Platform Design in a Tipping Economy: Norm Development and the Role of Information Disclosure*

Current version: April 2025

Mahsa Paridar† Mina Ameri‡ Elisabeth Honka§

PRELIMINARY AND INCOMPLETE.
DO NOT CITE OR CIRCULATE WITHOUT THE AUTHORS' PERMISSION.

Abstract

Many social platforms have recently enabled users to tip online content creators. We investigate how individuals decide whether and how much to tip content creators. Our novel data come from an online board game platform, on which users create and consume content, react to content, and can sell and purchase board games. A unique aspect of our data is that we observe all tipping incidences since the introduction of tipping on the platform. We develop a model in which users make tipping decisions as a function of their beliefs about an evolving tipping norm, content quality, user characteristics, and other factors. Users' beliefs about the tipping norm are based on two types of signals: tips they received themselves and tips they observed other users to give to the focal piece of content. A new feature of our model is that we allow the signals to be correlated within a type, across types, and across time. Users incorporate the signals into their belief about the tipping norm via Bayesian updating. Our results show that both types of signals impact users' perception of the current tipping norm with the tips they personally received being a more informative type of signal. Tip amounts are primarily driven by users' beliefs about the tipping norm followed by content quality and user characteristics. Using prediction exercises, we show that users tip smaller amounts but much more often when tips given in the broader community are not visible, increasing the total tip amount by 39%. Our predictions also demonstrate that tipping behavior is sticky after a change in platform's information disclosure, even in the medium-run and especially as it relates to the breadth of tipping. We discuss the implications of different information disclosure levels for platform design.

Keywords: Online Tipping, Platform Design, Information Disclosure, Social Norms, Bayesian Updating.

JEL Classification: D83, L82

*We thank Anand Bodapati, Randy Bucklin, Pradeep Chintagunta, Brett Hollenbeck, Tai Lam, Peter Rossi, and Robert Zeithammer for their suggestions. We thank seminar participants at Indian School of Business, San Diego State University, and Texas A&M for their comments. All errors are our own.

†University of California Los Angeles, mahsaparidar@ucla.edu.

‡University of Pittsburgh, mina.ameri@pitt.edu.

§University of California Los Angeles, elisabeth.honka@anderson.ucla.edu.

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, we 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.

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

[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.

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, we first investigate how different information provision strategies affect the development of a tipping norm and users' tipping decisions. We 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?

We 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. We study users' decisions for 22 months following the introduction of tipping. During this time, users who gave tips, on average,

[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.

gave 5.73 tips and users who received tips, on average, received 15.60 tips annually. Our data show that the standard deviation of tip amounts decreased over time, suggesting that users tipped more similar amounts as time progressed.

In our 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 sources: 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 we allow the signals to be correlated when deriving the posterior distribution. More specifically, signals can be correlated within a source and a time period, across sources 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.

Our 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 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, we separate the portions of the utility that come from content quality, individuals' beliefs about the norm, and individual characteristics (via user fixed effects). We 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, we examine how information disclosure affects users' tipping behavior. We do so by implementing three prediction scenarios: in the first one, users update their perception

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

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. Our 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, we study how sticky the perceived tipping norm and tipping behavior are. We 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, our main model. Our 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, we 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, we 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, we shed light on how users decide to tip. We further show that, while users learn from self-experience and observing others' actions, these two sources of information play different roles in shaping users' beliefs about the norm. By examining how these factors interplay in shaping individual tipping behaviors, we 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, we review the relevant literature. In Section 3, we describe our data. We present our model in Section 4 and discuss the results in Section 5. In the following section, we perform prediction exercises and conclude in section 7.

2 Relevant Literature

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

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

2.2 Social Norms

Social norms are the unwritten codes and informal understandings that define what others expect of us and what we 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 use a game-theoretic or computational approach (e.g., Young 1993; Sen and Airiau 2007; Epstein 2001). To the best of our 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.

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

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

2.3 Special Interest Communities

Lastly, our 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 (2018) 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 (2023) study how strangers become friends on an anime platform. To the best of our knowledge, no empirical study has investigated the board game industry.

3 Data

Our 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 1 shows the number of users joining BGG over time.

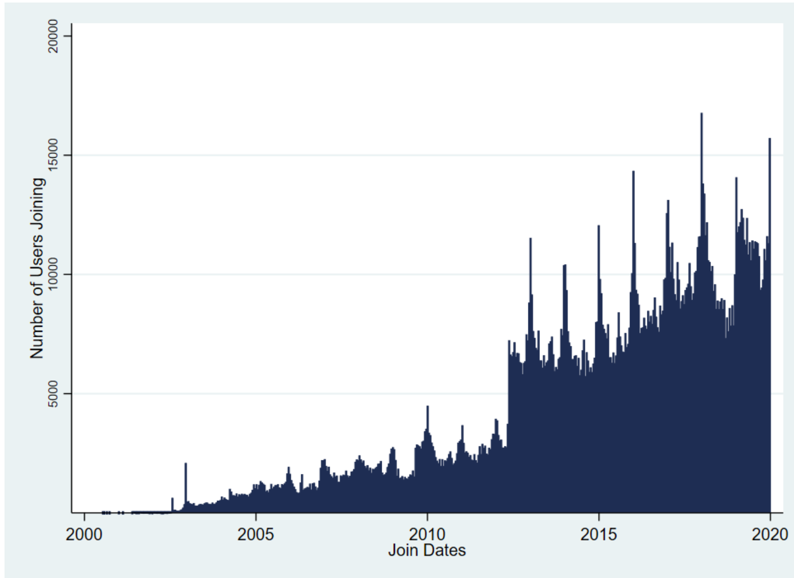

Figure 1: Number of Users Joining BGG Over Time

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 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 2 shows a post for which the content creator received both likes and tips from other users. Users can see who tipped and

[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.

the amount of each tip by clicking on the cent icon and who liked the content by clicking on the thumbs-up icon.

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

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. We 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] We drop 109 tip incidences with tip amounts of more than 20 GG.

For our sample of users, we collected all the content they created, all tips they gave, and information on other spending activities such as purchasing symbolic badges. Furthermore, we 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 our data are that we 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 we make: In 100% of the tipping incidences, users also

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

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

engaged in at least one other activity on the platform e.g., linking content, buying a badge etc. Therefore, we 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, we focus on content for which users had a non-monetary reaction.

In our data, we observe that users typically tip on the same day or on days following a non-monetary reaction. Therefore, we 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. Our panel contains 3.9 million user-content-day observations.

3.2 Data Description

Figure 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.

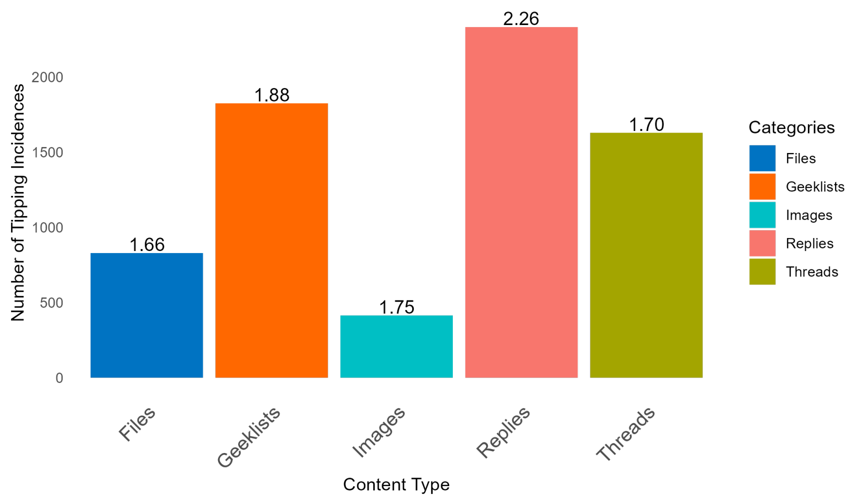

Figure 3: Number of Tip Incidences and Average Tip Amount by Content Category

Table 1 summarizes key statistics of our 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: Descriptive Statistics

Figure 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.

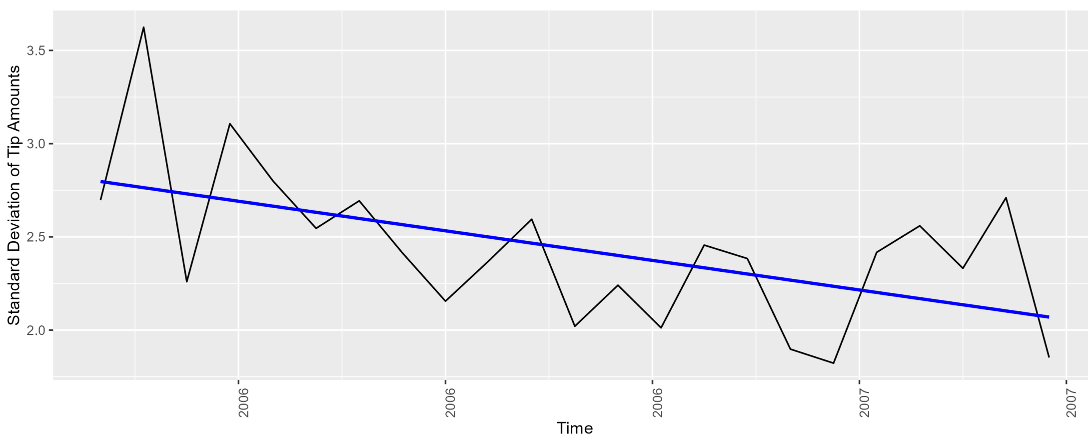

Figure 4: Standard Deviation of Tips Given in a Month Over Time

Figure 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 four times.

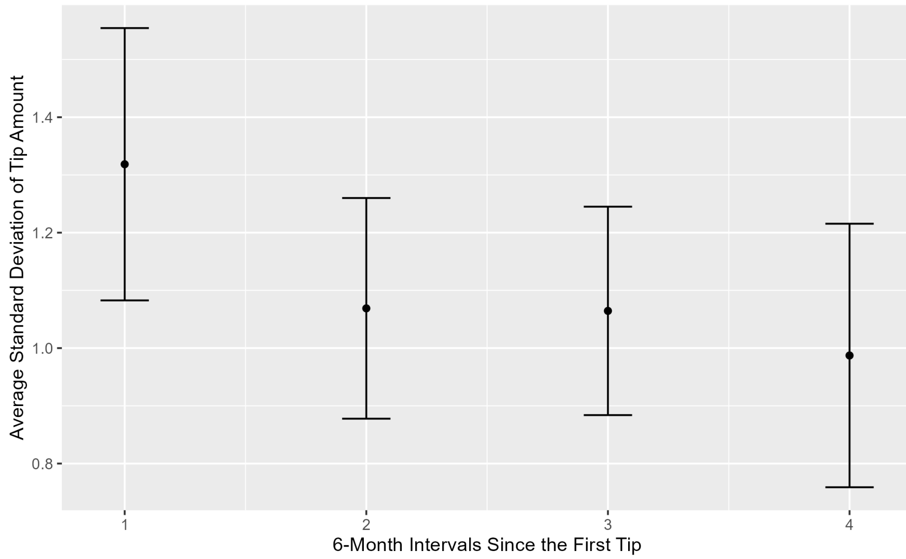

Figure 5: Within-User Standard Deviation of Tip Amounts Over Time

Users might have different tendencies to tip because of their nationality or culture. Table 2 shows the number of tip incidences, average tip amount per tip incidence, and the percentage of users in our data coming from each country.

Table 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

4 Model and Estimation

4.1 Model

4.1.1 Assumptions

We 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 our empirical context, the average ratio of budget to tip for users is 126, we 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 our data, only 2% of tips are given by a pair of users to each other, suggesting that reciprocity is not a major driver in our setting. Therefore, we do not include it in our 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 we neither observe logins nor which content users see, we make the following assumptions based on data patterns. First, our 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, we 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, our data show that in 95% of tipping incidences, users also had a non-tipping reaction to the piece of content. Therefore, we 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, we 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.

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]

$$U_{ijt} = \alpha_i + \beta\mu_{it} + \gamma' Q_{ijt} + \eta' C_{ijt} + \epsilon_{ijt}, \tag{1}$$

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 encourage more (high-quality) content creation in the future (Azar 2007; Paridar, Ameri, and Honka 2023). Q_{ijt} contains variables capturing content quality.

[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.

For textual content, we use the length of the text, operationalized as number of sentences, that 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). We 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. We 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, we 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, we 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, we also include the number of content pieces created by all users on the platform on day t.

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 sources 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, we 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.,

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

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). We 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.,

$$p\{s_{itn}^c \mid \mu, \sigma_c^2\} \sim \mathcal{N}(\mu, \sigma_c^2). \tag{3}$$

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 sources and the dependent variable are the same in nature, signals are a sample from the pool of tipping decisions over different time periods. Thus, we do not need to assume that the variances of the dependent variable and the signals are the same.

Furthermore, because we not only observe when an individual receives a signal but also the value of the signal, we can remain agnostic about whether the signals come from a distribution with constant or time-varying mean. Thus, we 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 we do not observe the exact time of the day when a user makes her tipping decision for each post j, we 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, we 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.

4.1.4 Bayesian Updating with Correlated Signals

In our 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., we can no longer assume that signals come from univariate normal distributions. Instead, we assume that the signals come from a multivariate normal distribution given by

$$p\{s^p_{i1,1}, s^p_{i1,2}, \ldots, s^p_{it,M^p_{it}}, s^c_{i1,1}, s^c_{i2,2}, \ldots, s^c_{it,M^c_{it}} \mid \mu, \Sigma\} \sim \mathcal{MVN}(\mu, \Sigma) \tag{4}$$

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

[12] We assume that a day starts at 6 AM instead of 12 AM to account for users' activities in late hours 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, we model a user's tipping decisions happening at 12 PM.

or noise 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 source as well as between signals from different sources. 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]

$$\mu_{it}=\frac{1}{\omega_{it}}\left(\omega_0\mu_0+\sum_{k,z=1,1}^{M_{it}^{p}+M_{it}^{c}}\frac{\omega_{k,z}(s_{itk}+s_{itz})}{2}\right), \tag{5a}$$

$$\omega_{it}=\omega_0+\sum_{k,z=1,1}^{M_{it}^{p}+M_{it}^{c}}\omega_{kz}, \tag{5b}$$

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 5a and 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, we impose more structure on the precision matrix that results in the estimation of three partial correlations and three decay factors. More specifically, we 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, we first describe the structure of the precision matrix that results in the desired correlation structure. We then present the formulas for the posterior mean and precision. All derivations

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

are shown in detail 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

$$\Omega_t = \begin{bmatrix} \Omega_p & \Omega_{pc}^T \\ \Omega_{pc} & \Omega_c \end{bmatrix} \tag{6}$$

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

$$\Omega_{|\omega,\lambda} \in \{\Omega_{p|\omega_p,\lambda_p}, \Omega_{c|\omega_c,\lambda_c}\} = \begin{bmatrix} \Omega_{11} & \delta^1\Omega_{12} & \cdots & \delta^t\Omega_{1t} \\ \delta^1\Omega_{21} & \Omega_{22} & \cdots & \delta^{t-1}\Omega_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ \delta^t\Omega_{t1} & \delta^{t-1}\Omega_{t2} & \cdots & \Omega_{tt} \end{bmatrix} \tag{7}$$

[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 source 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:

$$\Omega_{kk} = \begin{bmatrix} \omega & \lambda & \cdots & \lambda \\ \lambda & \omega & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \omega \end{bmatrix}_{N_{ik} \times N_{ik}} , \quad \Omega_{zk} = \begin{bmatrix} \lambda & \lambda & \cdots & \lambda \\ \vdots & \vdots & \ddots & \vdots \\ \lambda & \lambda & \cdots & \lambda \end{bmatrix}_{N_{iz} \times N_{ik}} \tag{8}$$

We now turn to the matrix capturing the correlations between signals from different sources, 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 $\delta_{pc}^{|z-k|} \lambda_{pc}$:

$$\Omega_{pc} = \lambda_{pc} \begin{bmatrix} J_{N_{i1}^p \times N_{i1}^c} & \delta_{pc}^1 J_{N_{i1}^p \times N_{i2}^c} & \cdots & \delta_{pc}^t J_{N_{i1}^p \times N_{it}^c} \\ \delta_{pc}^1 J_{N_{i2}^p \times N_{i1}^c} & J_{N_{i2}^p \times N_{i2}^c} & \cdots & \delta_{pc}^{t-1} J_{N_{i2}^p \times N_{it}^c} \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{pc}^t \, J_{N_{it}^p \times N_{i1}^c} & \delta_{pc}^{t-1} J_{N_{it}^p \times N_{i2}^c} & \cdots & J_{N_{it}^p \times N_{it}^c} \end{bmatrix} \tag{9}$$

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

$$\omega_t = \omega_0 + (\omega_p - \lambda_p) \sum_{k=1}^{t} N_{ik}^p + (\omega_c - \lambda_c) \sum_{k=1}^{t} N_{ik}^c + \lambda_p \sum_{k,z=1}^{t} \delta_p^{|k-z|} N_{iz}^p \times N_{ik}^p + \lambda_c \sum_{k,z=1}^{t} \delta_c^{|k-z|} N_{iz}^c \times N_{ik}^c + 2\lambda_{pc} \sum_{k,z=1,1}^{t} \delta_{pc}^{|k-z|} N_{iz}^p \times N_{ik}^c \tag{10}$$

and the posterior mean is given by

$$\mu_{it} = \frac{1}{\omega_t} \Bigg(\omega_0 \mu_0 + (\omega_p - \lambda_p) \sum_{k=1}^{t} \mathbb{S}_{ik}^p + (\omega_c - \lambda_c) \sum_{k=1}^{t} \mathbb{S}_{ik}^c + \lambda_p \sum_{k,z=1}^{t} \left(\delta_p^{|k-z|} N_{ik}^p \mathbb{S}_{iz}^p \right) + \lambda_c \sum_{k,z=1}^{t} \left(\delta_c^{|k-z|} N_{ik}^c \mathbb{S}_{iz}^c \right) + \frac{\lambda_{pc}}{2} \sum_{k,z=1}^{t} \left(\delta_{pc}^{|k-z|} N_{ik}^c \mathbb{S}_{iz}^p \right) + \frac{\lambda_{pc}}{2} \sum_{k,z=1}^{t} \left(\delta_{pc}^{|k-z|} N_{iz}^p \mathbb{S}_{ik}^c \right) \Bigg) \tag{11}$$

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.

4.2 Estimation

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

$$\log L(\theta|y,\Psi,Q,C) = \sum_{i=1}^{N}\sum_{\tau=1}^{\mathcal{T}}\sum_{j=1}^{J}\Bigg\{ I(y_{ijt}>0)\log\left[\frac{1}{\sigma}\phi\left(\frac{y_{ijt}-(\alpha_i+\beta\mu_{it}+\gamma' Q_{ijt}+\eta' C_{ijt})}{\sigma}\right)\right] + I(y_{ijt}=0)\log\left[\Phi\left(\frac{\alpha_i+\beta\mu_{it}+\gamma' Q_{ijt}+\eta' C_{ijt}}{\sigma}\right)\right]\Bigg\} \quad (12)$$

ϕ is the standard normal probability density function, Φ is the standard normal cumulative distribution function, and $I(\cdot)$ is an indicator function. We 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, we show how we rewrite the formulas for the posterior mean and posterior precision to avoid calculating permutations and speed up their computations during the estimation. Our data contain about 4.2 million observations and we estimate about 1,800 individual fixed effects.[15] Even though we estimate a non-linear model with normally distributed errors, we do not face the incidental parameter problem in our empirical setting because of the large T, i.e., we 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, we 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.

5 Results

We present the estimation results in Table 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), we assume that signals are independent. Column (iii) depicts the results for our 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 we 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 we transition to our 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, we focus on our main model shown in column (iii).

Table 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 four 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 our 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 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 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 6(b)). We 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 Web Appendix B.

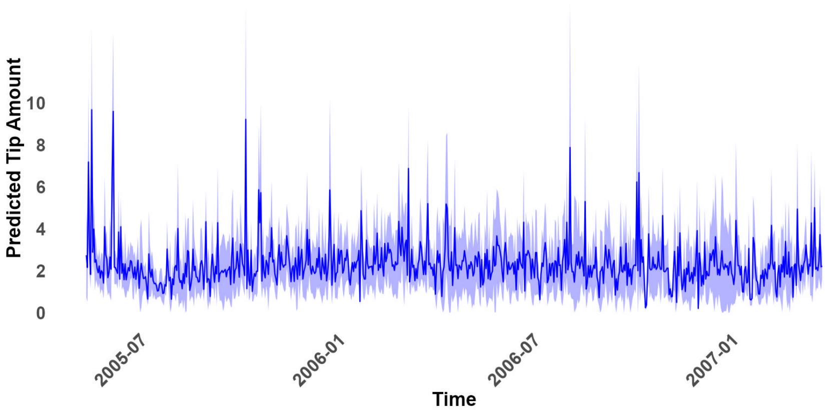

(a) Predicted Tip Amount Per Tip Incidence

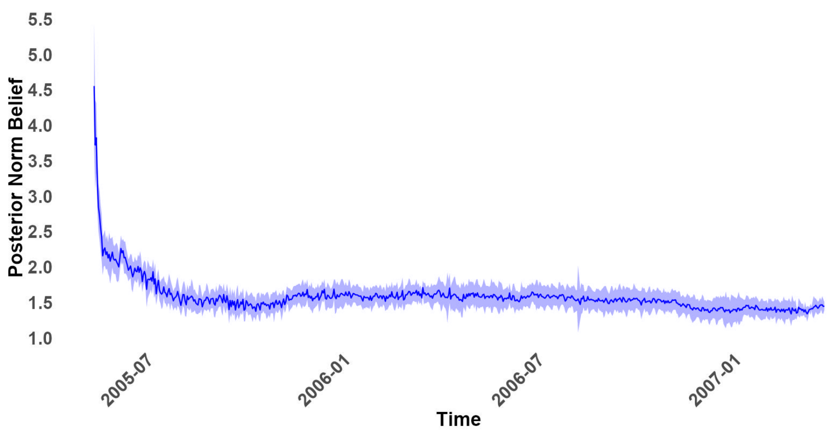

(b) Posterior Belief about Norm

Figure 6: Predicted Tips and Posterior Norm Belief with 95% Confidence Intervals

5.1 Model Fit

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

6 Prediction Exercises

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

6.1 Tip Decomposition

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

Table 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, our results indicate that tipper characteristics only play a minor role in individuals' tipping decisions.

Table 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 7, we examine how the contributions of the four 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.

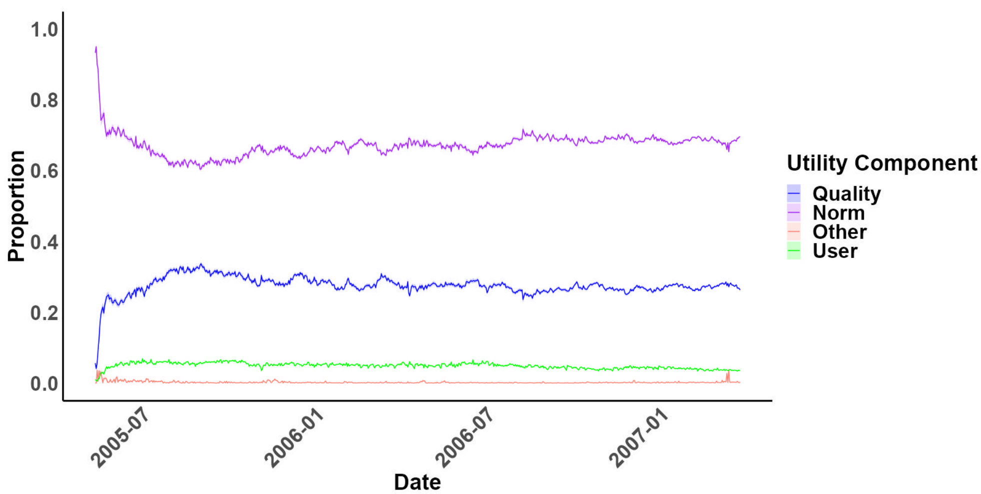

Figure 7: Tip Amount Decomposition over Time with 95% Confidence Intervals

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, we empirically evaluate how different information disclosure strategies, i.e., tip visibility, impact users' tipping behavior.

To investigate the effects of information disclosure, we 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, we 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, we repeat the prediction calculations 100 times (with different error draws) and then compute the average predictions. We then compare the average predictions from these three scenarios to the average predictions from our main model.

We report the average percentage changes for each scenario compared to our baseline main model in Table 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] We 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 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 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 8, we plot users' average posterior norm beliefs for all four scenarios over time. The variation in average posterior norm beliefs is larger when users do not observe or only partially observe community signals.

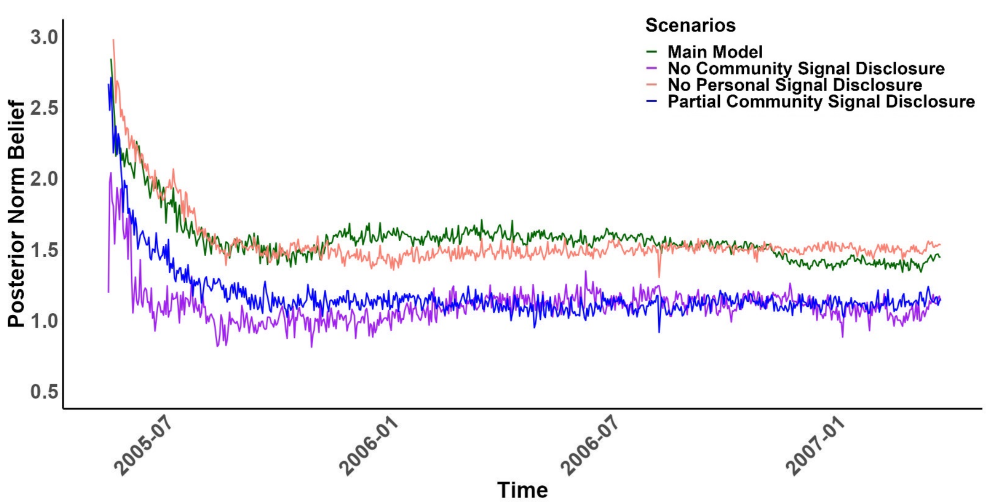

Figure 8: Posterior Norm Over Time

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 source varies significantly in our 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 3), in practice, community signals result in an effect that is 6 times greater than that of personal signals.[17]

6.3 Tipping and Norm Stickiness

In this section, we examine how "sticky" social norms and tipping behavior are. More specifically, we 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 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, we 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 our main model) in the first 11 months of the study period. We then compare the predictions from these three scenarios to those from our main model, where the platform discloses tipping information.

Figure 9 displays the tipping norm development for the four 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 8) because users have already accumulated information about tipping amounts and formed a perceived tipping

[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, we repeat the prediction calculations 100 times (with different error draws) and then compute the average predictions. We then compare the average predictions from these three scenarios to the average predictions from our main model.

norm in 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 our main model with full information disclosure. Generally, the differences in tipping norms between the different scenarios and our main model are smaller compared with the end of the study period in Figure 8 indicating a degree of stickiness that persists even 11 months after a change in platform information disclosure.

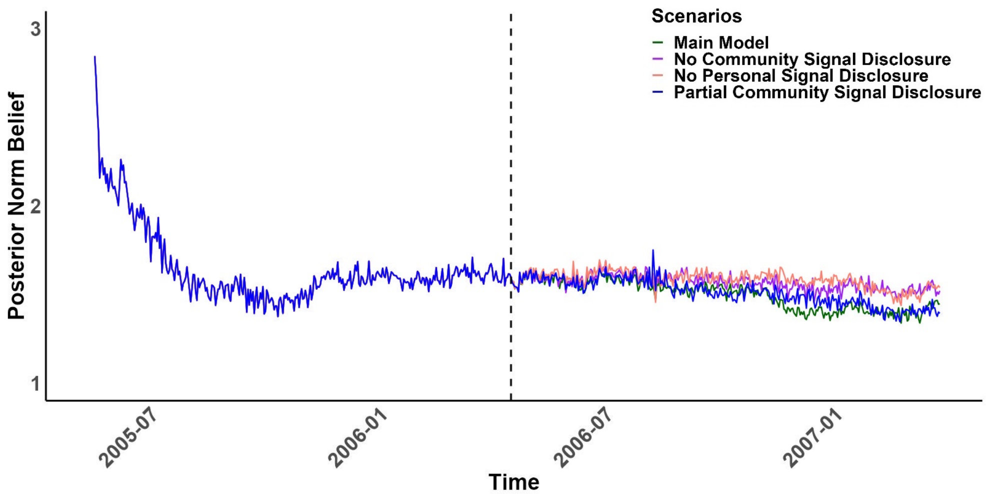

Figure 9: Posterior Norm Over Time

The results for tipping decisions are shown in Table 6. The percentage changes in Table 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 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 5 with the exception of a small increase in the number of unique tippers instead of a decrease.

Table 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, we have shown that there is some stickiness in users' norm belief and tipping behavior. Next, we want to quantify the amount of stickiness in terms of its medium-run effects on tipping behavior. We do so as follows: for each of the three scenarios, we 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, we 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 7.

Table 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 7). Here, it is important to keep in mind that the results in Table 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 we 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 7.

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. We study how these beliefs about the current tipping norm, combined with content quality and other factors, influence tipping behavior.

Our 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, we 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.

We then analyze the impact of different information provision strategies by partially or fully removing the visibility of personal and community signals. Our 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, we also examine how sticky the perceived tipping norm and tipping behavior is after a change in the platform's information disclosure. We 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 our findings, platforms can ensure a steady stream of income from tipping activities. Furthermore, our 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.

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 (2023), "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.

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.

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

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, 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.

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

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.

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

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.

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.

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

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.

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

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.

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

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

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.

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.

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.

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.

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.

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.

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

Web Appendix A: Derivation of Bayesian Updating Formulas

A.1 Bayesian Updating with Independent Signals

In this section of the appendix, we 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, we 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, we 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:

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

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

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

Simplifying the exponentials and combining terms involving μ, we derive the expression:

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

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, we 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

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

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

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, we 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 source as well as between signals from different sources. We 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.

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

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

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

$$\omega_{it} = \omega_0 + \sum_{k,z=1,1}^{M^p_{it}+M^c_{it}} \omega_{k,z} \ , \tag{A16}$$

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

A.3 Bayesian Updating with Desired Correlation Structure

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

$$\begin{aligned} \psi_{it} = [&S^p_{i1,1}, S^p_{i1,2}, \ldots, S^p_{i1,N^p_{i1}}, S^p_{i2,1}, S^p_{i2,2}, \ldots, S^p_{i2,N^p_{i2}}, \ldots, S^p_{it,1}, S^p_{it,2}, \ldots, S^p_{it,N^p_{it}}, \\ &S^c_{i1,1}, S^c_{i1,2}, \ldots, S^c_{i1,N^c_{i1}}, S^c_{i2,1}, S^c_{i2,2}, \ldots, S^c_{i2,N^c_{i2}}, \ldots, S^c_{it,1}, S^c_{it,2}, \ldots, S^c_{it,N^c_{it}}] \end{aligned} \tag{A18}$$

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

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-source correlations among signals and Ω_{pc} corresponds to the correlations between signals from different sources. 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 source 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 sources. 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 $N^p_{iz} \times N^c_{ik}$ for personal signals of time period z and community signals of time period k with all elements equal to $\delta_{pc}^{|z-k|}\lambda_{pc}$:

$$\Omega_{pc} = \lambda_{pc} \begin{bmatrix} J_{N^p_{i1} \times N^c_{i1}} & \delta^1_{pc} J_{N^p_{i1} \times N^c_{i2}} & \cdots & \delta^t_{pc} J_{N^p_{i1} \times N^c_{it}} \\ \delta^1_{pc} J_{N^p_{i2} \times N^c_{i1}} & J_{N^p_{i2} \times N^c_{i2}} & \cdots & \delta^{t-1}_{pc} J_{N^p_{i2} \times N^c_{it}} \\ \vdots & \vdots & \ddots & \vdots \\ \delta^t_{pc}\ J_{N^p_{it} \times N^c_{i1}} & \delta^{t-1}_{pc} J_{N^p_{it} \times N^c_{i2}} & \cdots & J_{N^p_{it} \times N^c_{it}} \end{bmatrix} \tag{A22}$$

We now proceed to calculate the posterior mean and precision given the structure for Ω, beginning with ω_t. First, we calculate the sum of elements in Ω_p and Ω_c. We 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:

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

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

$$N^p_{iz} \times N^p_{ik} \times \lambda_p. \tag{A24}$$

Therefore, the total sum for Ω_p is given by

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

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

$$\lambda_{pc} \sum_{k,z=1,1}^{t,t} \delta_{pc}^{|k-z|} N^p_{iz} \times N^c_{ik}. \tag{A26}$$

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, we turn to the calculation of the posterior mean μ_t. First, we 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, we have:

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

For the off-diagonal blocks Ω^p_{zk} in Ω_p, we have:

$$\begin{aligned}\sum_{r=1}^{N^p_{iz}}\sum_{q=1}^{N^p_{ik}}\omega_{rq}(s^p_{iz,r}+s^p_{ik,q}) =&\lambda_p(s^p_{iz,1}+s^p_{ik,1})+\lambda_p(s^p_{iz,1}+s^p_{ik,2})+\cdots+\lambda_p(s^p_{iz,1}+s^p_{ik,N^p_{ik}})+\\ &\lambda_p(s^p_{iz,2}+s^p_{ik,1})+\lambda_p(s^p_{iz,2}+s^p_{ik,2})+\cdots+\lambda_p(s^p_{iz,2}+s^p_{ik,N^p_{ik}})+\\ &\vdots\\ &\lambda_p(s^p_{iz,N^p_{iz}}+s^p_{ik,1})+\lambda_p(s^p_{iz,N^p_{iz}}+s^p_{ik,2})+\cdots+\lambda_p(s^p_{iz,N^p_{iz}}+s^p_{ik,N^p_{ik}})\\ =&N^p_{ik}\lambda_p\mathbb{S}^p_{iz}+N^p_{iz}\lambda_p\mathbb{S}^p_{ik}.\end{aligned} \tag{A29}$$

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

$$\omega_p\sum_{k=1}^{t}\mathbb{S}^p_{ik}+\lambda_p\sum_{k=1}^{t}(N^p_{ik}-1)\mathbb{S}^p_{ik}+\frac{\lambda_p}{2}\sum_{k,z=1,k\neq z}^{t}\delta_p^{|k-z|}\left(N^p_{ik}\mathbb{S}^p_{iz}+N^p_{iz}\mathbb{S}^p_{ik}\right) \tag{A30}$$

.

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

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

.

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

$$\sum_{r=1}^{N^p_{iz}}\sum_{q=1}^{N^c_{ik}}\omega_{rq}(s^p_{iz,r}+s^c_{ik,q})=\lambda_{pc}\sum_{k,z=1}^{t}\delta_{pc}^{|k-z|}\left(N^c_{ik}\mathbb{S}^p_{iz}+N^p_{iz}\mathbb{S}^c_{ik}\right). \tag{A32}$$

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

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

to compute the posterior more efficiently in the optimization process, we 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 $\sum_{k,z=1}^{t}\left(\delta_p^{|k-z|}N_{ik}^{p}\mathbb{S}_{iz}^{p}\right)$ 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

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

Web Appendix B: Supplements for Predictive Exercises

We 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. We 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)

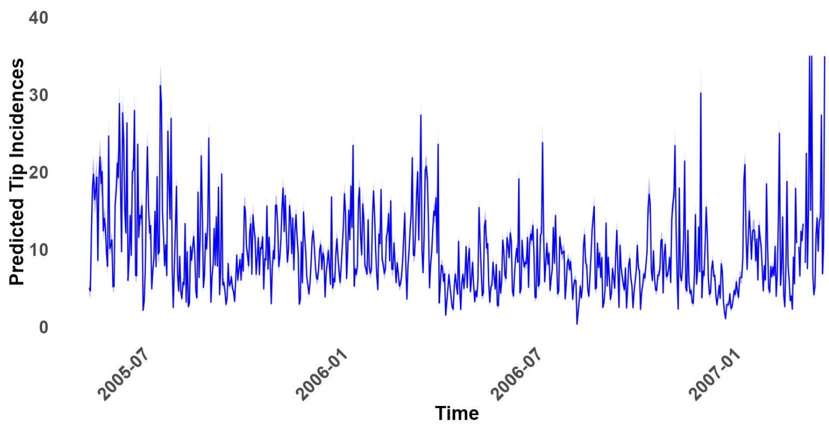

(a) Predicted Number of Tip Incidences

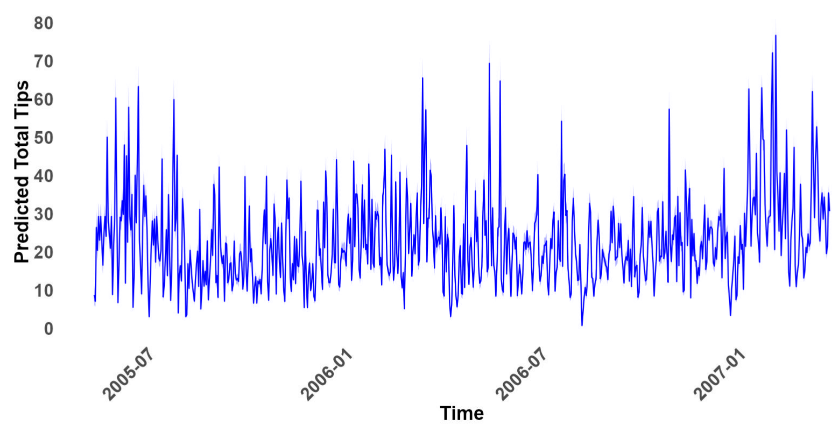

(b) Predicted Total Tips Amount

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