' intrinsic motives.

Compensation has a positive effect on UGC length, but shows no impact on the informativeness of posts. To put it differently, when users receive monetary rewards from the platform, they tend to write less frequently, but the content they produce is of higher quality. Specifically, it is longer and contains more detailed information. This indicates that the platform is achieving its goal of encmy aging users to write higher-quality content, although in fewer numbers.[18] I conclude that monetary platform rewards are not a suitable tool to increase the quantity of produced UGC in the short-run. However, these findings do not imply that platforms should not use monetary rewards: these types of rewards incentivize

[18] I also estimated a model with UGC length operationalized as the number of words written in a day and find compensation to also have a negative effect on UGC length in that model. To put it differently, although the number of words in each post increases, the increase is not large enough to make up for the decrease in the number of words due to a decrease in UGC quantity.

different aspects of user behavior that might be desired by a platform as discussed in the next section.

Comparing the effects of tips and compensation shows that receiving 1 GG as compensation is 3.4 times more effective at encmy aging longer posts than receiving 1 GG in tips. This difference becomes even more pronounced for contents like reviews, where users, on average, receive higher compensation than tips. For example, given that reviews receive an average of 2.4 GG in compensation and 1.23 GG in tips, the overall impact of compensation on review length is more than five times that of tips.

Receiving a badge does not show any immediate effect on UGC production, neither in terms of quantity nor quality. Note that the coefficient estimates in Table B-6 measure short-term effects, capturing the effects of receiving a badge on user behavior during the three days following a badge award. However, when I examine user behavior between badges or as users approach the next badge milestone, a distinct pattern emerges: the coefficient estimates capturing content generation behavior in these in-between stages reveal a U-shaped pattern. In Figure B-3, I illustrate how the distance from the next badge milestone – across the three most common badge categories: reviews, reports, and overall posts – affects both the quantity and length of users' posts. The figure presents predicted effects for the first six badges in each category. For example, in the case of review badges, content production initially declines following the receipt of the third badge, but then sharply increases as users approach the fmy th badge, particularly when they are within 15 posts of the required 250 posts needed to receive the next badge. This pattern aligns with goal-gradient effects observed in the hierarchical incentives literature: Goes, Guo, and Lin (2016) found negative effects of badges on content generation, while Lal and Srinivasan (1993) and Oyer (1998) found similar declines in effort within salesforce compensation structures. The post-badge decline is consistent with the idea that users temporarily "take a break" or decrease effort after reaching a milestone, reinforcing the importance of designing strategies to sustain engagement within gamified platforms. On the other hand, while the continuous effect of distance from the next badge may weaken at times, badges can serve other purposes for the platform, such as enhancing role clarity and providing structure to the community (Hanson, Jiang, and Dahl, 2019).

(a) Effect on Post Quantity **(b)** Effect on Post Length

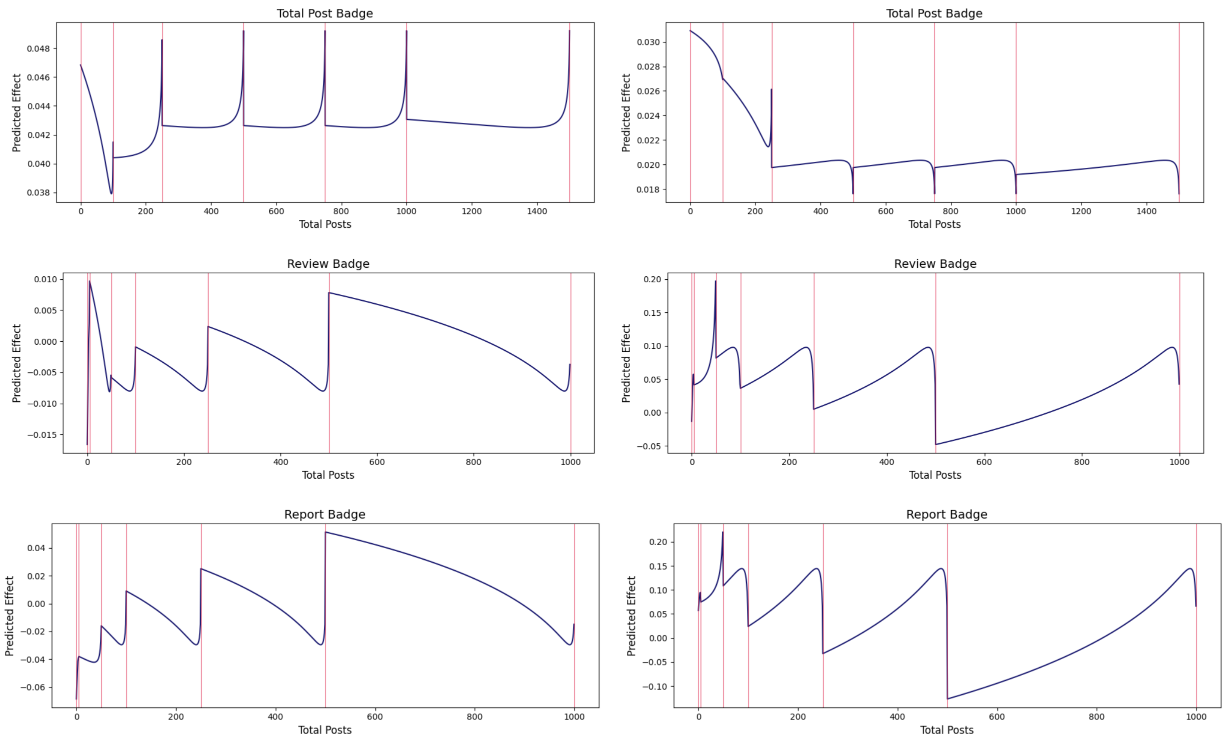

Figure B-3: Effect of Distance from the Next Badge Milestone (Red Vertical Lines Indicate Milestones)

Similar to the opposing effects of compensation on content quantity and length, I find that, while quantity follows a U-shaped trend, post length (and informativeness) exhibits an inverse-U shaped pattern as users get closer to the next badge milestone. To put it differently, after receiving a badge, users write longer (more informative) posts, but post length (informativeness) sharply declines when they get very close to the next milestone. This pattern suggests that users may be more deliberate and reflective in their content creation when they are further from the next badge but shift towards shorter, less detailed posts as they rush to complete the remaining contributions needed for the next badge. Overall, the minimal short-term impact of badges combined with the (inverse) U-shaped pattern suggest that platforms may wish to introduce more intermediate milestones or localized recognitions.

In summary, my findings highlight the distinct roles that different reward types play in driving content generation. Among the fmy reward types, monetary peer rewards (tips) stand out as the most effective reward type in increasing the number of posts on the platform, while also promoting higher-quality submissions. Although monetary compensation offered by the platform tends to decrease overall posting frequency, it emerges as the most powerful incentive for eliciting deeper, more substantive contributions. Taken together, these findings illustrate how tips serve as an effective reward for promoting frequent activity, whereas platform compensation offers stronger support for high-quality contributions. While likes have a modest impact on a per-like level, their sheer volume can create a substantial aggregate effect. Lastly, badges, despite having limited short-term effects, can function as a sustained, albeit fluctuating motivational force or serve other purposes such as providing role clarity (Hanson, Jiang, and Dahl, 2019).

2.6 Top Contributors vs. Casual Users

A common pattern in social media is that the majority of content on a platform is created by a small group of users. For example, the 90-9-1 rule suggests that 90% of users are rarely active, 9% of users are somewhat active, and 1% of users are very active. On social media, the Pareto Principle often manifests as an 80%-20% split between inactive and active users.[19] Although different platforms may exhibit variations in the exact proportions, the underlying theme remains that a small subset of users generates the majority of content. Because these high-volume creators shape much of the platform's engagement and community culture, understanding how to effectively reward them is pivotal for platform managers. Furthermore, one might wonder whether the platform can "create" top contributors, i.e., transform casual users into high-volume contributors, using rewards. I will revisit this possibility when I

[19] See, e.g., https://en.wikipedia.org/wiki/1%25_rule, https://www.nngroup.com/articles/participation-inequality/, https://stangarfield.medium.com/90-9-1-rule-of-thumb-fact-or-fiction-2377c12f3a79.

analyze how tenure interacts with user type and reward structures.

To capture the differential responses of high-activity and low-activity users, I divide users into two groups: users who are among the top 20% of individuals in my data in terms of average daily produced UGC volume over the observation period (or their membership if they joined during the observation period) are classified as *top contributors*, and I refer to all other users as *casual users*. Although top contributors comprise only one-fifth of the user base, their posts make up more than 80% of the content, underscoring their importance to the platform in my study.

I re-estimate my models separately for these two groups and report the results in Table B-7.[20] A glance at Table B-7 immediately reveals that rewards affect casual users and top contributors differently. Although both tips and likes have positive and statistically significant effects on posting activity for both groups, top contributors exhibit a larger response to tips but a smaller response to likes compared to casual users. In other words, highly active users are more influenced by monetary than by non-monetary peer rewards. A potential explanation is that top contributors, who likely have received numerous likes in the past, have become habituated to them since more exposure diminishes the effects of rewards (Thompson and Spencer, 1966; Zhang and Gao, 2016). By contrast, tips retain their effectiveness because they can be used in tangible ways, such as purchasing cosmetic badges or lottery tickets for board game prizes. Meanwhile, casual users, being less accustomed to frequent peer acknowledgment, still perceive likes as a meaningful signal of recognition.

Comparing the effect magnitudes, tips are about 3.5 times more effective than likes for casual users, but nearly 20 times more effective for top contributors. Given that top contributors, on average, produce far more content per day, the platform-level impact of increasing tips is even more pronounced among this select group. While top contributors constitute only 20% of the platform's users, a 10% increase in their tips produces 6.4 times more new content than the same percentage increase in tips would elicit from casual users.

[20] The full set of results is reported in Web Appendix C available at `https://minaameri.com/incentives-appendix`.

Table B-7: Effects of Rewards for Top Contributors vs. Casual Users

	Quantity		Quality			
			Length		Informativeness	
	Casual	Top	Casual	Top	Casual	Top
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Tips	0.0371***	0.0596***	0.0172**	0.0220***	0.0008***	0.0003
	(0.0105)	(0.0086)	(0.0018)	(0.0050)	(0.0002)	(0.0002)
Likes	0.0106***	0.0031**	0.0037**	-0.0033**	0.0000	0.0000
	(0.0014)	(0.0010)	(0.0001)	(0.0010)	(0.0000)	(0.0000)
Compensation	-0.0326	-0.0908**	0.1153***	0.0624	-0.0018**	0.0002
	(0.0208)	(0.0152)	(0.0001)	(0.0326)	(0.0007)	(0.0005)
Badge	-0.1786	0.1328	-0.3965*	-0.0534	-0.0044	-0.0028
	(0.1519)	(0.1272)	(0.1328)	(0.1809)	(0.0049)	(0.0037)
Controls	Yes		Yes		Yes	
Individual-Day Fixed Effects	Yes		Yes		Yes	
Individual-UGC-Year Fixed Effects	Yes		Yes		Yes	
Number of Observations	168,095,525	44,610,705	168,095,525	44,610,705	168,095,525	44,610,705
R^2	0.27	0.33	0.89	0.87	0.68	0.64

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Compensation from the platform reveals an interesting pattern. Among casual users, compensation does not affect how often they post, but encmy ages them to write longer posts. However, the effort involved in writing longer content does not translate into more information as compensation has a significantly negative effect on informativeness, suggesting that the lengthier posts casual users produce are more diluted. To put it differently, the additional word count does not mean that posts contain more valuable information. By contrast, for top contributors, compensation reduces the quantity of posts, but shows a marginally significant positive effect on post length with no discernible effect on informativeness. Put differently, top contributors respond by producing fewer but longer posts with more information when compensated.

Overall, monetary platform rewards boost either the amount or the length of content. Yet, they do not unilaterally elevate substantive depth for all users, underscoring the nuanced interplay between extrinsic rewards and intrinsic motivation in content creation. The effect of compensation suggests a crowding-out of intrinsic motivation to post for top users

– a finding that echoes prior evidence that experienced users may respond less favorably to paid incentives (Qiao et al., 2020). Crowding out typically occurs when a person has higher intrinsic motivation, such as top contributors who like producing content or engage deeply with the community (Bénabou and Tirole, 2006). For these users, platform monetary incentives can shift the meaning of their activity from something freely chosen to something externally driven, undermining the gratification they derived from contributing,

In contrast, casual users often have less intrinsic motivation (Bénabou and Tirole, 2006). Because they were not strongly driven to post in the first place, platform monetary rewards are less likely to interfere with their sense of autonomy or personal enjoyment. Instead, they can function as an extra smy ce of motivation without stripping away an existing internal drive. Effectively, because there is less intrinsic motivation to "crowd out" to begin with, casual users are more apt to see financial incentives as a net gain, rather than a threat to their autonomy or identity as contributors.

Furthermore, while compensation encmy ages longer posts by both groups, only top contributors, who have more experience in content production, are successful in improving the content of the posts as well. Hence, compensation has the potential to be effective for both types of users, but in different ways: some produce less but still high-caliber content, whereas others post at the same rate but with longer diluted posts.

These findings illuminate the reality that intrinsic and extrinsic motivators do not operate uniformly across different user segments. Top contributors, driven by a mix of intrinsic interest and deeper platform engagement, respond to rewards in ways that can differ sharply from casual users, who are more easily influenced by platform rewards or social signals, such as likes. In practical terms, when a small subset of users wields substantial influence due to their high output, platform managers must be mindful that certain rewards may either discmy age these top contributors (by crowding out intrinsic motivation) or fail to resonate if they become too common without providing more tangible benefits (as with likes). As a result, when designing a system that encmy ages a broader range of users to participate more regularly, platforms must be careful to avoid crowding out the high-quality content that top contributors reliably generate.

While these findings illustrate how different reward mechanisms impact user groups in their *current* state, they do not yet answer a central question: can casual users evolve into top contributors over time if given the right incentives? To put it differently, if top contributors are so pivotal to a platform's success, is it possible for a platform to strategically use rewards to "create" new top contributors from the casual user base? To investigate this question, I next turn to a temporal analysis of how top contributors and casual users respond to rewards as their membership tenure increases, illuminating whether – and how – platforms might effectively nurture tomorrow's power users.

2.6.1 Effects of Rewards over Top vs. Casual Users' Tenure

In this section, I investigate whether the duration of a top contributor's and a casual user's platform membership affects their reaction to rewards. Understanding whether the importance of rewards varies between top contributors and casual users in early membership periods can shed light on the question whether the platform can, in fact, "create" more top contributors through a strategic use of rewards. Specifically, it raises the possibility that early interventions might steer casual users toward higher levels of engagement.

To examine how the impact of rewards evolves over time, I interact the reward variables with the number of years a user has been a platform member ("tenure") and re-estimate the models for top and casual users separately. Table B-8 presents the estimated effects and reveals fmy patterns in how the effects of rewards evolve as contributors become more seasoned.[21] First, early in their membership, top contributors and casual users show no statistically significant differences in their responses to rewards. In terms of post quantity, both groups respond positively to monetary peer rewards and negatively to monetary platform rewards. Furthermore, while it seems that casual users change the informativeness of

[21]The full set of results is reported in Web Appendix C available at `https://minaameri.com/incentives-appendix`.

the posts after receiving rewards, the effect differences between top contributors and casual users are not statistically significant.

Table B-8: Effects of Rewards for Top and Casual Users over Their Tenure

	Quantity		Quality			
			Length		Informativeness	
	Casual	Top	Casual	Top	Casual	Top
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Tips	0.0432*	0.0819***	0.0271	0.0307*	0.0014*	0.0006
	(0.0217)	(0.0171)	(0.0153)	(0.0147)	(0.0006)	(0.0005)
Likes	0.0174***	0.0183***	0.0017	-0.0005	0.0002*	0.0001
	(0.0033)	(0.0022)	(0.0027)	(0.0024)	(0.0001)	(0.0001)
Compensation	-0.1890***	-0.1561***	0.0123	0.0847	-0.0035*	-0.0030*
	(0.0598)	(0.0367)	(0.0746)	(0.0786)	(0.0015)	(0.004)
Badge	-0.3783	0.4972	-0.4612	-0.5645	-0.0128	0.0048
	(0.4686)	(0.3589)	(0.4514)	(0.4307)	(0.0097)	(0.097)
Tips × Tenure	-0.0044	-0.0116	-0.0062	-0.0047	-0.0004	-0.0002
	(0.0141)	(0.0117)	(0.0077)	(0.0072)	(0.0003)	(0.0003)
Likes × Tenure	-0.0040*	-0.0081***	0.0011	-0.0015	-0.0001*	0.0000
	(0.0017)	(0.0012)	(0.0014)	(0.0012)	(0.0000)	(0.0000)
Compensation × Tenure	0.1038**	0.0355	0.0694	-0.0122	0.0011	0.0018**
	(0.0366)	(0.0213)	(0.0466)	(0.0438)	(0.0008)	(0.0007)
Badge × Tenure	0.1196	-0.2825	0.0425	0.3829	0.0050	-0.0056
	(0.2489)	(0.1904)	(0.2798)	(0.2540)	(0.0063)	(0.0056)
Controls	Yes		Yes		Yes	
Individual-Day Fixed Effects	Yes		Yes		Yes	
Individual-UGC-Year Fixed Effects	Yes		Yes		Yes	
Number of Observations	168,095,525	44,610,705	168,095,525	44,610,705	168,095,525	44,610,705
R^2	0.27	0.32	0.89	0.87	0.68	0.64

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Second, for both top contributors and casual users, likes lose their effectiveness. This result is similar to Burtch et al. (2022)'s finding that new users are more responsive to peer rewards. For new users, peer rewards help foster community ties and relational bonds (Ren, Kraut, and Kiesler, 2007), but users may become used to them over time, rendering likes less potent as motivators. In other words, platforms can encmy age frequent likes to jumpstart top and casual users, but should also recognize that these rewards lose momentum for long-term members.

Third, while the impact of compensation on new users is negative, the interaction effect between compensation and tenure is positive and significant. The baseline coefficients representing the effects of compensation for new users from both groups are not statistically different from each other. However, because the interaction term for casual users is substantially greater, they experience a diminishing crowding-out over time. In other words, casual users may initially respond negatively to overt monetary rewards, but as they spend more time on the platform, they begin to welcome the additional incentive without feeling coerced. Meanwhile, top contributors maintain a comparatively stronger negative response to compensation over the long run. Overall, these findings suggest that platforms might benefit from staging monetary incentives to accommodate how each user segment evolves, while recognizing that top contributors remain more susceptible to potential crowding-out.

And lastly, not only are both groups' responses similar early on, but they also change in similar ways over time. Even when differences emerge (see Figure B-4), top contributors either lose sensitivity to certain rewards (e.g., likes) or respond negatively to others (e.g., compensation). These findings suggest that rewards are unlikely to be a key factor in determining who ultimately becomes a top contributor. Instead, factors such as intrinsic motivation likely play a more critical role in reaching and maintaining this status.

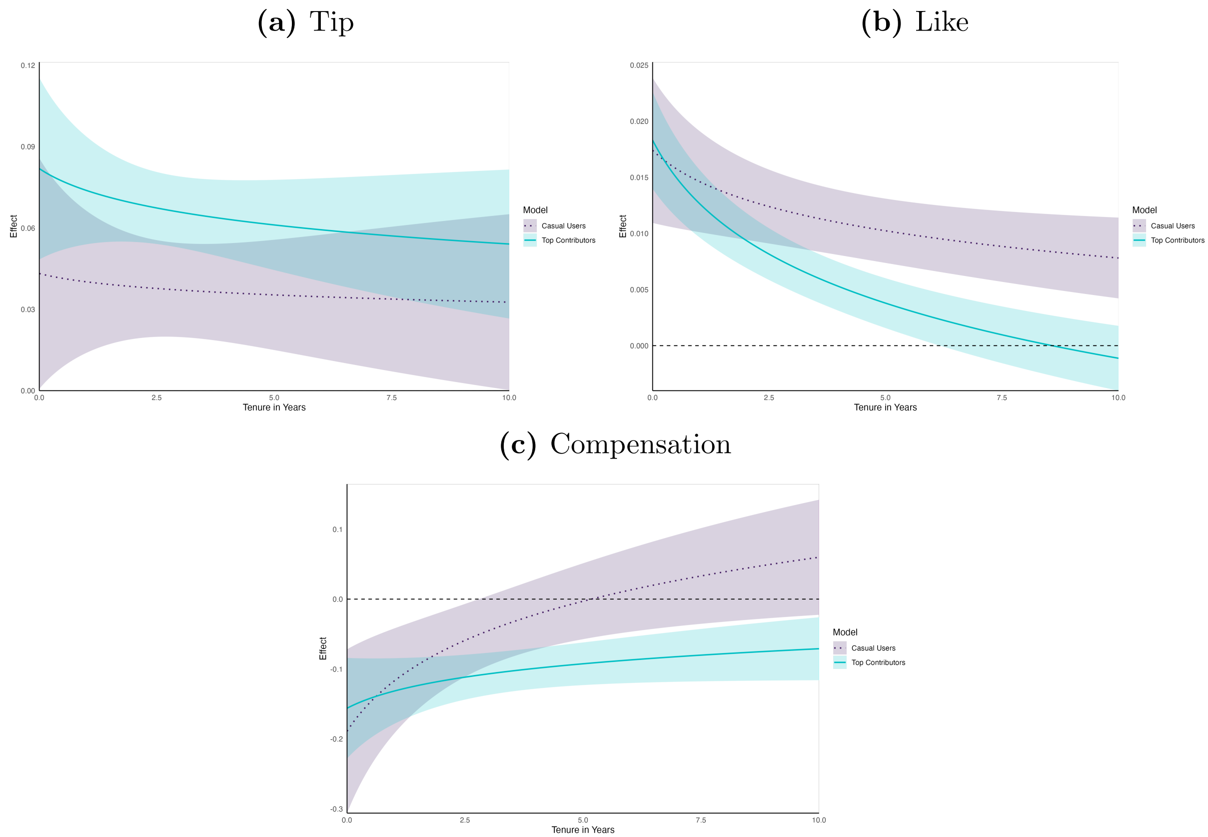

Figure B-4: Effects of Rewards on Post Quantity of Top Contributors and Casual Users Over Their Tenure

Exploring additional aspects of my data provides further evidence for these points. Recall that, in the previous analysis, I classified those users as top contributors who wrote the most posts per day over the *whole* observation period, i.e., top 20% of users in terms of average daily post quantity over the whole observation period ("all-time top contributors"). To see if very active users were consistently very active during their tenure, I adjust my definition of top contributors: instead of over the whole observation period, I examine each year separately, i.e., for each year, I identify the top 20% of users based on their average daily posting in that year and refer to them as that year's top contributors. More than 70% of those who ranked as top contributors in any given year were also all-time top contributors. About 60% of the all-time top contributors were top contributors for at least two-thirds of their membership years. Among the all-time top contributors, only 22% were not classified as a top contributor in their first membership year. However, even the users within this 22% subset posted an average of 44 posts in their first year, far above the six posts observed among average casual users. Indeed, more than half of these 22% were already among the upper half of the content producers in the first year after joining the platform, suggesting that they were on a trajectory toward prolific participation from the start.

Thus, while rewards encmy age casual users to increase their activity levels, the increases are not large enough to transform a sizable proportion of casual users into top contributors. Rather, the dominant pattern is that people who become top contributors are already inclined to higher participation levels early on. Rewards, such as tips and compensation, might accelerate or refine engagement, but they do not appear to be its primary driver. However, what works at the beginning of a user's membership may need adjustments as that user becomes more experienced, regardless of user type. New users, for instance, may respond enthusiastically to likes or find certain types of compensation off-putting. Meanwhile, long-time members might require more sustaining forms of appreciation – like tips, which retain their perceived value – while being wary of direct compensation if it crowds out their intrinsic joy in posting. Thus, while platforms may not be able to design reward strategies to push users toward becoming top contributors, they can utilize strategies that target new and seasoned users differently. For example, platforms can highlight newcomers' posts, increasing their visibility and likelihood of receiving rewards such as likes.

2.7 Robustness Checks

I conduct multiple checks to show the robustness of my results. The results of all robustness checks are shown in Web Appendix D available at `https://minaameri.com/incentives-appendix`. First, I estimate my model with 1-day and 7-day time windows. Recall that I use a 3-day time window in my main specification. The results are qualitatively robust. Second, I estimate my model with the number of tips and number of compensation rewards a user received (using a 3-day window). Recall that I use the amount (in GG) of tips and compensation the user received in my main specification. The results are qualitatively robust.

Third, I estimate my model using less granular individual-week instead of individual-day fixed effects. The results are qualitatively robust. Fmy th, I estimate my model without any fixed effects, i.e., I drop the individual-day and individual-year-UGC type fixed effects. The results are directionally robust. And lastly, I estimate the regressions only including one type of reward at a time. The results are robust. I conclude that my results are robust to a variety of alternative specifications.

2.8 Conclusion

How to encmy age users to write more content and at the desired quality level is a crucial question for the survival and success of many social media platforms. In this paper, I investigate which rewards can be used to achieve these goals. My results show that monetary peer rewards (tips) are effective in increasing the total volume of content and encmy aging users to invest greater effort in their posts. By contrast, monetary rewards from the platform encmy age users to increase the quality of their posts while reducing the total volume of their contributions. Non-monetary peer rewards (likes) show a small positive effect individually in increasing post quantity, but their sheer volume can have a substantial impact. Badges minimally affect short-term posting behavior and exhibit a U-shaped pattern in how they motivate users over longer stretches.

Another important dimension is how these incentives resonate differently with various user segments and over time. Specifically, platforms are keenly interested in attracting, retaining, incentivizing, and potentially “creating” top contributors, who produce the majority of a platform’s content, to remain active. Accordingly, I investigate how rewards affect the behavior of these top contributors compared to the behavior of casual users and whether the platform can “create” more top contributors through strategic use of rewards and early interventions that might steer casual users toward higher levels of engagement.

Top contributors and casual users show no significant differences early in their membership in their responses to rewards and change similarly over time as well. These results indicate that rewards are not a particularly effective tool managers can use to create top contributors. Furthermore, even when differences emerge, top contributors either lose sensitivity to certain rewards or respond more negatively to others than casual users. The negative effect of monetary rewards from the platform on the number of posts is especially pronounced for top contributors, who are likely intrinsically motivated. For these highly active individuals, platform-based compensation can undermine their sense of autonomy, leading to a reduction in the total number of posts they create, even though the quality of each individual post may rise. Further, likes lose much of their value for top contributors and are mostly effective for either newer or casual contributors who continue to perceive likes as a signal of acknowledge.

From a managerial standpoint, it follows that aligning the type of reward with the platform's immediate and long-term goals is crucial. Platforms seeking to maximize short-run content volume may benefit from fostering peer-to-peer monetary rewards, such as tips, which encmy age more frequent and lengthier posts, while non-monetary recognition (such as likes) can be particularly useful for onboarding newer users and keeping casual users engaged. The goal of generating detailed, in-depth reviews or analyses over a longer horizon may, however, call for a more strategic use of platform-based compensation.

By offering monetary rewards selectively, for instance, unlocking additional rewards after a specific number of posts or temporary monetary boosts (e.g., "post three times this week to unlock a small reward") for less active users, platforms can target specific groups of users to impact the content quality or quantity. Nevertheless, managers must remain cognizant of potential crowding-out effects: top contributors, whose content often anchors a platform's community and draws significant user traffic, may be less responsive or even negatively affected by these direct forms of compensation if it erodes their intrinsic desire to contribute.

To mitigate this, platforms can design payment structures that feel less transactional, such as integrating rewards into broader recognition systems (e.g., funding for user-led initiatives or charitable donations in their name), thereby preserving intrinsic motivation while still providing tangible benefits (Cassar, 2019).

Non-monetary peer rewards, such as likes, often lose their effect on more experienced users, necessitating a shift toward incentives that convey higher perceived value but do not diminish autonomy. For instance, experienced contributors may be motivated by exclusive perks, such as limited-time ad-free experiences or increased visibility for their content. At the same time, compensation's crowding-out effect may weaken with tenure for less active users, allowing them to eventually respond more positively to monetary rewards. This indicates that the most effective incentive systems pay close attention not only to user type, i.e., top contributor versus casual user, but also to how these users evolve over their membership and posting history. The interplay between intrinsic and extrinsic motivations remains delicate and can have significant implications for the quality and quantity of UGC.

When it comes to the quality of UGC, it is important to keep in mind that posts that measure higher on the two quality dimensions are not necessarily better for a platform. Whether a platform would benefit from, e.g., more of less informative posts, also depends on the starting point, i.e., the current quality level of posts, and the topic of platform. In balancing short-term objectives against long-term community health, managers must also weigh the trade-offs between maintaining a high volume of user submissions and fostering detailed, in-depth contributions, and design a flexible reward scheme that combines monetary elements with social acknowledgement according to users' tenure and activity levels. Newcomers may benefit from frequent positive feedback to develop the habit of participation, while seasoned members may crave more substantial forms of recognition, either through peer-to-peer tipping or carefully designed compensation schemes. My findings ultimately underscore that one-size-fits-all reward programs are ill-suited to sustain user engagement and UGC quality. A thoughtful mix of monetary and non-monetary mechanisms, aligned with different user stages and tailored to the platform's overarching goals, is necessary to encmy age both the volume and the depth of contributions that can sustain an active, thriving community.

My research is not without limitations. First, I focus on UGC in text form and do not examine other forms of UGC, e.g., videos. This limitation is driven by BGG not using visual content. It is left for future research to examine whether my findings carry over for other forms of UGC. Second, I mainly focus on short-term effects of rewards, i.e., how receiving a reward affects user behavior in the following three days. While I test the robustness of my results with a longer time window of seven days and find that the effects of rewards decrease, I leave studying longer-term effects for future research. Third, I focus on two dimensions of text quality, i.e., length and informativeness. While existing literature on content quality has found that length of a text is highly correlated with, and thus representative of, several quality aspects of the text, there are other text aspects that can also reflect quality, e.g., relevance of the images and links used in the text. Future research can further examine the impact of rewards on this aspect of content quality.

Fmy th, I do not examine the effects of platform monetary rewards on replies because the platform does not provide any monetary rewards for this type of UGC. As a result, I are unable to analyze the potential impact of such rewards on the quantity and quality of replies. It would be interesting for future research to explore the effects of rewards from the platform for replies, and to compare these effects to those of peer rewards. And lastly, the quantity and quality of the content on the platform can also impact the platform's appeal to new visitors and their inclination towards becoming a member. I do not model platform growth. I leave it for future research to study how different types of incentives impact member acquisition and characteristics of these new members.

2.A Data Details

2.A.1 Badges

Users receive badges for writing a certain number of posts of specific types, i.e., any post in the discussion forum, reviews, game session reports, and game lists (Geeklists). Each type of content has its own chain of milestones and badges. A list of these badges and their corresponding thresholds is shown in Figure A-1.

Figure A-1: Badge Levels for Different UGC Types

- **Poster Level Badges (for each domain: BGG, RPGG, VGG)**
  - - Level 01 BGG/RPGG/VGG Poster (100-249 posts)
  - - Level 02 BGG/RPGG/VGG Poster (250-499 posts)
  - - Level 03 BGG/RPGG/VGG Poster (500-749 posts)
  - - Level 04 BGG/RPGG/VGG Poster (750-999 posts)
  - - Level 05 BGG/RPGG/VGG Poster (1000-1499 posts)
  - - Level 06 BGG/RPGG/VGG Poster (1500-1999 posts)
  - - Level 07 BGG/RPGG/VGG Poster (2000-2499 posts)
  - - Level 08 BGG/RPGG/VGG Poster (2500-2999 posts)
  - - Level 09 BGG/RPGG/VGG Poster (3000-3499 posts)
  - - Level 10 BGG/RPGG/VGG Poster (3500-3999 posts)
  - - Level 11 BGG/RPGG/VGG Poster (4000-4499 posts)
  - - Level 12 BGG/RPGG/VGG Poster (4500-4999 posts)
  - - Level 13 BGG/RPGG/VGG Poster (5000-5499 posts)
  - - Level 14 BGG/RPGG/VGG Poster (5500-5999 posts)
  - - Level 15 BGG/RPGG/VGG Poster (6000-6499 posts)
  - - Level 16 BGG/RPGG/VGG Poster (6500-6999 posts)
  - - Level 17 BGG/RPGG/VGG Poster (7000-7999 posts)
  - - Level 18 BGG/RPGG/VGG Poster (8000-8999 posts)
  - - Level 19 BGG/RPGG/VGG Poster (9000-9999 posts)
  - - Level 20 BGG/RPGG/VGG Poster (10000+ posts)
- **Game Reviews**
  - Copper - 5 Reviews
  - Silver - 50 Reviews
  - Gold - 100 Reviews
  - Platinum - 250 Reviews
  - Herculean - 500 Reviews
  - Ultimate - 1000 Reviews
- **Session Reports**
  - Copper - 5 Sessions
  - Silver - 50 Sessions
  - Gold - 100 Sessions
  - Platinum - 250 Sessions
  - Herculean - 500 Sessions
  - Ultimate - 1000 Sessions
- **GeekLists**
  - Copper - 5 GeekLists
  - Silver - 25 GeekLists
  - Gold - 50 GeekLists

Table A-1 shows summary statistics for the number of earned badges by the users in our sample within a year. The 'BGG Poster' badge is awarded for forum posts related to board games. Notably, reviews, session reports, and replies to other forum posts are also counted as forum posts for the purpose of earning this badge. Similarly, 'RPGG Poster' and 'VGG Poster' are for all the forum posts related to role-playing and digital games, respectively.

Table A-1: Number of Badges Earned by Users in Our Sample Within a Year

Badge Group	Thread	Related UGC Review	Report	Reply	List	Mean	SD	Min	Median	Max	N
BGG Poster	✓	✓	✓	✓		0.19	0.40	0.00	0.00	8.00	16,688
RPGG Poster	✓	✓	✓	✓		0.01	0.12	0.00	0.00	5.00	16,688
VGG Poster	✓	✓	✓	✓		0.00^{+}	0.04	0.00	0.00	2.00	16,688
Reviews		✓				0.01	0.12	0.00	0.00	13.00	16,688
Session Reports			✓			0.00^{+}	0.05	0.00	0.00	4.00	16,688
GeekLists					✓	0.01	0.04	0.00	0.00	1.00	16,688

2.A.2 Likes

Table A-2 shows the average number of likes users receive per post for each UGC type, conditional on receiving a like. Users receive the fewest number of likes for replies with an average of 4 likes per post and receive the most likes for reviews followed by reports with averages of 32 and 17 likes, respectively.

Table A-2: Number of Likes Received Per Post

	Mean	SD	Min	Median	Max	N
Threads	8.88	19.39	1.00	4.00	2,041.00	285,249
Reviews	31.64	44.95	1.00	19.00	1,254.00	23,799
Session Reports	17.39	20.18	1.00	12.00	422.00	14,475
Replies	3.89	5.81	1.00	2.00	425.00	5,220,679
Lists	8.09	21.77	1.00	3.00	2,137.00	695,809

Note: 00^{+} is a very small number greater than 0.

2.A.3 Tips and Compensation per Post

Table A-3 reports the summary statistics for the quantity and amount of tips and compensation users receive per post, conditional on receiving a reward. Users, on average, receive 1-3 tips for each thread, review, report, and reply they write, while an average list receives 9 likes. On average, users receive a total of 5.72 GG in tips for writing a thread, 3.04 GG for writing a review, 2.70 GG for writing a report, 13.00 GG for writing a list, and 1.31 GG for writing a reply. Additionally, users, on average, earn the total of 3.04 compensation per thread, 2.40 compensation per review, and 2.23 compensation per report. Note that users do not receive compensation from BGG for writing replies or lists.

Table A-3: Tips and Compensation Per Post

	Mean	SD	Min	Median	Max	N
Tips Per Post						
Threads						
Quantity	2.35	3.90	1.00	1.00	129.00	27,028
Amount	5.71	40.51	0.00^{+}	1.00	2,536.38	27,028
Reviews						
Quantity	2.74	3.60	1.00	2.00	78.00	8,318
Amount	3.04	7.20	0.01	1.00	152.08	8,318
Session Reports						
Quantity	2.25	2.24	1.00	1.00	33.00	3,862
Amount	2.70	8.79	0.01	1.00	300.07	3,862
Replies						
Quantity	1.34	1.10	1.00	1.00	161.00	350,816
Amount	1.31	9.46	0.00^{+}	0.25	2,655.01	350,816
Lists						
Quantity	8.91	44.01	1.00	3.00	1924.00	13,630
Amount	13.00	83.38	0.01	1.25	3,171.00	13,630
Compensation Per Post						
Threads						
Amount	3.04	1.12	1.00	3.00	5.00	1,526
Reviews						
Amount	2.40	0.86	1.00	2.20	7.00	18,405
Session Reports						
Amount	2.23	0.80	1.00	2.10	7.50	11,373

Note: 00^{+} is a very small number greater than 0.

2.A.4 Construction of Quality Control Measures

We employ measures capturing structure, content, and style dimensions of text quality. These measures are commonly used in the literature and applicable to our context (e.g., Stvilia et al. 2005; Hasan Dalip et al. 2009; Shah and Pomerantz 2010). Structural features are captured by the number of words, the number of sentences in a post, the number of words per sentence, the number of characters, the number of monosyllable and polysyllable words, and reading time (Blumenstock 2008; Demberg and Keller 2008; Hasan Dalip et al. 2009;

Anderka, Stein, and Lipka 2012). Reading time is operationalized as the time an average person needs to read a text, typically about 14.69 ms per character (Demberg and Keller 2008).

The Flesch Easing Read Index (FERI) (Kincaid et al., 1975), the Gunning Fog Index (GFI) (Gunning 1952), Automated Readability Index (Smith and Senter, 1967), and Coleman-Liau Index (Coleman and Liau, 1975) are the four content-related measures we employ. They reflect the complexity of the text. FERI is a readability/complexity score, typically between 0-100, that indicates the difficulty of understanding a passage in English (Kincaid et al. 1975), with higher scores corresponding to easier texts. The GFI, Automated Readability Index, and Coleman-Liau Index measure the readability of a text by estimating the number of years of formal education a person needs to understand a text when reading it for the first time using different formulas. For instance, a GFI of 12 indicates that a person must be a high school senior (around 18 years old) to understand a text.

We capture the linguistic style of the text in politeness using domain-independent lexical and syntactic features. These features operationalize key components of politeness theory, such as indirection, deference, impersonalization, and modality (Danescu-Niculescu-Mizil et al., 2013). To quantify politeness, we employ the classifier proposed by Danescu-Niculescu-Mizil et al. (2013), which produces a continuous score reflecting the likelihood that a text will be perceived as impolite by others. Lastly, we include two separate measures for the number of images that are included in the post and the number of hyperlinks referring to other internal or external web pages.

While these measures are used to capture different aspects of a text, some of them reflect similar underlying constructs and are highly correlated with each other. As a result, we conduct a factor analysis to combine these measures into orthogonal factors. The results from the factor analysis suggest using four factors. Table A-4 shows the factor loadings of the fourtheen quality measures on each of the factors. The number of words, number of monosyllable and polysyllable words, number of characters, number of sentences, number of words per sentence, and reading time are grouped into one factor, which we call *Length*, reflecting the extensiveness of the text (Hong et al. 2017). In our main analysis, the dependent variable reflecting the quality of UGC in terms of length is operationalized as the number of words to facilitate interpretation.

The next set of variables are grouped in a factor we refer to as *Complexity* since they captures the difficulty with which a reader can understand a written text. Note that the FREI loading has an opposite sign compared to the other three variables, as higher FREI scores indicate easier texts, whereas higher scores on the other three measures indicate more difficult texts. The number of images and the number of hyperlinks are grouped and we refer to this group as *Number of Linked Content*. Lastly, *Politeness*, as a separate factor, captures the linguistic style of the text.

In addition, we also quantify the informativeness of a text as another content-related measure (Sun, Han, and Feng 2019). The construction of this variable is similar to our main dependent variable, i.e., the ratio of the extracted number of facts in the text using ReVerb Open Information Extraction (Fader, Soderland, and Etzioni, 2011).

Table A-4: Rotated Factor Analysis Loadings

Variable	Factors				Uniqueness
	Length	Complexity	Number of Linked Content	Politeness	
Number of Words	0.9730	0.0640	0.0288	−0.0027	0.0484
Number of Monosyllable Words	0.9604	0.0510	0.0228	0.0048	0.0744
Number of Characters	0.9410	0.1058	0.0283	−0.0234	0.1019
Reading Time	0.9410	0.1058	0.0283	−0.0234	0.1019
Number of Sentences	0.8957	−0.0063	0.0711	0.0366	0.1912
Number of Words per Sentence	0.8697	0.1958	0.0087	0.0124	0.2050
Number of Polysyllable Words	0.5968	0.0488	−0.0524	−0.0502	0.6362
Automated Readability Index	0.0974	0.9051	0.0214	−0.0239	0.1703
Coleman-Liau Index	0.0719	0.8767	0.0299	−0.0138	0.2252
FREI	−0.0785	−0.8506	0.0103	−0.0305	0.2693
GFI	0.1452	0.8480	−0.0106	0.0167	0.2594
Number of Images	0.0318	0.0236	0.7722	−0.1099	0.3901
Number of Hyperlinks	0.1412	0.0120	0.6604	0.1343	0.5258
Politeness	−0.0075	0.0018	−0.0131	0.9832	0.0331

2.B Complete Sets of Results

Table B-1: Effects of Rewards on UGC Quantity and Quality

	(i)	(ii)	(iii)
	Quantity	Quality	
		Length	Informativeness
Tips	0.0572***	0.0221***	0.0003
	(0.0086)	(0.0050)	(0.0002)
Likes	0.0088***	-0.0020*	0.0000***
	(0.0011)	(0.0010)	(0.0000)
Compensation	-0.0841***	0.0751**	-0.0001
	(0.0147)	(0.0326)	(0.0005)
Badge	0.0914	-0.1029	-0.0030
	(0.1234)	(0.1809)	(0.0036)
Received Tips for Other UGC Type	0.0390***	0.0205***	0.0005**
	(0.0087)	(0.0052)	(0.0002)
Received Likes for Other UGC Type	0.0139***	-0.0037***	-0.0000***
	(0.0012)	(0.0011)	(0.0000)
Received Compensation for Other UGC Type	-0.0579***	0.0200	-0.0003
	(0.0141)	(0.0324)	(0.0005)
Received Badge for Other UGC Type	0.0770	-0.1080	-0.0030
	(0.1234)	(0.1809)	(0.0036)
Number of Days Since i's Last Post of UGC Type j	0.4793***	0.0642***	0.0002*
	(0.0045)	(0.0044)	(0.0001)
Number of Days Since i's Last Post of UGC Type j Squared	-2.1371***	-0.3205***	-0.0007
	(0.0194)	(0.0192)	(0.0005)
Number of Days Since i's Last Post in Other UGC Type	-0.1639***	0.1183***	-0.0009***
	(0.0039)	(0.0056)	(0.0001)
Number of Days Since i's Last Post in Other UGC Type Squared	0.5964***	-0.5474***	0.0034***
	(0.0165)	(0.0242)	(0.0006)
Number of user i's Written Posts of UGC Type j	0.0202***	-0.0068***	-0.0001***
	(0.0008)	(0.0007)	(0.0000)
Number of All Written Posts of UGC Type j	0.1685***	0.2567***	0.0002*
	(0.0015)	(0.0043)	(0.0001)
Number of All Written Posts of UGC Type j During 30 Days Prior	-0.0040***	0.0013***	0.0001***
	(0.0002)	(0.0002)	(0.0000)
Number of Written Posts in Other UGC Type	-0.0129***	0.0028*	0.0000
	(0.0012)	(0.0014)	(0.0001)
Number of Written Posts of UGC Type j During...			
... 3 Days Prior	-0.0100***	-0.0216***	-0.0004***
	(0.0030)	(0.0025)	(0.0001)
... 7 Days Prior	0.0233***	-0.0109***	0.0002*
	(0.0020)	(0.0015)	(0.0001)
... 30 Days Prior	-0.0380***	-0.0160***	0.0000***
	(0.0007)	(0.0008)	(0.0000)
... 365 Days Prior	-0.0106***	-0.0030***	0.0000***
	(0.0003)	(0.0003)	(0.0000)
Number of Written Posts in Other UGC Type During...			
... 3 Days Prior	0.0232***	-0.0321***	0.0004***
	(0.0033)	(0.0028)	(0.0001)
... 7 Days Prior	-0.0042	-0.0198***	0.0002*
	(0.0023)	(0.0017)	(0.0001)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-2: Effects of Rewards on UGC Quantity and Quality (Cont. 1)

	(i)	(ii)	(iii)
	Quantity	Quality	
		Length	Informativeness
...7 Days Prior	-0.0864***	-0.0018	-0.0003***
	(0.0016)	(0.0010)	(0.0000)
... 30 Days Prior	0.0159***	0.0056***	0.0001***
	(0.0004)	(0.0005)	(0.0000)
If Ever Wrote Post in Other UGC Type Dummy	0.0110***	0.0021*	0.0001***
	(0.0012)	(0.0010)	(0.0000)
If Wrote Post in Other UGC Type Dummy During...			
... 3 Days Prior	-0.0025***	0.0031***	-0.0000***
	(0.0004)	(0.0004)	(0.0000)
... 7 Days Prior	0.0114***	0.0068***	0.0001***
	(0.0014)	(0.0010)	(0.0000)
... 30 Days Prior	-0.0158***	0.0058***	-0.0001***
	(0.0006)	(0.0005)	(0.0000)
If Wrote Post of UGC Type j at t Dummy		3.6206***	0.0858***
		(0.0073)	(0.0002)
If Next BGG Poster Badge Milestone Exists Dummy	0.0040	-0.0049	0.0000
	(0.0032)	(0.0035)	(0.0001)
0 BGG Poster Badge Earned Dummy	0.0419**	0.0269**	0.0006
	(0.0152)	(0.0106)	(0.0006)
1 BGG Poster Badge Earned Dummy	0.0489***	0.0265**	0.0006
	(0.0152)	(0.0108)	(0.0006)
2 to 5 BGG Poster Badges Earned Dummy	-0.0029***	0.0015	0.0000***
	(0.0009)	(0.0012)	(0.0000)
6 to 15 BGG Poster Badges Earned Dummy	0.0468***	0.0114	0.0003
	(0.0156)	(0.0116)	(0.0006)
Number of Posts Needed To Next BGG Poster Badge ×...			
... 0 BGG Poster Badge Earned Dummy	-0.0043***	0.0001	0.0001***
	(0.0008)	(0.0010)	(0.0000)
... 1 BGG Poster Badge Earned Dummy	-0.0032*	-0.0042***	-0.0000
	(0.0015)	(0.0014)	(0.0001)
... 2 to 5 BGG Poster Badges Earned Dummy	0.0003***	-0.0002	-0.0000***
	(0.0001)	(0.0002)	(0.0000)
... 6 to 15 BGG Poster Badges Earned Dummy	-0.0070***	-0.0004	-0.0000
	(0.0022)	(0.0027)	(0.0001)
... 16 to 20 BGG Poster Badges Earned Dummy	0.0094	-0.0015	0.0000
	(0.0055)	(0.0044)	(0.0002)
Number of Posts Needed To Next BGG Poster Badge Squared ×...			
... 0 BGG Poster Badge Earned Dummy	0.0012***	0.0002	-0.0000***
	(0.0001)	(0.0002)	(0.0000)
... 1 BGG Poster Badge Earned Dummy	0.0003	0.0008***	0.0000***
	(0.0002)	(0.0002)	(0.0000)
... 2 to 5 BGG Poster Badges Earned Dummy	0.0495***	0.0175	0.0005
	(0.0151)	(0.0109)	(0.0006)
... 6 to 15 BGG Poster Badges Earned Dummy	0.0008**	0.0001	0.0000***
	(0.0003)	(0.0003)	(0.0000)
... 16 to 20 BGG Poster Badges Dummy	-0.0012*	0.0004	-0.0000***
	(0.0006)	(0.0006)	(0.0000)
If Next RPGG Badge Milestone Exists Dummy	0.0053	0.0058	0.0002*
	(0.0034)	(0.0044)	(0.0001)
0 RPGG Poster Badge Earned Dummy	-0.0415***	-0.0049	0.0004
	(0.0126)	(0.0127)	(0.0005)
1 RPGG Poster Badge Earned Dummy	-0.0035	-0.0185	-0.0012
	(0.0125)	(0.0170)	(0.0008)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-3: Effects of Rewards on UGC Quantity and Quality (Cont. 2)

	(i)	(ii)	(iii)
	Quantity	Quality	
		Length	Informativeness
Number of Posts Needed To Next RPGG Poster Badge Squared ×...			
... 0 RPGG Poster Badge Earned Dummy	-0.0022	-0.0047**	-0.0002*
	(0.0013)	(0.0017)	(0.0001)
... 1 RPGG Poster Badge Earned Dummy	0.0007	-0.0012	-0.0001
	(0.0017)	(0.0018)	(0.0001)
... 2 to 20 RPGG Poster Badges Earned Dummy	0.0010*	0.0001	-0.0000***
	(0.0005)	(0.0008)	(0.0000)
If Next VGG Poster Badge Milestone Exists Dummy	-0.0006	0.0031	-0.0002
	(0.0036)	(0.0041)	(0.0002)
0 VGG Poster Badge Earned Dummy	0.0125	0.0476	0.0016
	(0.0258)	(0.0299)	(0.0009)
1 VGG Poster Badge Earned Dummy	0.0159	0.0048	0.0003
	(0.0263)	(0.0332)	(0.0012)
Number of Posts Needed To Next VGG Poster Badge ×...			
... 0 VGG Poster Badge Earned Dummy	0.0209**	-0.0070	-0.0001
	(0.0084)	(0.0105)	(0.0003)
... 1 VGG Poster Badge Earned Dummy	0.0099	-0.0372**	-0.0003
	(0.0125)	(0.0146)	(0.0005)
... 2 to 20 VGG Poster Badges Earned Dummy	0.0059	0.0041	0.0004
	(0.0077)	(0.0121)	(0.0005)
Number of Posts Needed To Next VGG Poster Badge Squared ×...			
... 0 VGG Poster Badge Earned Dummy	-0.0044**	0.0016	0.0000
	(0.0018)	(0.0023)	(0.0001)
... 1 VGG Poster Badge Earned Dummy	-0.0019	0.0085***	0.0001
	(0.0024)	(0.0027)	(0.0001)
... 2 to 20 VGG Poster Badges Earned Dummy	-0.0008	-0.0009	-0.0001
	(0.0010)	(0.0016)	(0.0001)
0 Review Badge Earned Dummy	-0.0059	-0.0095	0.0013*
	(0.0151)	(0.0616)	(0.0006)
1 Review Badge Earned Dummy	0.0005	0.2642***	0.0016**
	(0.0196)	(0.0663)	(0.0006)
Number of Posts Needed To Next Review Badge ×...			
... 0 Review Badge Earned Dummy	0.0243*	0.1586***	-0.0004*
	(0.0109)	(0.0190)	(0.0002)
... 1 Review Badge Earned Dummy	-0.0110	-0.1054**	-0.0001
	(0.0132)	(0.0374)	(0.0004)
... 2 to 6 Review Badges Earned Dummy	-0.0062	0.0698	0.0008*
	(0.0082)	(0.0459)	(0.0004)
Number of Posts Needed To Next Review Badge Squared ×...			
... 0 Review Badge Earned Dummy	-0.0169**	-0.0898***	0.0002*
	(0.0061)	(0.0106)	(0.0001)
... 1 Review Badge Earned Dummy	0.0035	0.0123	-0.0000
	(0.0028)	(0.0082)	(0.0001)
... 2 to 6 Review Badges Earned Dummy	0.0012	-0.0125	-0.0001
	(0.0015)	(0.0087)	(0.0001)
0 Session Report Badge Earned Dummy	-0.0590***	0.0648	0.0005
	(0.0188)	(0.0701)	(0.0008)
1 Session Report Badge Earned Dummy	-0.0150	0.2728***	0.0013
	(0.0298)	(0.0780)	(0.0008)
2 to 6 Session Report Badges Earned Dummy	-0.0253*	0.1101	0.0006
	(0.0125)	(0.0591)	(0.0006)
Number of Posts Needed To Next Session Report Badge ×...			
... 0 Session Report Badge Earned Dummy	0.0500***	0.0729***	0.0001
	(0.0122)	(0.0159)	(0.0003)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-4: Effects of Rewards on UGC Quantity and Quality (Cont. 3)

	(i) Quantity	(ii) Quality: Length	(iii) Quality: Informativeness
... 1 Session Report Badge Earned Dummy	0.0036	0.0076	-0.0000
	(0.0031)	(0.0075)	(0.0001)
0 Geeklist Badge Earned Dummy	0.0224**	-0.0440***	-0.0001
	(0.0091)	(0.0120)	(0.0003)
1 to 3 Geeklist Badges Earned Dummy	0.0123	-0.0045	-0.0001
	(0.0092)	(0.0125)	(0.0004)
Number of Posts Needed To Next Geeklist Badge ×...			
... 0 Geeklist Badge Earned Dummy	-0.0213**	0.0614***	0.0003
	(0.0091)	(0.0124)	(0.0004)
... 1 to 3 Geeklist Badges Earned Dummy	-0.0031	0.0002	0.0000
	(0.0024)	(0.0033)	(0.0001)
Number of Posts Needed To Next Geeklist Badge Squared ×...			
... 0 Geeklist Badge Earned Dummy	0.0117**	-0.0327***	-0.0002
	(0.0051)	(0.0071)	(0.0002)
Length per Written Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0341***	-0.0364*	-0.0006*
	(0.0059)	(0.0170)	(0.0003)
... Review UGC Type Dummy	0.0299***	-0.0838***	-0.0001
	(0.0075)	(0.0238)	(0.0003)
... Session Report UGC Type Dummy	0.0361***	-0.0317	-0.0003
	(0.0078)	(0.0243)	(0.0003)
... Reply UGC Type Dummy	0.1778***	0.2746***	0.0038***
	(0.0075)	(0.0180)	(0.0003)
... (Baseline) List UGC Type	-0.0416***	0.0457**	0.0005
	(0.0055)	(0.0173)	(0.0003)
Readability per Written Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0433	-1.3528***	-0.0085**
	(0.0675)	(0.1619)	(0.0034)
... Review UGC Type Dummy	-0.3875**	-0.4759	-0.0054
	(0.1710)	(0.4182)	(0.0053)
... Session Report UGC Type Dummy	-0.0184	-0.6867	-0.0165***
	(0.1196)	(0.4233)	(0.0054)
... Reply UGC Type Dummy	-0.0756	-0.0486	0.0053
	(0.0676)	(0.1422)	(0.0037)
... (Baseline) List UGC Type	0.0478	1.2207***	0.0094***
	(0.0632)	(0.1952)	(0.0028)
Politeness per Written Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.1530	2.3253***	0.0245*
	(0.2029)	(0.6253)	(0.0119)
... Review UGC Type Dummy	-1.0458***	-1.4521	0.0454***
	(0.3146)	(1.5695)	(0.0109)
... Session Report UGC Type Dummy	-0.9511*	2.0406**	0.0446***
	(0.4462)	(0.8480)	(0.0125)
... Reply UGC Type Dummy	2.5666***	0.6317	-0.0271
	(0.4762)	(0.6158)	(0.0167)
... (Baseline) List UGC Type	-0.2774	-1.8025***	-0.0424***
	(0.1431)	(0.3849)	(0.0069)
Number of Linked Content Per Written Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.7685***	-0.4043**	-0.0218***
	(0.0867)	(0.1535)	(0.0035)
... Review UGC Type Dummy	-0.8866***	-1.3971***	-0.0257***
	(0.1319)	(0.4629)	(0.0050)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-5: Effects of Rewards on UGC Quantity and Quality (Cont. 4)

	(i) Quantity	(ii) Quality: Length	(iii) Quality: Informativeness
... Review UGC Type Dummy	0.7953***	0.5107***	-0.0003
	(0.0492)	(0.1144)	(0.0016)
... Session Report UGC Type Dummy	0.8217***	0.0219	-0.0009
	(0.0440)	(0.0983)	(0.0015)
... Reply UGC Type Dummy	-0.4562***	0.2172***	0.0087***
	(0.0220)	(0.0260)	(0.0008)
... (Baseline) List UGC Type	0.0061	0.0023	0.0001
	(0.0043)	(0.0087)	(0.0001)
Length per Written Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0027***	0.0151***	0.0001***
	(0.0008)	(0.0018)	(0.0000)
... Review UGC Type Dummy	-0.0002	-0.0014	-0.0001***
	(0.0008)	(0.0018)	(0.0000)
... Session Report UGC Type Dummy	-0.0001	0.0014	-0.0000***
	(0.0007)	(0.0016)	(0.0000)
... Reply UGC Type Dummy	0.0432***	0.1045***	0.0014***
	(0.0026)	(0.0052)	(0.0001)
... (Baseline) List UGC Type	-0.1149**	-0.2774**	-0.0059**
	(0.0485)	(0.1195)	(0.0023)
Readability per Written Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0403***	0.0372***	-0.0006*
	(0.0064)	(0.0096)	(0.0003)
... Review UGC Type Dummy	0.0243***	-0.0098	-0.0005**
	(0.0054)	(0.0094)	(0.0002)
... Session Report UGC Type Dummy	0.0257***	-0.0069	-0.0005**
	(0.0054)	(0.0095)	(0.0002)
... Reply UGC Type Dummy	0.2033***	0.6319***	0.0002
	(0.0358)	(0.0518)	(0.0020)
... (Baseline) List UGC Type	0.4818**	-0.8500**	-0.0042
	(0.1658)	(0.2918)	(0.0068)
Politeness per Written Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.2139**	-0.0555	0.0034
	(0.0738)	(0.1165)	(0.0029)
... Review UGC Type Dummy	-0.1572	-0.0039	0.0063*
	(0.0907)	(0.0559)	(0.0029)
... Session Report UGC Type Dummy	-0.1594	-0.0324	0.0063*
	(0.0899)	(0.0559)	(0.0029)
... Reply UGC Type Dummy	0.0039	0.0817	0.0036
	(0.1412)	(0.2092)	(0.0054)
... (Baseline) List UGC Type	-0.6972***	0.1041	-0.0008
	(0.0484)	(0.0587)	(0.0022)
Number of Linked Content Per Written Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.1263***	-0.0487***	-0.0007
	(0.0094)	(0.0148)	(0.0005)
... Review UGC Type Dummy	-0.0078	-0.0415***	-0.0005
	(0.0069)	(0.0102)	(0.0003)
... Session Report UGC Type Dummy	-0.0075	-0.0443***	-0.0005
	(0.0068)	(0.0097)	(0.0003)
... Reply UGC Type Dummy	0.1183**	-0.3404***	-0.0178***
	(0.0407)	(0.0531)	(0.0023)
... (Baseline) List UGC Type	-0.5404***	-0.0462*	0.0036***
	(0.0236)	(0.0225)	(0.0007)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-6: Effects of Rewards on UGC Quantity and Quality (Cont. 5)

	(i) Quantity	(ii) Quality	(iii)
		Length	Informativeness
... Reply UGC Type Dummy	0.6000***	-0.0411***	0.0083***
	(0.0101)	(0.0103)	(0.0006)
... (Baseline) List UGC Type	-0.3748***	-0.0005	-0.0016**
	(0.0126)	(0.0122)	(0.0006)
Length per Rewarded Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0127***	0.0249***	0.0006***
	(0.0026)	(0.0061)	(0.0001)
... Review UGC Type Dummy	0.0179***	0.0316***	0.0005***
	(0.0026)	(0.0064)	(0.0001)
... Session Report UGC Type Dummy	0.0145***	0.0409***	0.0003***
	(0.0031)	(0.0078)	(0.0001)
... Reply UGC Type Dummy	-0.0015	0.0848***	0.0009***
	(0.0035)	(0.0072)	(0.0002)
... (Baseline) List UGC Type	-0.0101***	-0.0150**	-0.0004***
	(0.0025)	(0.0060)	(0.0001)
Readability per Rewarded Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0579*	0.0438	0.0031**
	(0.0276)	(0.0444)	(0.0011)
... Review UGC Type Dummy	0.1170***	0.2176***	0.0023
	(0.0300)	(0.0590)	(0.0013)
... Session Report UGC Type Dummy	0.1065**	0.2003**	-0.0001
	(0.0374)	(0.0700)	(0.0015)
... Reply UGC Type Dummy	-0.0447	0.3507***	0.0039***
	(0.0268)	(0.0437)	(0.0013)
... (Baseline) List UGC Type	-0.0091	-0.0434	-0.0039***
	(0.0242)	(0.0481)	(0.0010)
Politeness per Rewarded Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0216	0.1171	-0.0005
	(0.0574)	(0.0828)	(0.0023)
... Review UGC Type Dummy	-0.0657	0.3813	-0.0012
	(0.0790)	(0.1958)	(0.0028)
... Session Report UGC Type Dummy	-0.0562	0.0431	-0.0024
	(0.0745)	(0.1785)	(0.0028)
... Reply UGC Type Dummy	-0.3598*	0.0041	0.0197*
	(0.1636)	(0.3131)	(0.0089)
... (Baseline) List UGC Type	0.0501	-0.0260	-0.0014
	(0.0606)	(0.1034)	(0.0025)
Number of Linked Content Per Rewarded Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.1381**	-0.1389	-0.0003
	(0.0539)	(0.0771)	(0.0020)
... Review UGC Type Dummy	-0.2205***	-0.0729	-0.0023
	(0.0528)	(0.0911)	(0.0021)
... Session Report UGC Type Dummy	-0.1791**	-0.1186	-0.0014
	(0.0670)	(0.1313)	(0.0028)
... Reply UGC Type Dummy	-0.3203***	-0.1999**	-0.0046
	(0.0546)	(0.0796)	(0.0024)
... (Baseline) List UGC Type	0.1792***	0.1339	-0.0015
	(0.0521)	(0.0781)	(0.0021)
Info per Rewarded Post of UGC Type j During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0723***	-0.0228	-0.0005
	(0.0112)	(0.0144)	(0.0004)
... Review UGC Type Dummy	0.0322 **	-0.0071	0.0002
	(0.0131)	(0.0203)	(0.0005)

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table B-7: Effects of Rewards on UGC Quantity and Quality (Cont.6)

	(i)	(ii)	(iii)
	Quantity	Quality	
		Length	Informativeness
... Session Report UGC Type Dummy	-0.0196***	0.0038*	0.0002***
	(0.0012)	(0.0017)	(0.0000)
... Reply UGC Type Dummy	0.6000***	-0.0411***	0.0083***
	(0.0101)	(0.0103)	(0.0006)
... (Baseline) List UGC Type	-0.3748***	-0.0005	-0.0016**
	(0.0126)	(0.0122)	(0.0006)
Length per Rewarded Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0003	0.0019***	0.0000***
	(0.0004)	(0.0006)	(0.0000)
... Review UGC Type Dummy	-0.0004	0.0005	0.0000***
	(0.0004)	(0.0006)	(0.0000)
... Session Report UGC Type Dummy	-0.0002	0.0015***	0.0000***
	(0.0003)	(0.0005)	(0.0000)
Readability per Rewarded Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	0.0034	0.0254***	0.0003
	(0.0046)	(0.0067)	(0.0002)
... Review UGC Type Dummy	0.0056	0.0129**	0.0002
	(0.0040)	(0.0057)	(0.0002)
... Session Report UGC Type Dummy	0.0056	0.0174***	0.0003
	(0.0040)	(0.0057)	(0.0002)
... Reply UGC Type Dummy	0.0303	0.1379***	-0.0021**
	(0.0162)	(0.0194)	(0.0008)
... (Baseline) List UGC Type	0.1277	0.0597	-0.0033
	(0.0661)	(0.0989)	(0.0031)
Politeness per Rewarded Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0039	0.0522	0.0011
	(0.0276)	(0.0424)	(0.0015)
Number of Linked Content Per Rewarded Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0409***	-0.0073	0.0003
	(0.0100)	(0.0136)	(0.0005)
... Review UGC Type Dummy	-0.0283***	-0.0273**	0.0001
	(0.0081)	(0.0107)	(0.0003)
... Session Report UGC Type Dummy	-0.0310***	-0.0194	-0.0001
	(0.0079)	(0.0104)	(0.0003)
... Reply UGC Type Dummy	0.0218	0.0003	-0.0001
	(0.0284)	(0.0321)	(0.0016)
... (Baseline) List UGC Type	-0.0781***	-0.0090	-0.0001
	(0.0125)	(0.0139)	(0.0004)
Info per Rewarded Post in Other UGC Type During 3 Days Prior ×...			
... Thread UGC Type Dummy	-0.0174***	-0.0057***	0.0000
	(0.0012)	(0.0017)	(0.0001)
... Review UGC Type Dummy	-0.0073***	-0.0057***	-0.0000***
	(0.0011)	(0.0013)	(0.0000)
... Session Report UGC Type Dummy	-0.0078***	-0.0054***	0.0000***
	(0.0010)	(0.0013)	(0.0000)
... Reply UGC Type Dummy	0.0199***	-0.0047	0.0023***
	(0.0053)	(0.0053)	(0.0003)
... (Baseline) List UGC Type	-0.1274***	0.0066	0.0006*
	(0.0077)	(0.0079)	(0.0003)
Constant	-0.0000***	0.0000***	-0.0000***
	(0.0000)	(0.0000)	(0.0000)
Individual-Day Fixed Effects	Yes	Yes	Yes
Individual-UGC-Year Fixed Effects	Yes	Yes	Yes
Quality of Past Posts	Yes	Yes	Yes
Number of Observations	212,706,230	212,706,230	212,706,230
R^2	0.29	0.88	0.65

Clustered standard errors in parentheses.
The dependent and independent variables are in logarithmic form.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

Akerlof, George A (1980), "A Theory of Social Custom, of Which Unemployment May Be One Consequence," *The Quarterly Journal of Economics*, 94 (4), 749–775.

Ameri, Mina, Elisabeth Honka, and Ying Xie (2019), "Learning from Online Social Ties," *Marketing Science*, 38 (4), 567–583.

— (2023a), "From Strangers to Friends: Tie Formations and Online Activities in an Evolving Social Network," *Journal of Marketing Research*, 60 (2), 329–354.

— (2023b), "From Strangers to Friends: Tie Formations and Online Activities in an Evolving Social Network," *Journal of Marketing Research*, 60 (2), 329–354.

Anderka, Maik, Benno Stein, and Nedim Lipka (2012), "Predicting Quality Flaws in User-Generated Content: The Case of Wikipedia," in "Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval," 981–990.

Angrist, Joshua and Joern-Steffen Pischke (2009), *Mostly Harmless Econometrics*, Princeton University Press.

Arellano, Manuel and Jinyong Hahn (2007), *Understanding Bias in Nonlinear Panel Models: Some Recent Developments*, Cambridge University Press, 381–409, Econometric Society Monographs.

Azar, Ofer (2004), "What Sustains Social Norms and How They Evolve? The Case of Tipping," *Journal of Economic Behavior & Organization*, 54 (1), 49–64.

— (2007), "The Social Norm of Tipping: A Review," *Journal of Applied Social Psychology*, 37 (2), 380–402.

— (2010), "Tipping Motivations and Behavior in the US and Israel," *Journal of Applied Social Psychology*, 40 (2), 421–457.

Azar, Ofer H (2020), "The Economics of Tipping," *Journal of Economic Perspectives*, 34 (2), 215–236.

Bateman, Patrick, Peter Gray, and Brian Butler (2011), "The Impact of Community Commitment on Participation in Online Communities," *Information System Research*, 22 (4), 841–854.

Bénabou, Roland and Jean Tirole (2006), "Incentives and prosocial behavior," *American economic review*, 96 (5), 1652–1678.

Berndt, Ernst, Bronwyn Hall, Robert Hall, and Jerry Hausman (1974), "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement*, 3/4 (1), 653–665.

Bernheim, B Douglas (1994), "A Theory of Conformity," *Journal of Political Economy*, 102 (5), 841–877.

Bicchieri, Cristina, Ryan Muldoon, and Alessandro Sontuoso (2011), "Social Norms," *Stanford Encyclopedia of Philosophy*.

Blumenstock, Joshua (2008), "Size Matters: Word Count as A Measure of Quality on Wikipedia," in "Proceedings of the 17th International Conference on World Wide Web," 1095–1096.

Burtch, Gordon, Qinglai He, Yili Hong, and Dokyun Lee (2022), "How Do Peer Awards Motivate Creative Content? Experimental Evidence from Reddit," *Management Science*, 68 (5), 3488–3506.

Burtch, Gordon, Yili Hong, Ravi Bapna, and Vladas Griskevicius (2018), "Stimulating Online Reviews by Combining Financial Incentives and Social Norms," *Management Science*, 64 (5), 2065–2082.

Cabral, Luis and Lingfang Li (2015), "A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay," *Management Science*, 61 (9), 2052–2063.

Cassar, Lea (2019), "Job mission as a substitute for monetary incentives: Benefits and limits," *Management Science*, 65 (2), 896–912.

Chandar, Bharat, Uri Gneezy, John List, and Ian Muir (2019), "The Drivers of Social Preferences: Evidence From a Nationwide Tipping Field Experiment," *Working Paper*.

Chen, Yan, Rosta Farzan, Robert Kraut, Iman YeckehZaare, and Ark Fangzhou Zhang (2024), "Motivating experts to contribute to digital public goods: A personalized field experiment on Wikipedia," *Management Science*, 70 (5), 3264–3280.

Chen, Yan, Maxwell Harper, Joseph Konstan, and Sherry Li (2010), "Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens," *American Economic Review*, 100 (4), 1358–1398.

Chen, Yubo, Qi Wang, and Jinhong Xie (2011), "Online Social Interactions: A Natural Experiment on Word of Mouth Versus Observational Learning," *Journal of Marketing Research*, 48 (2), 238–254.

Chevalier, Judy and Dina Mayzlin (2006), "The Effect Of Word-of-Mouth On Sales: Online Book Reviews," *Journal of Marketing Research*, 43 (3), 345–354.

Coleman, Meri and Ta Lin Liau (1975), "A computer readability formula designed for machine scoring." *Journal of Applied Psychology*, 60 (2), 283.

Conlin, Michael, Michael Lynn, and Ted O'Donoghue (2003), "The Norm of Restaurant Tipping," *Journal of Economic Behavior & Organization*, 52 (3), 297–321.

Danescu-Niculescu-Mizil, Cristian, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts (2013), "A computational approach to politeness with application to social factors," in "Proceedings of ACL," .

Deci, Edward L, Richard Koestner, and Richard M Ryan (1999), "A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation." *Psychological bulletin*, 125 (6), 627.

Deci, Edward L and Richard M Ryan (2013), *Intrinsic motivation and self-determination in human behavior*, Springer Science & Business Media.

Demberg, Vera and Frank Keller (2008), "Data From Eye-Tracking Corpora as Evidence for Theories of Syntactic Processing Complexity," *Cognition*, 109 (2), 193–210.

Epstein, Joshua (2001), "Learning To Be Thoughtless: Social Norms and Individual Computation," *Computational Economics*, 18, 9–24.

Everett, Jim, Lucius Caviola, Guy Kahane, Julian Savulescu, and Nadira Faber (2015), "Doing Good by Doing Nothing? The Role of Social Norms in Explaining Default Effects in Altruistic Contexts," *European Journal of Social Psychology*, 45 (2), 230–241.

Fader, Anthony, Stephen Soderland, and Oren Etzioni (2011), "Identifying Relations for Open Information Extraction," in "Proceedings of the Conference of Empirical Methods in Natural Language Processing (EMNLP '11)," Edinburgh, Scotland, UK.

Fehr, Ernst and Simon Gächter (2000), "Fairness and Retaliation: The Economics of Reciprocity," *Journal of Economic Perspectives*, 14 (3), 159–182.

Frey, Bruno S and Reto Jegen (2001), "Motivation crowding theory," *Journal of economic surveys*, 15 (5), 589–611.

Gallus, Jana (2017), "Fostering Public Good Contributions with Symbolic Award: A Large Scale Natural Field Experiment at Wikipedia," *Management Science*, 63 (12), 3999–4015.

Gallus, Jana, Olivia S Jung, and Karim R Lakhani (2020), "Recognition Incentives for Internal Crowdsourcing: A Field Experiment at NASA," Working Paper.

Garrod, Simon and Gwyneth Doherty (1993), "Conversation, Co-Ordination and Convention: An Empirical Investigation of How Groups Establish Linguistic Conventions," *Cognition*, 53 (3), 181–215.

Gneezy, Uri and Aldo Rustichini (2000), "Pay Enough or Don't Pay at all," *The Quarterly Journal of Economics*, 115 (3), 791–810.

Goes, Paulo, Chenhui Guo, and Mingfeng Lin (2016), "Do Incentive Hierarchies Induce User Effort? Evidence from An Online Knowledge Exchange," *Information Systems Research*, 27 (3), 497–516.

Goldstein, Noah J., Robert B. Cialdini, and Vladas Griskevicius (2008), "A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels," *Journal of Consumer Research*, 35 (3), 472–482.

Gunning, Robert (1952), *Technique of Clear Writing*, McGraw-Hill.

Haggag, Kareem and Giovanni Paci (2014), "Default Tips," *American Economic Journal*, 6 (3), 1–19.

Hanson, Sara, Lan Jiang, and Darren Dahl (2019), "Enhancing Consumer Engagement in an Online Brand Community via User Reputation Signals: A Multi-Method Analysis," *Journal of the Academy of Marketing Science*, 47 (5), 349–367.

Hasan Dalip, Daniel, Marcos André Gonçalves, Marco Cristo, and Pável Calado (2009), "Automatic Quality Assessment of Content Created Collaboratively by Web Communities: a Case Study of Wikipedia," in "Proceedings of the 9th ACM/IEEE-CS Joint Conference on Digital Libraries," 295–304.

He, Liuyi, Jifeng Luo, Yisong Tang, Zhiyan Wu, and Han Zhang (2023), "Motivating User-Generated Content: The Unintended Consequences of Incentive Thresholds," *Management Information Systems Quarterly*, 47 (3), 1015–1044.

Hendricks, Ken and Alan Sorensen (2009), "Information and the Skewness of Music Sales," *Journal of Political Economy*, 117 (2), 324–369.

Hennig-Thurau, Thorsten, Kevin P. Gwinner, Gianfranco Walsh, and Dwayne D. Gremler (2004), "Electronic Word-Of-Mouth via Consumer-Opinion Platforms: What Motivates Consumers to Articulate Themselves on the Internet?" *Journal of Interactive Marketing*, 18 (1), 38–52.

Homburg, Christian, Jan Wieseke, and Christina Kuehnl (2010), "Social Influence on Salespeople's Adoption of Sales Technology: A Multilevel Analysis," *Journal of the Academy of Marketing Science*, 38 (2), 159–168.

Hong, Hong, Di Xu, Alan Wang, and Weiguo Fan (2017), "Understanding the Determinants of Online Review Helpfulness: A Meta-Analytic Investigation," *Decision Support Systems*, 102 (C), 1–11.

Horn, Christopher, Alisa Zhila, Alexander Gelbukh, Roman Kern, and Elisabeth Lex (2013), "Using Factual Density to Measure Informativeness of Web Documents," in "Proceedings of the 19th Nordic Conference of Computational Linguistics (NODALIDA 2013)," 227–238.

Huang, Justin T, Rupali Kaul, and Sridhar Narayanan (2022), "The Causal Effect of Attention and Recognition on the Nature of User-Generated Content: Experimental Results from an Image-Sharing Social Network," Working Paper.

Huang, Ni, Gordon Burtch, Bin Gu, Yili Hong, Chen Liang, Kanliang Wang, Dongpa Fu, and Bo Yahg (2019), "Motivating User-Generated Content with Performance Feedback: Evidence from Randomized Field Experiments," *Management Science*, 65 (1), 327–345.

Huang, Ni, Yili Hong, and Gordon Burtch (2017), "Social network integration and user content generation," *MIS quarterly*, 41 (4), 1035–1058.

Khern-am nuai, Warut, Karthik Kannan, and Hossein Ghasemkhani (2018), "Extrinsic versus Intrinsic Rewards for Contributing Reviews in an Online Platform," *Information System Research*, 29 (4), 871–892.

Kim, Seung Hyun, On Amir, and Kenneth C. Wilbur (2023), "Tipping in a Digital Services Marketplace," Technical report, Working Paper.

Kincaid, J Peter, Robert Fishburne Jr, Richard Rogers, and Brad Chissom (1975), "Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel," Technical report, Naval Technical Training Command Millington TN Research Branch.

Kozinets, Robert (1999), "E-Tribalized Marketing?: The Strategic Implications of Virtual Communities of Consumption," *European Management Journal*, 17 (2), 252–264.

Kreps, David M (1997), "Intrinsic Motivation and Extrinsic Incentives," *The American Economic Review*, 87 (2), 359–364.

Kuang, Lini, Ni Huang, Yili Hong, and Zhijun Yan (2019), "Spillover Effects of Financial Incentives on Non-Incentivized User Engagement: Evidence From an Online Knowledge Exchange Platform," *Journal of Management Information Systems*, 36 (1), 289–320.

Lal, Rajiv and Venkataraman Srinivasan (1993), "Compensation Plans for Single- And Multi-Product Salesforces: An Application of the Holmstrom-Milgrom Model," *Management Science*, 39 (7), 777–793.

Lee, Richard, Jamie Murphy, and Larry Neale (2009), "The Interactions of Consumption Characteristics on Social Norms," *Journal of Consumer Marketing*, 26 (4), 277–285.

Leung, Louis (2009), "User-Generated Content on the Internet: An Examination of Gratifications, Civic Engagement and Psychological Empowerment," *New Media & Society*, 11 (8), 1327–1347.

Lex, Elisabeth, Michael Voelske, Marcelo Errecalde, Edgardo Ferretti, Leticia Cagnina, Christopher Horn, Benno Stein, and Michael Granitzer (2012), "Measuring The Quality of Web Content Using Factual Information," in "Proceedings of the 2nd joint WICOW/AIRWeb workshop on web quality," 7–10.

Li, Xinxin and Lorin Hitt (2008), "Self-Selection and Information Role of Online Product Reviews," *Information Systems Research*, 19 (4), 456–474.

Liu, Yuewen and Juan Feng (2021), "Does Money Talk? The Impact of Monetary Incentives on User-Generated Content Contributions," *Information Systems Research*, 32 (2), 394–409.

Lu, Shijie, Ying Xie, and Xingyu Chen (2023), "Immediate and enduring effects of digital badges on online content consumption and generation," *International Journal of Research in Marketing*, 40 (1), 146–163.

Lu, Shijie, Dai Yao, Xingyu Chen, and Rajdeep Grewal (2021), "Do Larger Audiences Generate Greater Revenues Under Pay What You Want? Evidence from a Live Streaming Platform," *Marketing Science*, 40 (5), 964–984.

Lynn, Michael (2014), "The Contribution of Norm Familiarity to Race Differences in Tipping: A Replication and Extension," *Journal of Hospitality & Tourism Research*, 38 (3), 414–425.

— (2016), "Why Are We More Likely to Tip Some Service Occupations Than Others? Theory, Evidence, and Implications," *Journal of Economic Psychology*, 54, 134–150.

— (2018), "How Motivations for Tipping Vary with Occupational Differences in Descriptive Tipping Norms," *Journal of Behavioral and Experimental Economics*, 77, 1–10.

Lynn, Michael and Michael Sturman (2010), "Tipping and Service Quality: A Within-Subjects Analysis," *Journal of Hospitality & Tourism Research*, 34 (2), 269–275.

Melnyk, Vladimir, Francois Carrillat, and Valentyna Melnyk (2022), "The Influence of Social Norms on Consumer Behavior: A Meta-Analysis," *Journal of Marketing*, 86 (3), 98–120.

Moe, Wendy and Michael Trusov (2011), "The Value of Social Dynamics in Online Product Ratings Forums," *Journal of Marketing Research*, 48 (3), 444–456.

Mudambi, Suzan and David Schuff (2010), "What Makes a Helpful Online Review? A Study of Customer Reviews on Amazon.com," *MIS Quarterly*, 34 (1), 185–200.

Muldoon, Ryan, Chiara Lisciandra, Cristina Bicchieri, Stephen Hartmann, and Jan Sprenger (2013), "On the Emergence of Descriptive Norms," *Politics, Philosophy & Economics*, 13 (1), 3–22.

Mummalaneni, Simha, Hema Yoganarasimhan, and Varad Pathak (2024), "How Do Content Producers Respond to Engagement on Social Media Platforms?" *Available at SSRN*.

Nevskaya, Yulia and Paulo Albuquerque (2019), "How Should Firms Manage Excessive Product Use? A Continuous-Time Demand Model to Test Reward Schedules, Notifications, and Time Limits," *Journal of Marketing Research*, 56 (3), 379–400.

Neyman, Jerzy and Elizabeth Scott (1948), "Consistent Estimates Based on Partially Consistent Observations," *Econometrica*, 16 (1), 1–32.

Oyer, Paul (1998), "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *The Quarterly Journal of Economics*, 113 (1), 149–185.

Paridar, Mahsa, Mina Ameri, and Elisabeth Honka (2023), "The Impact of Rewards on User-Generated Content," Technical report, Working Paper.

Parrett, Matt (2011), "Do People with Food Service Experience Tip Better?" *The Journal of Socio-Economics*, 40 (4), 464–471.

Qiao, Dandan, Shun-Yang Lee, Andrew B Whinston, and Qiang Wei (2020), "Financial Incentives Dampen Altruism in Online Prosocial contributions: A Study of Online Reviews," *Information Systems Research*, 31 (4), 1361–1375.

Qiao, Dandan and Huaxia Rui (2023), "Text performance on the vine stage? The effect of incentive on product review text quality," *Information Systems Research*, 34 (2), 676–697.

Ren, Yuqing, Robert Kraut, and Sara Kiesler (2007), "Applying common identity and bond theory to design of online communities," *Organization studies*, 28 (3), 377–408.

Restivo, Michael and Arnout van de Rijt (2014), "No Praise Without Effort: Experimental Evidence on How Rewards Affect Wikipedia's Contributor Community," *Information, Communication & Society*, 17 (4), 451–462.

Roberts, Jeffrey A, Il-Horn Hann, and Sandra A Slaughter (2006), "Understanding The Motivations, Participation, And Performance of Open Source Software Developers: A Longitudinal Study of The Apache Projects," *Management Science*, 52 (7), 984–999.

Schuster, Lisa, Krzysztof Kubacki, and Sharyn Rundle-Thiele (2016), "Community-Based Social Marketing: Effects on Social Norms," *Journal of Social Marketing*, 6 (2), 193–210.

Sen, Sandip and Stephane Airiau (2007), "Emergence of Norms Through Social Learning," *Proceedings of the International Joint Conference on Artificial Intelligence*, 1507–1512.

Shah, Chirag and Jefferey Pomerantz (2010), "Evaluating and Predicting Answer Quality in Community QA," in "Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval," 411–418.

Smith, Edgar A and RJ Senter (1967), *Automated readability index*, volume 66, Aerospace Medical Research Laboratories, Aerospace Medical Division, Air.

Stvilia, Besiki, Michael Twidale, Les Gasser, and Linda Smith (2005), "Assessing Information Quality of a Community-Based Encyclopedia," in "Knowledge Management: Nurturing Culture, Innovation, and Technology," World Scientific, 101–113.

Sun, Xinyu, Maoxin Han, and Juan Feng (2019), "Helpfulness of Online Reviews: Examining Review Informativeness And Classification Thresholds by Search Products And Experience Products," *Decision Support Systems*, 124, 113099.

Sun, Yacheng, Xiaojing Dong, and Shelby McIntyre (2017), "Motivation of User-generated Content: Social Connectedness Moderates the Effects of Monetary Rewards," *Marketing Science*, 36 (3), 329–337.

Thompson, Richard F and William A Spencer (1966), "Habituation: a model phenomenon for the study of neuronal substrates of behavior." *Psychological review*, 73 (1), 16.

Toubia, Olivier and Andrew T. Stephen (2013), "Intrinsic vs. Image-Related Utility in Social Media: Why Do People Contribute Content to Twitter," *Marketing Science*, 32 (3), 368–392.

Wang, Qili, Liangfei Qiu, and Wei Xu (2024), "Informal payments and doctor engagement in an online health community: An empirical investigation using generalized synthetic control," *Information Systems Research*, 35 (2), 706–726.

Wasko, Molly McLure and Samer Faraj (2005), "Why Should I Share? Examining Social Capital and Knowledge Contribution in Electronic Networks of Practice," *MIS quarterly*, 29 (1), 35–57.

Woolley, Kaitlin and Marissa A Sharif (2021), "Incentives Increase Relative Positivity of Review Content and Enjoyment of Review Writing," *Journal of Marketing Research*, 58 (3), 539–558.

Yang, Mochen, Yuqing Ren, and Gediminas Adomavicious (2019), "Understanding User-Generated Content and Customer Engagement on Facebook Business Pages," *Information System Research*, 30 (3), 839–855.

Young, Peyton (1993), "The Evolution of Conventions," *Econometrica*, 61 (1), 57–84.

— (2015), "The Evolution of Social Norms," *Annual Review of Economics*, 7 (1), 359–387.

Yu, Yinan, Warut Khern-am nuai, and Alain Pinsonneault (2022), "When Paying for Reviews Pays Off: The Case of Performance-Contingent Monetary Rewards." *MIS Quarterly*, 46 (1), 609–626.

Zhang, Xiaoquan (Michael) and Feng Zhu (2011), "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia," *American Economic Review*, 101 (4), 1601–1615.

Zhang, Yan and Leilei Gao (2016), "Wanting ever more: Acquisition procedure motivates continued reward acquisition," *Journal of Consumer Research*, 43 (2), 230–245.

Zhang, Yuchi and David Godes (2018a), "Learning from Online Social Ties," *Marketing Science*, 37 (3), 424–444.

— (2018b), "Learning from Online Social Ties," *Marketing Science*, 37 (3), 424–444.