Valuing influence with social learning

Hyun-Soo Ahn[1], Christopher Thomas Ryan[2], Joline Uichanco[3], and Mengzhenyu Zhang[4]

[1]University of Michigan, Ross School of Business, hsahn@umich.edu
[2]University of British Columbia, Sauder School of Business, chris.ryan@sauder.ubc.ca
[3]New York University, Tandon School of Engineering, ju441@nyu.edu
[4]University College London, School of Management, zhenyu.zhangmeng@ucl.ac.uk

Abstract

Problem definition: Influencer marketing has become a prevalent strategy to promote products through social media. This paper examines the value of influencer marketing when followers not only learn from the influencer's signal but can also engage in social learning by observing peers' purchase behaviors and reviews. **Methodology/results:** We adopt an information design framework to analyze how a firm should value an influencer based on two key dimensions: the accuracy of the influencer's past recommendations (*informativeness*) and the extent to which followers rely exclusively on the influencer versus learning from peers (*charisma*). **Managerial implications:** Our model uncovers insights about the interaction between information design and social learning. First, the naive intuition that the influencer is less valuable with social learning does not always hold. The influencer holds greater value under the social learning context when customers have a moderate intention to buy, as her endorsement reinforces customer convictions, making them resilient against later negative feedback from other followers. Second, when the firm can strategically select an influencer, the optimal information structure is biased towards the positive signals: always endorse good products (true positive rate of 1) but sometimes endorse bad products (nonzero false positive rate). Third, the optimal influencer when social learning exists has a lower false positive rate than the one without social learning, meaning that, when there exists subsequent social learning, it becomes even more important to have an influencer whose positive endorsement is trustworthy. In other words, the optimal influencer should be able to reveal more information with social learning than without.

1 Introduction

Promoting products through social media influencers has grown dramatically over the past decade. Influencers create content and share their impressions of products with followers on platforms such as YouTube, Instagram, TikTok, and Facebook. Influencer Marketing Hub, a platform that collects data on influencer marketing, reports that the influencer marketing industry grew to $21.1 billion in 2023, up from $1.7 billion at the beginning of 2016 (InfluencerMarketingHub, 2024). One source reports that the top ten highest-paid content creators (influencers) on YouTube earned $10.5 million to $29.5 million in 2024 (Bennett, 2024).

Companies increasingly recognize the value of influencers' ability to drive sales. A key reason for partnering with influencers rather than communicating directly with consumers is credibility (Lin and Liu, 2024). Customers often perceive companies as overly optimistic and willing to promote even poor products. As a result, company-generated messages are frequently regarded as cheap talk and are heavily discounted by consumers. In contrast, messages from reputable influencers—those who have built their credibility through informative and honest reviews—can carry significantly greater persuasive power.[1] Influencers can also communicate with their followers in ways that companies cannot, such as showcasing a makeup kit's pros and cons through a live demonstration on Instagram. For these reasons, many companies invest heavily in influencer partnerships despite substantial costs.

However, finding effective influencers remains challenging. A survey found that 61% of firms struggled to find the "right" influencers. The difficulty arises because the influencer landscape is vast and diverse. In 2020, there were over 50 million influencers worldwide (Gagliese, 2022). These influencers greatly differ product category, channels, follower size, number of active followers, informativeness, and credibility (Geyser, 2024). This challenge has led to specialized agencies and platforms that identify and recruit influencers. For example, Upfluence.com maintains a database of three million influencers with over 20 search parameters. Despite these tools, quantifying an influencer's value remains challenging. Our paper addresses this gap by answering three questions:

(Q1) When is it advantageous to use an influencer, and when is it not? How much benefit can a firm derive from working with an influencer with specific characteristics?

(Q2) How does social learning among followers—where customers observe peers' purchases and reviews—affect an influencer's value? Does organic social learning make the influencer more or less valuable?

(Q3) Who is the ideal influencer for a given product? Should the ideal influencer in environments with strong social learning reveal more or less information than in environments without social learning?

While characteristics such as follower count and platform are easily observable, much less is known about an influencer's *informational value*. Our focus is to study this informational value and how it changes when customers update beliefs through social learning.

We model an influencer's value along two dimensions of persuasiveness. First, *informativeness* represents the influencer's precision in revealing the true product quality. An influencer who positively evaluates every product (even poor ones) has different informational value than one who is more selective. We explicitly model how followers form beliefs based on this informativeness. Second, *charisma* captures the extent to which followers rely exclusively on the influencer's endorsement versus seeking additional information. Some followers are *devotees* who listen solely to the influencer, while others are *skeptics* who also gather information from peers through purchase behaviors and reviews—a process we call *social learning*.[2] Charisma measures

[1] Of course, not all influencers are truthful and highly informative. The characteristics of an influencer are the main features of our model.

[2] Although the influencer is part of the social network, we do not consider the influencer's signal as social

the fraction of devotees among followers, reflecting how much the influencer's initial signal is diluted by social learning.

These two dimensions of persuasiveness—charisma and informativeness—jointly determine how follower beliefs evolve in response to the influencer's message and subsequent peer actions, ultimately impacting the firm's revenue. We provide a theoretical framework to quantify the value of an influencer in promoting a product based on these dimensions. To our knowledge, this is the first model that analyzes how the informational value of an influencer's message changes in these two dimensions, going beyond easily measurable factors such as follower count and message format. Our paper also extends the Bayesian persuasion literature (Kamenica and Gentzkow, 2011) to settings with organic social learning.

Main findings. Our analysis yields insights about when and why influencers create value:

(F1) A firm should consider not only how positive an influencer's messaging is, but also how discerning and selective she is in her endorsements. This insight emerges from our model, which shows that an influencer's false positive and false negative errors have an asymmetric impact on the value of influence. Specifically, a high rate of false positives (i.e., endorsing a bad product) is more detrimental than false negatives (i.e., not recommending a product that is actually good).

(F2) When customers need to be persuaded to purchase, it is beneficial for a firm to collaborate with an influencer, regardless of how skeptical her followers are (Proposition 2, Lemma 6, and Theorem 1).

(F3) Surprisingly, when customers already hold a high prior belief, sponsoring an influencer can still be beneficial. However, this is only the case when a sufficiently large proportion of her followers are skeptics who engage in *social learning*. Otherwise, if a large proportion are devotees who exclusively rely on the influencer's endorsement, then the influencer can actually decrease the firm's expected revenue. In the case of social learning, when customers' prior belief is moderately high, value creation happens because the influencer's signal strengthens followers' beliefs so that they are robust against any subsequent negative signals by other followers (Theorem 1). This finding is counterintuitive. One might expect that in the presence of social learning, the influencer's signal value might be diminished because of the large amount of information from social learning. Our results show that the instances that make the influencer valuable are even richer than in the no-social-learning environment.

(F4) Regardless of the composition of followers, it is always beneficial for a firm to partner with an influencer who is effective in lowering followers' beliefs about a product's quality after sending a negative signal (Proposition 2, Proposition 5, and Theorem 2). In practical terms, this means that an influencer who rarely makes false negative errors (i.e., one who is rarely dismissive of good products) is more likely to be valuable to the firm.

(F5) When firms can choose among influencers, the optimal information structure has an h-

learning because we view the influencer as implementing the firm's information policy rather than being socially motivated.

biased structure: *always* endorse a "good" product but *sometimes* endorse a "bad" product (Theorem 2). Surprisingly, the optimal rate of false positives (endorsing a bad product) is *lower* with social learning than without, meaning firms should seek more accurate influencers when social learning is present. This occurs because greater belief spread after the message of an accurate influencer interacts beneficially with the diffusion process. One might expect abundant organic information to make precise influencer signals less important, but the opposite holds: informativeness becomes more valuable precisely when social learning is active.

A critical factor that makes our model insightful (and challenging) is modeling the dynamic interactions through which skeptics update beliefs via social learning. We handle this complexity using a diffusion approximation (Glynn, 1990), which makes the model tractable without losing the effect of social learning on the revenue. This tractable approximation makes comparison among different influencers possible. A static model fails to capture the natural evolution of followers' beliefs over time as their experience with the product accumulates, and consequently, it cannot provide our insights (see Appendix H).

In summary, this paper makes four contributions. First, we provide a tractable framework for quantifying influencer value that accounts for both the influencer's characteristics (informativeness and charisma) and the information environment (presence or absence of social learning). Second, we leverage asymptotic analysis via diffusion limit to extract valuable insights on the value of influencers with social learning, a problem that is generally not well-behaved and does not lead to analytical results. Third, we characterize optimal information design in the presence of social learning, extending Bayesian persuasion to dynamic environments where receivers learn from each other. Finally, we analyze the setting where the firm has a private signal of the product state and show that private information forces full disclosure, thereby undermining the strategic value of information design.

The rest of the paper proceeds as follows. Section 2 reviews related literature. Section 3 introduces the model framework. Section 4 analyzes how to value a given influencer. Section 5 examines how the degree of social learning changes influencer value. Section 6 characterizes optimal influencer selection. Section 7 studies privately informed firms and shows how private information forces full disclosure. Section 8 concludes.

2 Connections to existing literature

Our model and results are closely related to two areas: (i) the marketing and economics literature on information transmission from firms to customers and (ii) operations and computer science literature on algorithmic information design and optimal signaling mechanisms.

The problem of how firms use information to manipulate the opinions of customers has been extensively studied in the marketing literature. This includes papers that have explored mechanisms for shaping the opinions of their customers. A prime example is the manipulation of online reviews by firms (e.g., Mayzlin 2006 and Dellarocas 2006). However, today's consumers are becoming more discerning and aware of the potential manipulation of anonymous information and online opinions. This awareness partially explains the rise in prominence of social media influencers whose success relies on maintaining credibility in online spaces. Consequently, it is the goal of firms not to manipulate the messages of influencers but rather to select those whose genuine and unaltered messaging naturally aligns with the firm's interests. Jalloul and Kostami (2022) recognize the effectiveness of influencer marketing using a two-period model.

Influencers, however, do not operate in a vacuum. It is important to consider their followers' effect on each other. Kozinets et al. (2010) studied word-of-mouth marketing among bloggers—influencers who create written content on a blog—and found that, after the initial campaign by the blogger, the network co-production of the marketing message by the blog readers can shape the narrative. Using empirical evidence, Hamami (2019) found that product reviews from early adopters can change the beliefs about product quality among later-arriving customers. Effects of this type are captured in our model by the existence of skeptic followers who "shape" the influencer's message and must be considered in valuing an influencer's post.

Our framework also contributes to the economics literature on information transmission and persuasion. The economics literature has long understood the importance of credibility in communication (Crawford and Sobel, 1982). To effectively model information transmission, researcher have developed models of communication with commitment power, as explored in the Bayesian persuasion literature, which hold great relevance for studying word-of-mouth marketing. Kamenica and Gentzkow (2011) and Rayo and Segal (2010) provide conditions under which a persuasive signal is effective, i.e., the sender achieves a higher expected payoff by committing to a signaling strategy. Our work is closely related to the persuasion literature because the influencer's signaling strategy cannot be changed in the short run, thereby adhering to the notion of commitment. The firm only benefits from sponsoring an influencer if the influencer's signaling strategy can "persuade" more followers to purchase. Moreover, our setting does not suffer from the implementability issue faced by traditional Bayesian persuasion models as different influencers are practical vehicles to implement the information policy.

Since the seminal work of Kamenica and Gentzkow (2011), the study of Bayesian persuasion has developed rapidly. The most relevant strand to our setting is generalizations to multiple receivers. This literature can be divided into two settings, whether the sender sends private signals or public signals. Arieli and Babichenko (2019) and Candogan and Drakopoulos (2020) study the optimal signaling mechanism in the setting where the sender sends private (possibly different) signals to different receivers. Our focus is on public signaling mechanisms. In practice, a social media influencer sends public messages to all of her followers. This public display of "content" is the key driver of follower engagement. For Bayesian persuasion literature with public signals, Alonso and Câmara (2016a) and Alonso and Câmara (2016b) consider the public signaling mechanism in a voting setting. The most significant distinction with our model is that we allow a certain group of receivers to learn not only from the signal sent by the sender but also from the purchasing behavior of others in the social network. Candogan (2019) studies a persuasion game where receivers are socially connected via a network. He proposes tractable algorithms to solve for the static equilibrium. Each receiver's utility depends on the state of the world and the decisions of others in the network. Our work differs from his by specifying the dynamic learning of receivers from the organic information generated within the system. Specifically, we model decision dependency among receivers via their dynamic learning behavior.

Applications of the Bayesian persuasion framework in marketing and operations management applications have developed rapidly in recent years. Boyaci et al. (2022) is the most relevant and recent paper using these tools to study influencer marketing. They also use the Kamenica and Gentzkow (2011) framework, where the firm can commit to a static information structure. However, our model differs from theirs because we model the dynamic social learning behavior among followers, which is triggered by the influencer's signal. In a relevant context but not adopting the Bayesian persuasion framework, Fainmesser and Galeotti (2021) analyze the equilibrium in a market with many influencers, followers, and marketers where there is competition between intermediaries (influencers) to send the signals. Mostagir and Siderius (2023) study the dynamic decision of reviewers over time to maintain their reputation while accepting "bribes". Simplifying assumptions are made to study the competition among many influencers and the dynamic choice of reviewers. Fainmesser and Galeotti (2021) use a "reduced-form" approach to model the relationship between an influencer and a follower, and Mostagir and Siderius (2023) model reviewers to have either a high or low type in terms of "skill". Our work (and similarly, Boyaci et al. 2022) comprehensively analyzes the information transmission from the influencer to the followers. Küçükgül et al. (2022) use a dynamic Bayesian persuasion framework to study the optimal information policy when customer arrivals are deterministic and cannot directly observe the actions/reviews realized by previous customers. To the best of our knowledge, we are the first to consider a Bayesian persuasion sender (the influencer) followed by stochastically and sequentially arriving receivers (followers) that can learn from one another (social learning), and strive to understand how social learning distorts the initial information policy.

Indeed, a firm's various decisions can be impacted by social learning. Ifrach et al. (2019) considers the pricing decision when customers sequentially learn from reviews in a Bayesian manner using the asymptotic analysis of the underlying stochastic system. Product ranking decision is considered in Maglaras et al. (2022). Besbes and Scarsini (2018) characterize the statistical properties of customers' beliefs under social learning by assuming different degrees of observations, ranging from observing all past reviews (fully rational) to only observing the average of past reviews (boundedly rational). Papanastasiou and Savva (2016) study dynamic pricing in environments where demand is shaped by social learning. In their model, consumers revise their beliefs after observing the purchasing behavior of others, affecting their decision to buy immediately or wait. In contrast, Debo and Secomandi (2018) examine pricing when customers update their beliefs about product value based on private signals that remain unobserved by both other consumers and the seller. We are the first to consider optimal information disclosure in the presence of social learning.

In our setting, we assume that influencers have the commitment power to a static information structure. Consumers observe the influencer's historical reviews and learn how accurate she is by comparing her reviews to the ex-post-realized success of the products. These consumers' beliefs

Table 1: Utility U for different combinations of decision and state

Decision \ State	H	L
Purchase	$g-p$	$-p$
Do not purchase	0	0

about the influencer's informativeness cannot be established within one or two reviews but rather over a long period of time. Therefore, practically, the assumption that the influencer can commit to dramatically changing "informativeness" for the review of a single product could be fragile. Given that little has been done to understand the value of influence where the signal from the influencer triggers a series of social learning among followers, we focus on the static information policy with subsequent social learning in this paper. It would be interesting to explore a *dynamic* information disclosure policy in the presence of social learning under an applicable context in future research.

3 Model

A monopolist firm sells a new product at a fixed price p during a selling season of length T. Before the season begins, the firm can sponsor an influencer to promote the product to her followers. The product's true value is uncertain at the outset: the state $w \in \Omega \triangleq \{H, L\}$ is either high (H), meaning the product is highly valued by the followers (a "good" product), or low (L), meaning the product has low value (a "bad" product).

The true state is unknown to the firm, the influencer, and the followers. This uncertainty arises because no one knows whether the product's technology is sufficiently mature, whether it will achieve widespread adoption, or, in the context of fashion products, whether it will gain social acceptance. At time zero, the firm and the followers share a common prior belief about the state: $\big(\mu^0(H), \mu^0(L)\big)$, where $\mu^0(H)$ is the probability that the product is in the high state and $\mu^0(L)$ is the probability that it is in the low state.

3.1 Customer purchase decisions

Customers arrive according to a Poisson process with rate λ. Upon arrival, each customer decides whether to purchase based on her expected utility, which depends on her belief about the state and the utility structure shown in Table 1. We normalize the utility from not purchasing to zero. A customer who purchases receives utility $g - p$ if the state is H and utility $-p$ if the state is L, where $g > 0$ represents the value of a good product (a bad product has 0 value).

Consider a customer who arrives with belief y that the state is H. Her expected utility from purchasing is $-p + yg$, while her expected utility from not purchasing is 0. She purchases if and only if $-p + yg \geq 0$, or equivalently, $y \geq p/g$:

$$\text{Decision}(y) = \begin{cases} \text{purchase}, & \text{if } y \geq p/g, \\ \text{not purchase}, & \text{otherwise.} \end{cases} \tag{1}$$

The threshold p/g is the price per unit value of a good product, and the customer will purchase her belief exceeds this threshold. The belief y is formed by updating the customer's prior μ^0 using any information the customer has received, such as a signal from a sponsored influencer or signals generated by other customers' purchases.

3.2 Influencer signal and Bayesian updating

If the firm sponsors the influencer, the influencer sends a signal to her followers at time zero, before customers begin arriving. In practice, sponsorship involves sending a product sample and paying a fee in exchange for a post (e.g., a blog entry, YouTube video, TikTok clip, or Instagram post) featuring the product. After evaluating the sample, the influencer creates her post based on her experience, which provides a signal about the product's state.

We model the influencer's post as a signal $s \in S \triangleq \{h, \ell\}$ that is observable to all her followers, where h represents an endorsement of the product (high signal) and ℓ represents no endorsement (low signal). The signal $s \in \{h, \ell\}$ is realized with probability $\pi(s|w)$ when the true state is $w \in \{H, L\}$, where the distribution π is publicly known from the influencer's posting history. We represent π as a two-by-two matrix:

$$\pi \triangleq \begin{pmatrix} \pi_{hH} & \pi_{\ell H} \\ \pi_{hL} & \pi_{\ell L} \end{pmatrix} \triangleq \begin{pmatrix} \pi(h|H) & \pi(\ell|H) \\ \pi(h|L) & \pi(\ell|L) \end{pmatrix} \tag{2}$$

where the middle notation is condensed for convenience. We call π the *informativeness* of the influencer and refer to an influencer with informativeness π as a π-influencer. The matrix π corresponds to a statistical experiment in the sense of Blackwell and Girshick (1979), which is also called an information structure or signal structure in the information design literature (Kamenica and Gentzkow, 2011).

The term "informativeness" reflects that π captures the historical accuracy of the influencer's endorsements as observed from past data. For example, an influencer who has correctly identified "good" products 80% of the time and "bad" products 70% of the time would have $\pi_{hH} = 0.8$, $\pi_{\ell H} = 0.2$, $\pi_{\ell L} = 0.7$, and $\pi_{hL} = 0.3$. Observers can assess an influencer's informativeness by comparing her past signals with the ex-post realized states of previously reviewed products.

A π-influencer's signal updates followers' beliefs from the prior to a post-signal prior[3], calculated according to Bayes' rule. Given the realized signal $s \in \{h, \ell\}$, the post-signal prior belief of state $w \in \{H, L\}$ is:

$$\mu(w|s) \triangleq \frac{\mu^0(w)\pi(s|w)}{\sum_{w' \in \Omega} \mu^0(w')\pi(s|w')} \tag{3}$$

Since $\mu(L|s) = 1 - \mu(H|s)$ for any signal s, it suffices to track only the belief in the high state. Doing so simplifies our notation. We denote the post-signal prior belief in the high state after a high signal as $\mu^h \triangleq \mu(H|h)$ and after a low signal as $\mu^\ell \triangleq \mu(H|\ell)$. Similarly, we use scalar μ^0 to denote the prior belief in the high state $\mu^0(H)$.

[3]We use the term "post-signal prior" instead of "posterior" to distinguish the fact that, in Section 4.3, the beliefs of followers can evolve after the influencer's signal is released. This makes the generic term "posterior" somewhat ambiguous. We maintain the word "prior" in "post-signal prior" to connote that this belief is prior to the dynamics of arriving followers.

Note that μ^h and μ^ℓ are functions of μ^0 via (3). Specifically, for signal $s \in \{h, \ell\}$, we have:

$$\mu^s = \frac{\mu^0 \pi_{sH}}{\mu^0 \pi_{sH} + (1-\mu^0)\pi_{sL}} \tag{4}$$

Throughout the paper, we assume

$$\mu^\ell \leq \mu^h, \tag{5}$$

which implies $\pi_{\ell L} + \pi_{hH} \geq 1$. This condition states that the belief in a good product is greater when receiving a high signal compared to receiving a low signal. This assumption is without loss of generality because, otherwise, followers can "bet against" the influencer's recommendation by interpreting ℓ as the high signal.

3.3 Post-signal dynamics

At time zero, all customers hold a common initial belief: the post-signal prior (μ^ℓ or μ^h) if the influencer is sponsored, or the initial prior (μ^0) otherwise. We now describe the subsequent dynamics during the selling season $[0, T]$ as customers arrive and make purchase decisions, where according to (1) a customer purchases if his belief at the time of arrival exceeds threshold p/g.

Customers are heterogeneous and fall into two categories: devotees and skeptics. An arriving follower is a devotee with probability α and a skeptic with probability $1-\alpha$. We call α the influencer's *charisma*, reflecting the extent to which followers adhere to the influencer's signal while ignoring information from others. *Devotees* do not engage in social learning and rely exclusively on the initial common belief when making purchasing decisions. *Skeptics* engage in social learning: they start with this common belief but update it dynamically based on additional information generated organically by other customers. We denote by μ_t the belief of a skeptic at time t. Consistent with the social learning literature (e.g., Ifrach et al. 2019), information is generated by followers who have bought the product (for example, the purchase behavior of other followers and feedback posted as comments on the influencer's endorsed posts). In this way, customers become co-producers of market signals (Kozinets et al., 2010), generating information that affects subsequent arrivals. This additional information is sometimes called *organic* (Fainmesser and Galeotti, 2021) because it evolves spontaneously among consumers.

Let $D_t(\mu)$ denote the cumulative number of purchases by time t when the initial common belief is μ, where $\mu \in \{\mu^0, \mu^h, \mu^\ell\}$. The dynamics of the demand process $D_t(\mu)$ depends on the influencer's charisma. Without social learning ($\alpha = 1$), $D_t(\mu)$ is a homogeneous Poisson process (Section 4.2). With social learning ($\alpha < 1$), the purchase rate is affected by skeptics' dynamically evolving beliefs μ_t, and the information generated from purchases feeds back into these beliefs, creating a feedback loop (Section 4.3).

Given the initial common belief $\mu \in \{\mu^0, \mu^h, \mu^\ell\}$, the firm's expected revenue is

$$v(\mu) \triangleq \mathbb{E}\left[\int_0^T p \mathrm{d}D_t(\mu)\right], \tag{6}$$

where p is the selling price. Hence, the expected revenue from sponsoring a π-influencer is:

$$\sum_{w \in \Omega} \mu^0(w) \sum_{s \in S} \pi(s|w) v(\mu^s). \tag{7}$$

This expectation has two layers: the outside layer is an expectation over the unknown state, and the inner layer is an expectation over the unknown signal of the influencer (which is itself conditional on the state).

Remark 1. Although we refer to a specific influencer as a π-influencer, note that there are three key characteristics that describe an influencer and her followers: informativeness π, charisma α, and follower size (embedded in the arrival rate λ).

4 Valuing an influencer

We study a firm's decision of whether to sponsor a specific influencer, which we refer to as the *sponsorship problem*. The firm will sponsor the influencer when the expected benefit outweighs the cost. The cost is in the form of influencer fees that are often exogenously determined by market rates (e.g., $500 per Instagram post for a micro-influencer, Hitchcock 2025). The sponsorship problem can be easily extended to a setting where the firm has to choose the influencer from a short list provided by a marketing agency. For instance, if the firm has to choose from k finite influencers with each costing c_k, the optimal influencer can be found by comparing the net benefits of the k influencers.

On the other hand, the benefit of sponsoring an influencer is the change in the firm's expected revenue with and without the influencer's signal. For a π-influencer, this can be calculated as:

$$\Pi(\pi, \mu^0) \triangleq \sum_{w \in \Omega} \mu^0(w) \sum_{s \in S} \pi(s|w) v(\mu^s) - v(\mu^0), \tag{8}$$

which represents the expected revenue gain from the influencer's signal updating the initial common belief from μ^0 to μ^s.

We refer to Π as the *value of the influencer*. Our goal in the remainder of this section is to characterize this value. Section 4.1 presents some preliminaries that simplify our analysis. We then gain insights into the interaction between influence and social learning by studying how the influencer's value changes under two scenarios: one without social learning (Section 4.2) and one with social learning (Section 4.3).

4.1 Some preliminaries

We first present preliminaries. For readers familiar with the Kamenica and Gentzkow (2011) framework, these preliminaries should be familiar. We include these preliminaries to make our work self-contained and to clarify our main analytical technique, which relies on visualizing the influencer's value Π by plotting the expected revenue function $v(\cdot)$ (see Figure 1 below). The technique allows our analysis to proceed visually (as will be seen in Figures 2, 4, 5, 8, 9 and 13 throughout the rest of the paper).

The technique is based on a simple transformation. Given the influencer's informativeness π and the prior μ^0, we define a distribution $\tau(\cdot)$ over the set of post-signal beliefs in $[0, 1]$. Since there are two possible signals $|S| = 2$, there are two possible post-signal beliefs μ^ℓ and μ^h (determined *a posteriori* after observing the signal). However, before the signal is realized, we

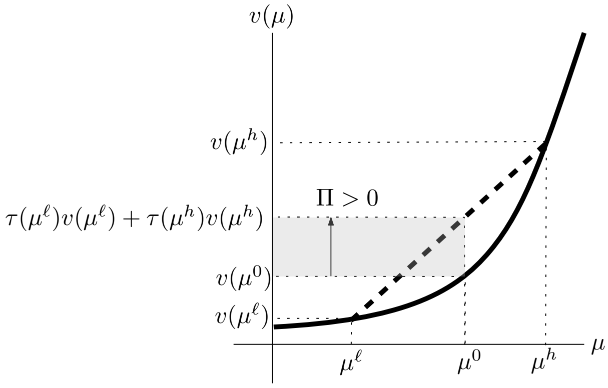

Figure 1: Visualization of Π in (10) as the vertical height of the gray-colored rectangle. By (10), the horizontal axis is the post-signal prior z.

can compute the *a priori* probability of each post-signal belief as:

$$\tau(\mu^s) \triangleq \sum_{w\in\Omega} \mu^0(w)\pi(s \mid w). \tag{9}$$

Hence, *a priori* of observing the signal, the expected revenue from sponsoring an influencer is $\tau(\mu^h)v(\mu^h) + \tau(\mu^\ell)v(\mu^\ell)$. This observation leads to the following lemma which reformulates the influencer's value Π using the distribution $\tau(\cdot)$.

Lemma 1. The value of the influencer in (8) is equivalent to:

$$\Pi(\pi, \mu^0) = \Pi(\tau, \mu^0) = \sum_{s\in\{\ell,h\}} \tau(\mu^s)v(\mu^s) - v(\mu^0) \tag{10}$$

where $\tau(\cdot)$ and μ^0 satisfy

$$\tau(\mu^\ell)\mu^\ell + \tau(\mu^h)\mu^h = \mu^0. \tag{11}$$

Proof. See Appendix D.1.

Equation (10) shows that the influencer's value can be expressed as a weighted average of revenues under different post-signal beliefs, minus the revenue without the influencer. The weights are the *a priori* probabilities $\tau(\mu^h)$ and $\tau(\mu^\ell)$ of each belief occurring. Condition (11), known as "Bayes plausibility," ensures consistency: the expected post-signal belief (weighted by the probabilities of each signal) must equal the prior belief μ^0. This prevents the influencer from systematically biasing beliefs upward or downward on average—any change in beliefs must balance out across the two possible signals. Interested readers are referred to Appendix B for a concrete numerical example.

Geometrically, this lemma allows us to visualize the influencer's value using the revenue function $v(\cdot)$. For given post-signal beliefs μ^h and μ^ℓ, the influencer's value Π is the vertical distance between the revenue function $v(\cdot)$ evaluated at the prior μ^0 and the secant line connecting the points $(\mu^\ell, v(\mu^\ell))$ and $(\mu^h, v(\mu^h))$. This is illustrated in Figure 1. Importantly, by the Bayes plausibility condition, the prior μ^0 always lies between μ^ℓ and μ^h, ensuring this geometric interpretation is well-defined.

This geometric perspective reveals when the influencer creates value. If $v(\cdot)$ is convex over the interval $[\mu^h, \mu^\ell]$, then $\Pi \geq 0$: the influencer is valuable when the revenue function curves upward between the two possible post-signal beliefs. Conversely, if $v(\cdot)$ is concave over this interval, then $\Pi \leq 0$: the influencer destroys value when the revenue function curves downward. The intuition is straightforward: a convex revenue function means the weighted average of revenues at the extreme beliefs exceeds the revenue at the average belief, while a concave function produces the opposite result.

4.2 Baseline setting: No social learning

We first analyze a baseline model where all customers are devotees ($\alpha = 1$). In this setting, customers form their beliefs exclusively from the influencer's signal and do not engage in social learning. We use the term "social learning" to refer to learning from additional signals or reviews besides the influencer's message. This baseline case allows us to obtain closed-form solutions for the influencer's value and establish a benchmark for comparison with the social learning case studied in Section 4.3.

For a given arrival rate λ and a common prior $\mu \in \{\mu^0, \mu^h, \mu^\ell\}$, the demand process is a Poisson process with a constant rate

$$\Lambda(\mu) \triangleq \begin{cases} \lambda & \text{if } \mu \geq p/g, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

This result comes directly from the threshold decision rule in (1). Using (6), the expected revenue $v(\mu)$ for a given prior belief μ is:

$$v(\mu) = p\mathbb{E}\left[\int_0^T \mathrm{d}D_t(\mu) \mid \mathcal{F}_0\right] = p\mathbb{E}\left[\int_0^T \mathbb{E}\left[\mathrm{d}D_t(\mu) \mid \mathcal{F}_t\right] \mid \mathcal{F}_0\right] = p\Lambda(\mu)T \tag{13}$$

where the first equality is by definition, the second equality uses the tower rule of conditional expectation, and the third evaluates the expected total arrivals in $[0, T]$ given a Poisson rate $\Lambda(\mu)$. Therefore, $v(\cdot)$ is a step function with a discontinuity at $\mu = p/g$, as illustrated in Figure 2.

We next examine how the value of an influencer is affected by the prior μ^0 and her informativeness π. Following the Bayesian persuasion literature, we express this value as a function of the post-signal priors, μ^ℓ and μ^h, which depend on μ^0 and π through (4).

Proposition 1 (Value of the influencer without social learning)**.** By (5) and (11), we have $\mu^\ell \leq \mu^0 \leq \mu^h$. When all customers are devotees that exclusively learn from the influencer's signal, the influencer's value is determined as follows:

(i) If $\mu^\ell \leq p/g \leq \mu^h$ and $\mu^0 \leq p/g$, then $\Pi \geq 0$ with

$$\Pi = p\lambda T \cdot \frac{\mu^0 - \mu^\ell}{\mu^h - \mu^\ell}. \tag{14}$$

(ii) If $\mu^\ell \leq p/g \leq \mu^h$ and $p/g \leq \mu^0$, then $\Pi \leq 0$ with:

$$\Pi = -p\lambda T \cdot \frac{\mu^h - \mu^0}{\mu^h - \mu^\ell}. \tag{15}$$

(iii) Otherwise, $\Pi = 0$.

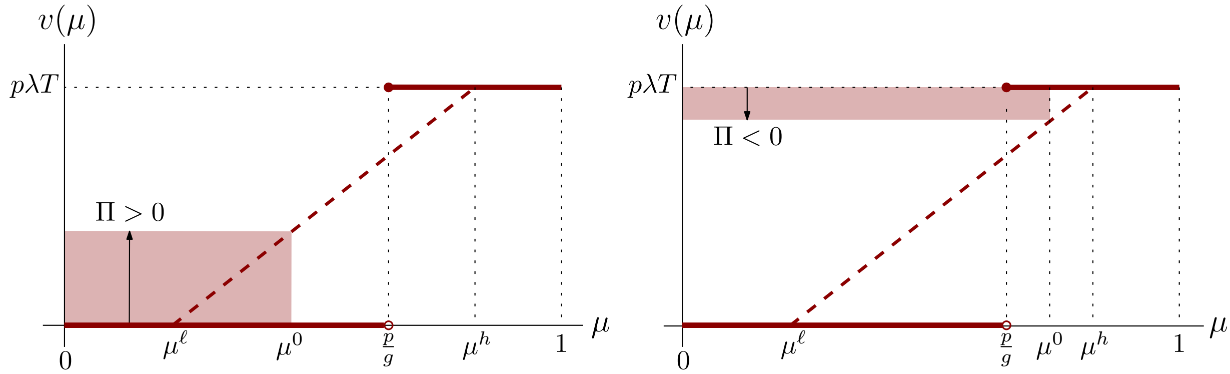

(a) Positive value of influencer when prior belief of high state μ^0 is low

(b) Negative value of influencer when prior belief of high state μ^0 is high

Figure 2: An illustration of Proposition 1

This proposition characterizes when an influencer creates value in a setting without social learning. The key requirement is that $\mu^\ell \leq p/g \leq \mu^h$, meaning that the influencer is sufficiently informative so that her signal changes purchase decisions. Followers will not purchase after a low signal but will purchase after a high signal. If this condition fails, the influencer creates no value because followers' decisions remain unchanged regardless of which signal she sends.

Given this requirement, the influencer's value depends on where the prior belief μ^0 sits relative to the purchase threshold p/g. Figure 2 illustrates cases (i) and (ii) geometrically, with the vertical height of the red rectangle representing the influencer's value Π.

Case (i) describes the scenario where influencers are valuable. When customers initially hold a low belief about the product ($\mu^0 \leq p/g$), they will not purchase without additional information. An influencer who can send a positive signal raises the belief above the threshold, leading to purchases that would not have occurred otherwise. Formula (14) shows that the influencer's value increases linearly with how far the prior μ^0 is from the lower post-signal belief μ^ℓ, normalized by the spread $\mu^h - \mu^\ell$.

Case (ii) reveals a cautionary finding: influencers can actually harm revenue when the prior belief is already high ($\mu^0 \geq p/g$). In this case, customers would purchase based on their prior belief alone. However, the influencer introduces risk. A negative signal could lower beliefs below the purchase threshold, converting buyers into non-buyers. The positive signal provides no additional benefit since customers were already going to purchase. This asymmetry makes the expected value negative.

These results are relatively standard in the applied information design literature (see, e.g., Boyaci et al. 2022; Alizamir et al. 2020). However, a less common area of investigation is how the influencer's value depends on her informativeness π. Our framework enables us to characterize this relationship in Proposition 2 below. The proof is in Appendix D.2. Since there are only two signals, the pair $(\pi_{hH}, \pi_{\ell L})$ fully characterizes the two-by-two informativeness matrix π.

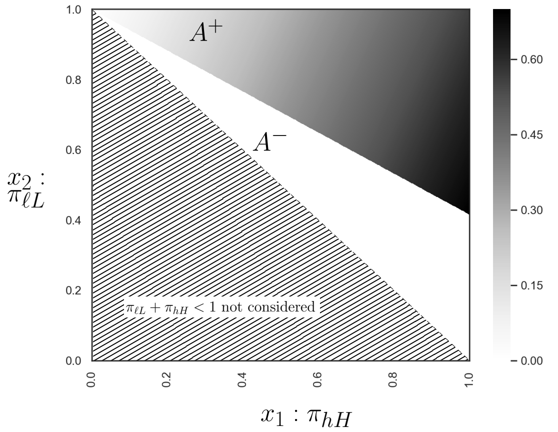

Figure 3: An illustration of Proposition 2. When $\mu^0 \leq p/g$ the region A^+ corresponds to the region where the value of the influencer is positive and the region A^- is the region where the influencer's value is zero. The heatmap shading shows the influencer's value, with darker colors indicating higher value.

Proposition 2 (Impact of π on influencer value)**.** Let $m \triangleq \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}$ and define the regions

$$A^+ \triangleq \left\{(x_1, x_2) \in [0,1]^2 \cap \{x_2 \leq 1 - x_1\} : x_2 \geq 1 - m x_1\right\}$$
$$A^- \triangleq \left\{(x_1, x_2) \in [0,1]^2 \cap \{x_2 \leq 1 - x_1\} : x_2 < 1 - m x_1\right\}.$$

The influencer's value Π is *positive* if $\mu^0 < p/g$ and $(\pi_{hH}, \pi_{\ell L}) \in A^+$, and nonpositive otherwise.

Figure 3 uses heatmap shading to illustrate the influencer's value as a function of informativeness, with regions A^+ and A^- marking the boundary between positive and zero value. Since customers do not purchase under the prior ($\mu^0 < p/g$), an information structure is valuable if and only if a positive signal can raise the belief μ^h above the purchase threshold p/g. When informativeness $(\pi_{hH}, \pi_{\ell L})$ lies in region A^+, this condition is satisfied and the resulting post-signal belief yields positive influencer value as in case (i) of Proposition 1. Otherwise, the post-signal belief μ^h remains below p/g, so the influencer has zero value as in case (iii).

We first discuss the impact of μ^0. Note that $m = \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}$ is the slope of the line separating region A^+ from A^-. As μ^0 increases toward p/g, customers are closer to the purchase threshold, making it easier for a positive signal to trigger purchases. This expands the set of information structures (region A^+) that can shift customers above this threshold.

Beyond the role of the prior, the figure also reveals how different types of errors affect value. A key insight is that false positive errors ($\pi_{hL} = 1-\pi_{\ell L}$) and false negative errors ($\pi_{\ell H} = 1-\pi_{hH}$) have asymmetric impacts. As the false negative rate increases, fewer information structures become valuable. However, for any false negative rate, there always exists an information structure with sufficiently high true negative rate that creates positive value. In contrast, when the false positive rate is too high, no information structure creates value, even one that perfectly detects good products ($\pi_{hH} = 1$).

The implication is that high false positive rate is more detrimental than high false negative rate. When an influencer frequently endorses bad products (high false positive rate), her positive signals lose credibility. Even if she perfectly detects good products, followers may treat the influencer's positive signal to be false, limiting its ability to raise beliefs above the purchase threshold. Such an influencer brings no value to the firm. Conversely, an influencer who frequently fails to endorse good products (high false negative rate) can still create value if she rarely endorses bad products. When false positives are uncommon, her endorsements remain highly credible and can push followers' beliefs above the purchase threshold, triggering purchases.

Finally, the heatmap reveals patterns in how value varies with informativeness. For fixed true negative rate $\pi_{\ell L}$, the influencer's value increases monotonically with the true positive rate π_{hH}. However, for fixed true positive rate π_{hH}, the maximum value occurs at an interior value of $\pi_{\ell L} \in (0, 1)$ rather than at the extreme $\pi_{\ell L} = 1$. This implies the ideal influencer sends a somewhat noisy signal in the bad state. The intuition is straightforward: an influencer with perfectly accurate signal generates zero revenue when the state is bad, since she never endorses bad products and thus never triggers purchases in that state. By introducing noise, the influencer occasionally endorses bad products, generating positive expected revenue even in the bad state. This benefit from occasional false positives must be balanced against the credibility loss discussed earlier, yielding an interior optimum.

4.3 A model with social learning

We now introduce social learning into the model by allowing a fraction of customers to be skeptics ($\alpha < 1$). Recall that devotees rely exclusively on the influencer's signal when making purchase decisions, while skeptics also incorporate information generated organically by other customers, such as observed purchase behaviors and posted reviews. This organic information accumulates over time as customers arrive and make purchase decisions, creating a dynamic feedback loop that fundamentally changes the structure of demand.

We model the additional organic information-generating processes as follows. With probability γ, each purchase generates a public signal $r \in \{h, \ell\}$, where h denotes a positive signal and ℓ denotes a negative signal. These signals represent organic feedback such as online reviews, ratings, or social media posts about the product. This signal r is drawn from a conditional probability $\pi^R(r|w)$ distribution, where $w \in \{H, L\}$ is the true state. For convenience, we write:

$$\pi^R \triangleq \begin{pmatrix} \pi^R_{hH} & \pi^R_{\ell H} \\ \pi^R_{hL} & \pi^R_{\ell L} \end{pmatrix} \triangleq \begin{pmatrix} \pi^R(h|H) & \pi^R(\ell|H) \\ \pi^R(h|L) & \pi^R(\ell|L) \end{pmatrix}$$

The distribution π^R represents the average accuracy of the organic signals, which can be interpreted as the historical reliability of the online community. Note the distinction between π^R (accuracy of organic signals) and π (informativeness of the influencer). Our approach of treating π^R as exogenously given by historical data is standard in the information economics literature (see, e.g., Banerjee 1992; Bikhchandani et al. 1992). We provide a micro-foundation of π^R in Appendix I.

Remark 2 (Justification of assumptions)**.** Our model makes three key assumptions about the organic signal generation process: (i) only buyers generate signals, (ii) the propensity to generate signals (γ) is independent of the true state, and (iii) the customer arrival rate (λ) is independent of the true state. These choices reflect institutional characteristics of most e-commerce platforms and standard modeling practices. Requiring verified purchases limits reviews from non-buyers, which supports (i). Separating the rate of signal generation from signal informativeness aligns with the information economics literature. This makes the problem tractable and allows us to emphasize how quality impacts review content rather than volume, which supports (ii). In many e-commerce sites, customer arrival is driven by platform algorithms or marketing, which are often independent of underlying quality, which supports (iii). In Appendix I, we examine the robustness of our analysis to alternative assumptions.

Demand process. By the decision rule in (1), a customer purchases if and only if her belief exceeds p/g. Let μ_t denote the belief of a skeptic at time t about the probability that the state is H. At time $t = 0$, all customers share the same initial belief: $\mu_0 = \mu^0$ if the influencer is not sponsored, or $\mu_0 = \mu^s$ if the influencer sends signal s. Devotees maintain this initial belief μ_0 throughout the selling season, yielding a constant purchase rate of $\lambda\alpha\mathbf{1}_{\mu_0 > p/g}$. In contrast, skeptics update their beliefs to μ_t as they observe organic signals over time, yielding a time-varying purchase rate of $\lambda(1-\alpha)\mathbf{1}_{\mu_t \geq p/g}$. Note that we assume skeptics do not purchase when exactly indifferent ($\mu_t = p/g$), while devotees do. This technical assumption ensures well-defined boundary conditions, but has negligible impact on the analysis.

The overall demand rate at time t is therefore $\lambda \cdot f(\alpha, \mu_0, \mu_t)$, where

$$f(\alpha, \mu_0, \mu_t) \triangleq \alpha\mathbf{1}_{\mu_0 \geq p/g} + (1-\alpha)\mathbf{1}_{\mu_t > p/g}. \tag{16}$$

Let D_t denote the cumulative demand up to time t, and let $\mathcal{F}_t$ denote the information available at time t. Since demand follows a nonhomogeneous Poisson process with time-varying rate $\lambda f(\alpha, \mu_0, \mu_t)$, conditional on $\mathcal{F}_t$, the transition probabilities for D_{t+h} over a time interval of length h are:

$$\mathbb{P}\left(D_{t+h} = D_t + k \mid \mathcal{F}_t\right) = \begin{cases} 1 - \lambda f(\alpha, \mu_0, \mu_t)h + o(h), & \text{if } k = 0, \\ \lambda f(\alpha, \mu_0, \mu_t)h + o(h), & \text{if } k = 1, \\ o(h), & \text{if } k \geq 2, \end{cases} \tag{17}$$

where $o(h)$ denotes a term where $\lim_{h\to 0} \frac{o(h)}{h} = 0$.

Belief dynamics. We now formalize how skeptics' beliefs μ_t evolve over time. We denote by $\mathcal{H}_t = \{R_{hu}, R_{\ell u}, 0 \leq u \leq t\}$ the information set available to skeptics at time t, where R_{hu} is the cumulative number of h (positive) signals by time u and $R_{\ell u}$ is the cumulative number of ℓ (negative) signals by time u. We let $\mathcal{F}_t \triangleq \sigma(R_{hu}, R_{\ell u}, 0 \leq u \leq t)$ denote the filtration generated by these signal processes.

Importantly, the purchase history $\{D_u, 0 \leq u \leq t\}$ does not provide additional information to skeptics. Because the arrival rate λ is not state-dependent, the rate of purchases reveals nothing about the state. Therefore, skeptics update their beliefs solely based on the signal history $\mathcal{H}_t$. When only purchase counts (without signals) are observed, the model reduces to the baseline setting in Section 4.2 where charisma α plays no role (see Appendix C).

The signal process $\{R_{ht} + R_{\ell t}, t \geq 0\}$ is a nonhomogeneous Poisson process with rate $\lambda\gamma f(\alpha, \mu_0, \mu_t)$ at time t, where γ is the probability that a buyer sends an organic signal. The likelihood of the signal history $\{R_{hu}, R_{\ell u}, u \leq t\}$ corresponds to the likelihood of inter-arrival times $\{t_1, t_2, \ldots, t_{R_{ht}+R_{\ell t}}\}$ in this Poisson process, where t_k is the time that the k^{th} signal occurs. Given state $w \in \{H, L\}$, the likelihood of the signal history at time t is:

$$\begin{aligned}
&\mathbb{P}\left(\{R_{hu}, R_{\ell u}, u \leq t\} \mid w\right) \\
&\quad = e^{\left(-\int_0^t \lambda\gamma f(\alpha,\mu_0,\mu_u)\mathrm{d}u\right)} \prod_{k=1}^{R_{ht}+R_{\ell t}} \left(\lambda\gamma f(\alpha, \mu_0, \mu_{t_k})\right) (\pi^R_{hw})^{R_{ht}} (\pi^R_{\ell w})^{R_{\ell t}}
\end{aligned} \tag{18}$$

This is the standard likelihood function for inter-arrival times in a nonhomogeneous Poisson process (Zhao and Xie, 1996). Given the prior belief μ_0 and the likelihood function (18), we derive the posterior belief μ_t using Bayes' rule. The proof is in Appendix D.3.

Lemma 2. A skeptic with prior belief $\mu \in \{\mu^0, \mu^s\}$ who observes signals $\mathcal{H}_t = \{R_{hu}, R_{\ell u}, u \leq t\}$ would have a posterior belief at time t equal to:

$$\mu_t = \frac{\mu}{\mu + (1-\mu)\left(\frac{\pi^R_{hL}}{\pi^R_{hH}}\right)^{R_{ht}} \left(\frac{\pi^R_{\ell L}}{\pi^R_{\ell H}}\right)^{R_{\ell t}}}. \tag{19}$$

The expression in (19) in intuitive. If no signal arrives, μ_t does not change because the signal generation rate is not state-dependent. By contrast, when a positive signal arrives and $\pi^R_{hL}/\pi^R_{hH} < 1$, skeptics strengthen their belief that the product is good. The ratio $\pi^R_{hL}/\pi^R_{hH} < 1$ means positive signals are more likely under the high state, so positive signals naturally increase the belief in product quality. In simple terms, skeptics become more confident about product quality when they see positive reviews, but only if positive reviews are more common for good products than bad ones. The updating is faster when skeptics are initially uncertain and slower when they already have strong convictions (μ close to 0 or 1).

Since μ_t depends on the discrete processes R_{ht} and $R_{\ell t}$ through the nonlinear function (19), a differential form for the dynamics of μ_t is not immediately obvious. However, we can apply Itô's lemma to derive a stochastic differential equation. Itô's lemma is a chain rule defined on stochastic processes. While many readers may be familiar with Itô's lemma applied to functions of Brownian motion, here we use a version that applies to jump processes. The proof is in Appendix D.4.

Lemma 3. The posterior belief μ_t satisfies the following stochastic differential equation:

$$\mathrm{d}\mu_t = \frac{\mu_{t-}(1-\mu_{t-})\left(\pi^R_{hH} - \pi^R_{hL}\right)}{\pi^R_{hH}\mu_{t-} + \pi^R_{hL}(1-\mu_{t-})}\mathrm{d}R_{ht} + \frac{\mu_{t-}(1-\mu_{t-})\left(\pi^R_{\ell H} - \pi^R_{\ell L}\right)}{\pi^R_{\ell H}\mu_{t-} + \pi^R_{\ell L}(1-\mu_{t-})}\mathrm{d}R_{\ell t}. \tag{20}$$

where μ_0 is the prior belief, and μ_{t-} denotes the left limit $\mu_{t-} = \lim_{h\to 0} \mu_{t-h}$.

These results characterize how customer beliefs evolve. Devotees maintain static beliefs, while skeptics' beliefs follow a stochastic jump process. The posterior belief μ_t jumps when signals arrive, with jump sizes that are small when skeptics already hold strong beliefs (either μ_{t-} close to 1 or close to 0). The jump magnitude reflects the accuracy of the online community:

for positive signals, it depends on π^R_{hH} and π^R_{hL}; for negative signals, on $\pi^R_{\ell H}$ and $\pi^R_{\ell L}$.

Expected revenue. Similar to Section 4.2, the key to computing the influencer's value is evaluating the expected revenue function $v(\mu)$ for the common initial prior μ where $\mu \in \{\mu^0, \mu^h, \mu^\ell\}$. When all followers are devotees, $v(\cdot)$ is a simple step function. However, when skeptics are present, computing $v(\cdot)$ is more challenging because of its dependence on skeptics' belief μ_t that evolves dynamically. We can see this from:

$$v(\mu) = \mathbb{E}\left[\int_0^T p\mathrm{d}D_t \;\middle|\; \mathcal{F}_0\right] = \mathbb{E}\left[\int_0^T p\mathbb{E}\left(\mathrm{d}D_t \mid \mathcal{F}_t\right) \;\middle|\; \mathcal{F}_0\right] = \mathbb{E}\left[\int_0^T p\lambda f(\alpha,\mu,\mu_t)\,\mathrm{d}t \;\middle|\; \mathcal{F}_0\right] \quad (21)$$

where the last equality uses the evolution of D_t from (17). In this expression, we assume that the firm's belief is the same as that of the skeptics since both have access to the same information.

Our approach is to utilize the dynamics of μ_t to evaluate the expected revenue using dynamic programming (DP). We define a revenue-to-go function with two state variables: (i) the posterior belief of skeptics $y \in [0,1]$, and (ii) the remaining time $\bar{t} \in [0,T]$ until the end of the selling season. (The reader should not confuse the DP state variables y and $\bar{t}$ with the state of nature w.) These state variables are sufficient based on the belief dynamics in Lemma 3. Let $V(y,\bar{t};z)$ denote the expected revenue-to-go when skeptics hold posterior belief y, the remaining time is $\bar{t}$, and devotees hold a fixed belief z. Note that the total expected revenue is then $v(\mu) = V(\mu,T;\mu)$, since skeptics and devotees start with the same belief at time $t=0$.

We derive the DP formulation of $V(y,\bar{t};z)$ by considering a outcomes during a small time interval of length h. From the transition probabilities in (17), a purchase will occur with probability $\lambda f(\alpha,z,y)h + o(h)$, earning a revenue p. If a purchase occurs, the buyer produces a signal with probability γ. Conditional on the current belief y, this signal is positive with probability $y\pi^R_{hH} + (1-y)\pi^R_{hL}$ and negative with probability $y\pi^R_{\ell H} + (1-y)\pi^R_{\ell L}$. The belief updates to $y + J_{hy}$ after a positive signal, to $y + J_{\ell y}$ after a negative signal, and remains y if no signal arrives, where J_{hy} and $J_{\ell y}$ are the jump sizes given a positive and negative signal, respectively. From Lemma 3:

$$J_{hy} \triangleq \frac{y(1-y)\left(\pi^R_{hH} - \pi^R_{hL}\right)}{\pi^R_{hH}y + \pi^R_{hL}(1-y)}, \quad \text{and} \quad J_{\ell y} \triangleq \frac{y(1-y)\left(\pi^R_{\ell H} - \pi^R_{\ell L}\right)}{\pi^R_{\ell H}y + \pi^R_{\ell L}(1-y)}. \quad (22)$$

Therefore, we can write $V(y,\bar{t};z)$ as:

$$\begin{aligned} V(y,\bar{t};z) = \lambda f(\alpha,z,y)h \cdot \Big[&p + \gamma\left(\pi^R_{hH}y + \pi^R_{hL}(1-y)\right)V(y+J_{hy},\bar{t}-h;z) \\ &+ \gamma\left(\pi^R_{\ell H}y + \pi^R_{\ell L}(1-y)\right)V(y+J_{\ell y},\bar{t}-h;z) + (1-\gamma)V(y,\bar{t}-h;z)\Big] \\ &+ (1-\lambda f(\alpha,z,y)h)\cdot V(y,\bar{t}-h;z) + o(h), \end{aligned} \quad (23)$$

Subtracting both sides by $V(y,\bar{t}-h;z)$, dividing by h, and taking the limit as $h \to 0$ yields the Bellman equation.

Lemma 4. The expected revenue-to-go function V satisfies the Bellman equation:

$$\begin{aligned} \frac{\partial V(y,\bar{t};z)}{\partial \bar{t}} =& \lambda f(\alpha,z,y)p + \Big[\left(\pi^R_{hH}y + \pi^R_{hL}(1-y)\right)V(y+J_{hy},\bar{t};z) \\ &+ \left(\pi^R_{\ell H}y + \pi^R_{\ell L}(1-y)\right)V(y+J_{\ell y},\bar{t};z) - V(y,\bar{t};z)\Big]\cdot\lambda\cdot f(\alpha,z,y)\gamma, \end{aligned} \quad (24)$$

where $J_{hy}, J_{\ell y}$ are as defined in (22). The boundary conditions are

$$V(y, 0; z) = 0, \quad \forall y \in [0, 1];$$
$$V(1, \bar{t}; z) = \lambda f(\alpha, z, 1)p\bar{t}, \quad \forall \bar{t} \in [0, T]$$
$$V(0, \bar{t}; z) = \lambda f(\alpha, z, 0)p\bar{t}, \quad \forall \bar{t} \in [0, T].$$

The Bellman equation (24) captures a fundamental trade-off: the firm earns immediate revenue from current purchases, but must also account for how these purchases generate information (through organic signals) that affects future sales. The first term represents the instantaneous revenue rate, while the second term captures the expected change in value due to information that updates skeptics' beliefs.

Asymptotic analysis. The Bellman equation in Lemma 4 is analytically intractable because the belief y jumps up by J_{hy} or down by $J_{\ell y}$, which are complicated nonlinear functions of y, as defined in (22). Thus, we instead characterize value function V through asymptotic analysis. Ifrach et al. (2019) similarly use asymptotic analysis to make the pricing problem under social learning tractable.

A natural first approach might be to use fluid approximation, which takes a "law of large numbers" limit where randomness is averaged out to reveal the underlying deterministic trends. This approach is widely used in analyzing state-dependent queueing networks (Bäuerle, 2002; Zychlinski, 2023). However, fluid approximation is not suitable for our model because the belief process $\{\mu_t\}$ is a martingale, a stochastic process with no systematic drift. From Lemma 3 we have $\mathbb{E}[\mathrm{d}\mu_t|\mathcal{F}_t] = 0$, meaning that in expectation, beliefs do not change. Since there is no deterministic trend, removing randomness would cause the belief process to remain constant at its initial value μ_0. This would eliminate the dynamic updating that defines social learning, reducing our model to the static baseline case without skeptics (see Appendix H).

We instead use the *diffusion limit* which approximates a discrete jump process with a continuous process (Glynn, 1990). Think of the belief μ_t as a particle that jumps randomly: it jumps up by $J_{h\mu_t}$ when a positive signal arrives, and down by $J_{\ell\mu_t}$ when a negative signal arrives. These jumps occur at random times following a Poisson process. The diffusion approximation is valid when jumps are frequent but individually small—that is, when many signals arrive over time, but each individual signal does not completely flip beliefs. In this regime, the accumulated effect of many small random jumps can be well-approximated by continuous random fluctuations, similar to how the discrete path of many coin flips converges to a smooth random walk in the limit. We validate through Monte Carlo simulations (Appendix K) that this approximation is highly accurate, particularly when the customer arrival rate λ is large.

Mathematically, this approximation is justified by the functional martingale central limit theorem, which states that appropriately scaled martingales converge to Brownian motion. The formal result is as follows (proof in Appendix D.5).

Lemma 5. Let $\mu_t^{(n)}$ denote the belief process when the arrival rate is λn with arrival jump size $1/\sqrt{n}$, and let $\{\langle\mu^{(n)}\rangle_t, t \geq 0\}$ denote the quadratic variation[4]. For $\mu_0 \geq p/g$, the following holds:

$$\frac{\mu_t^{(n)} - \mu_0}{\sqrt{\langle\mu^{(n)}\rangle_t}} \to^{d.} \mathcal{N}(0,1) \text{ as } n \to \infty$$

where $\to^{d.}$ denotes convergence in distribution and $\mathcal{N}(0,1)$ denotes standard Normal distribution.

This result says that the belief change, when properly normalized by its variability, converges to a normal distribution. This convergence justifies replacing the discrete jump dynamics with continuous Brownian motion.

Using Lemma 5, we can replace the Bellman equation involving discrete jumps with a partial differential equation involving a continuous diffusion term. The result is as follows whose proof is in Appendix D.6.

Proposition 3. Let $\tilde{V}$ be the diffusion approximation of the value function V. If $z \geq p/g$, then $\tilde{V}$ satisfies the partial differential equation:

$$\frac{\partial \tilde{V}(y,\bar{t};z)}{\partial \bar{t}} = \lambda f(\alpha,z,y)p + \frac{1}{2}\frac{\lambda f(\alpha,z,y)\gamma y^2(1-y)^2\left(\pi_{hH}^R - \pi_{hL}^R\right)^2}{\left(\pi_{hH}^R y + \pi_{hL}^R(1-y)\right)\left(\pi_{\ell H}^R y + \pi_{\ell L}^R(1-y)\right)}\frac{\partial^2 \tilde{V}(y,\bar{t};z)}{\partial y^2}, \tag{25}$$

with boundary conditions:

$$\begin{aligned}\tilde{V}(y,0;z) &= 0, \quad \forall y \in [0,1],\\ \tilde{V}(1,\bar{t};z) &= \lambda p\bar{t}, \quad \forall \bar{t} \in [0,T],\\ \tilde{V}(p/g,\bar{t};z) &= \alpha\lambda p\bar{t}, \quad \forall \bar{t} \in [0,T].\end{aligned}$$

The PDE in (25) is a forward parabolic equation on the domain $y \in [p/g, 1]$, $\bar{t} \in [0,T]$. The diffusion coefficient of $\frac{\partial^2 \tilde{V}}{\partial y^2}$ is positive and continuously differentiable for $y \in (p/g, 1)$ given our assumption that $\pi_{hH}^R > \pi_{hL}^R$ (organic signals are informative). The PDE is equipped with an initial condition (for $\bar{t} = 0$) and Dirichlet boundary conditions (for $y = 1$ and $y = p/g$). Standard PDE theory (see, e.g., Friedman 1964) guarantees existence and uniqueness of a classical solution under these conditions.

The first term in the PDE (25) represents the instantaneous revenue rate from current purchases. The second term captures how uncertainty in beliefs evolves over time. The coefficient of $\frac{\partial^2 \tilde{V}}{\partial y^2}$ can be interpreted as the volatility of the belief process: when organic signals are highly informative (large $|\pi_{hH}^R - \pi_{hL}^R|$), beliefs change more rapidly in response to signals, increasing this volatility. This second-order term, which is not present in the baseline model without social learning, captures the option value of information: beliefs can move in either direction, and this uncertainty affects the firm's expected revenue.

We approximate the expected revenue function by $\tilde{v}(\mu) \triangleq \tilde{V}(\mu, T; \mu)$. The tractability gained from the diffusion limit allows us to characterize the structure of $\tilde{v}(\cdot)$ analytically, which would be impossible with the original jump process formulation. The proof of the following result is in Appendix D.7.

[4]Quadratic variation of a martingale is defined as $\lim_{\Delta t\to 0}\sum_{i=1}^n (X_{t_i} - X_{t_{i-1}})^2$ where $0 = t_0 < t_1 < \ldots < t_n = t$ and $\Delta t \triangleq \max_i(t_i - t_{i-1})$.

Lemma 6. When $\pi^R_{hH} \neq \pi^R_{hL}$, the following results hold:

(i) For $\mu \in [0, p/g)$, we have $\tilde{v}(\mu) = 0$.

(ii) For $\mu \in [p/g, 1]$, the function $\tilde{v}(\mu)$ is concave and nondecreasing in μ.

Lemma 6 reveals a key structural property: with social learning, the value function is increasing for $\mu \in [p/g, 1]$. This contrasts with the baseline model without social learning (studied in Section 4.2), where the expected revenue is constant for $\mu \in [p/g, 1]$. The concavity reflects diminishing marginal returns: when the prior is already high, skeptics start with strong convictions, so organic signals produce smaller belief updates, limiting the value of social learning. When the prior is moderate (just above p/g), organic signals produce larger belief changes, creating more potential for positive value of information.

Figure 4 plots both $\tilde{v}(\mu)$ (the expected revenue with social learning) and $v^o(\mu)$ (the expected revenue without social learning). Both are zero for $\mu < p/g$ with a discontinuity at $\mu = p/g$. In $\mu \in [p/g, 1]$, $v^o(\mu)$ is a constant while $\tilde{v}(\mu)$ is increasing and concave. This visual comparison highlights how social learning creates sensitivity to the initial belief: unlike the baseline case where only crossing the threshold matters, with social learning the magnitude of the initial belief affects revenue through its influence on the subsequent information learning. In the next section, we explore these implications in detail, characterizing how social learning affects the value an influencer creates for the firm.

5 Impact of social learning on influencer value

The structure of $\tilde{v}(\mu)$ revealed in Lemma 6 has important implications for when influencers create value. To understand these implications, we must examine where the purchase threshold p/g falls relative to the three belief levels: μ^ℓ (belief after a low influencer signal), μ^0 (belief without influencer), and μ^h (belief after a high influencer signal). Since $\mu^\ell \leq \mu^0 \leq \mu^h$, there are four distinct cases depending on where p/g falls among these beliefs. Figure 4 illustrates these four cases, extending the comparison between $\tilde{v}(\mu)$ and $v^o(\mu)$ to show how the influencer's value varies across these scenarios.

In each panel of Figure 4, the black curve shows expected revenue with social learning $\tilde{v}$, while the red curve shows expected revenue without social learning v^o. The shaded areas represent influencer value: gray for Π_{SL} (with social learning) and red for Π_{NSL} (without social learning). These values equal the vertical heights of the respective shaded regions.

Cases (a) and (d): Influencer provides no value. When all belief levels fall on the same side of the threshold p/g, the influencer is not informative since her signal alone cannot change purchase decisions. In these cases, the influencer cannot create value. In panel (a), all beliefs exceed p/g. Without social learning, $\Pi_{\mathsf{NSL}} = 0$. With social learning, $\Pi_{\mathsf{SL}} < 0$: the influencer actually has a negative value, since a low signal weakens beliefs and makes consumers more susceptible to the subsequent signals. In panel (d), all beliefs are below p/g, so no one purchases even after a high signal. Since there are no purchases, there are no organic signals. Both with and without social learning, $\Pi_{\mathsf{SL}} = \Pi_{\mathsf{SL}} = 0$.

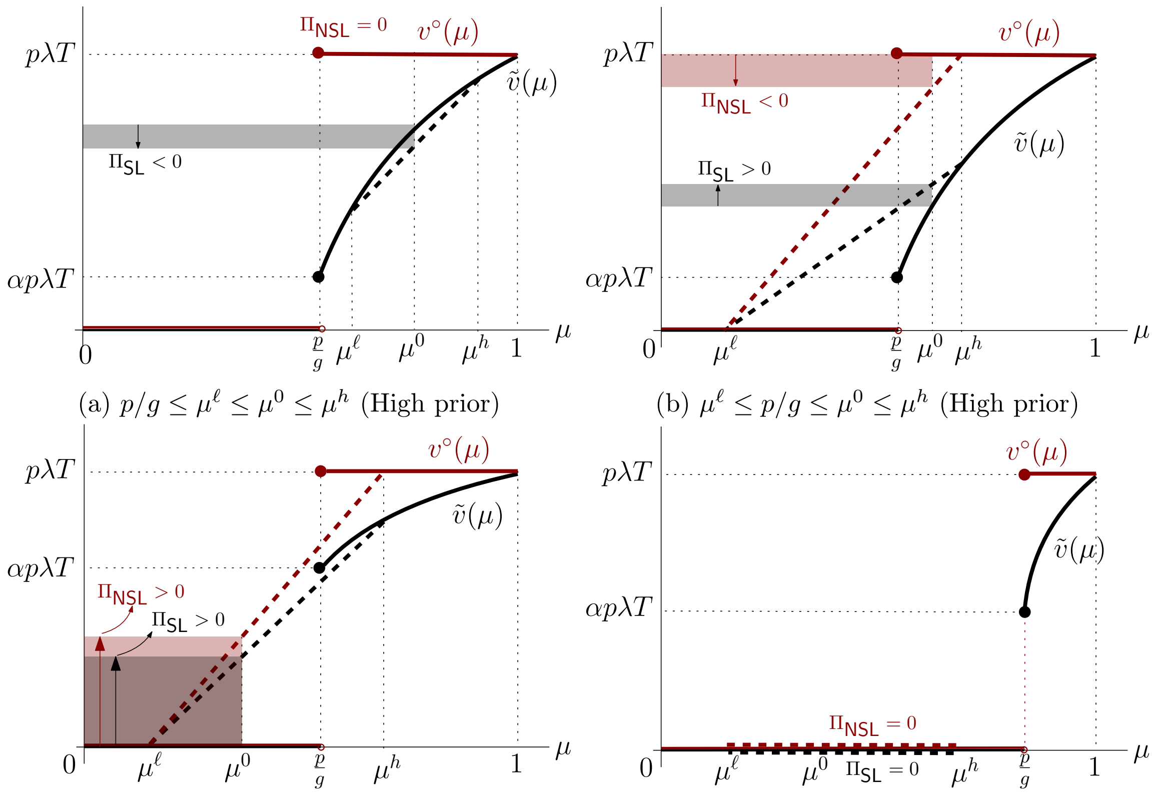

Figure 4: Four cases, (a) through (d), of the possible arrangement of μ^ℓ, μ^0, and μ^h with respect to p/g, and the impact on the value of the influencer with and without social learning.

Case (c): Social learning *discounts* value. When $\mu^0 < p/g < \mu^h$, customer purchases are initiated only after a positive signal from the influencer. Panel (c) shows that $\Pi_{\text{NSL}} > \Pi_{\text{SL}} > 0$. Without social learning, all followers who see the high signal purchase throughout the season. With social learning, skeptics may stop purchasing if subsequent organic signals are negative, reducing the influencer's value. This aligns with intuition: social learning dilutes the influencer's impact by introducing competing information sources.

Case (b): Social learning *boosts* value. When $\mu^\ell \leq p/g \leq \mu^0$, panel (b) shows a counter-intuitive result. Here, $\Pi_{\text{NSL}} < 0$ but $\Pi_{\text{SL}} > 0$. Without social learning, the influencer adds no value because customers already purchase at prior μ^0, so a high signal provides no benefit while a low signal stops all purchases. With social learning, however, a high signal raises skeptics' starting belief from μ^0 to μ^h. This higher starting point means beliefs are less likely to fall below p/g during the diffusion process (even after negative organic signals), generating more purchases over time. The influencer becomes valuable precisely because of social learning that introduces curvature of the value function due to the variance term in (25).

These cases reveal a nuanced picture: social learning's impact on influencer value depends critically on whether customers would purchase without the influencer's signal. When customers need persuasion to purchase (case c, where $\mu^0 < p/g$), social learning *reduces* influencer value

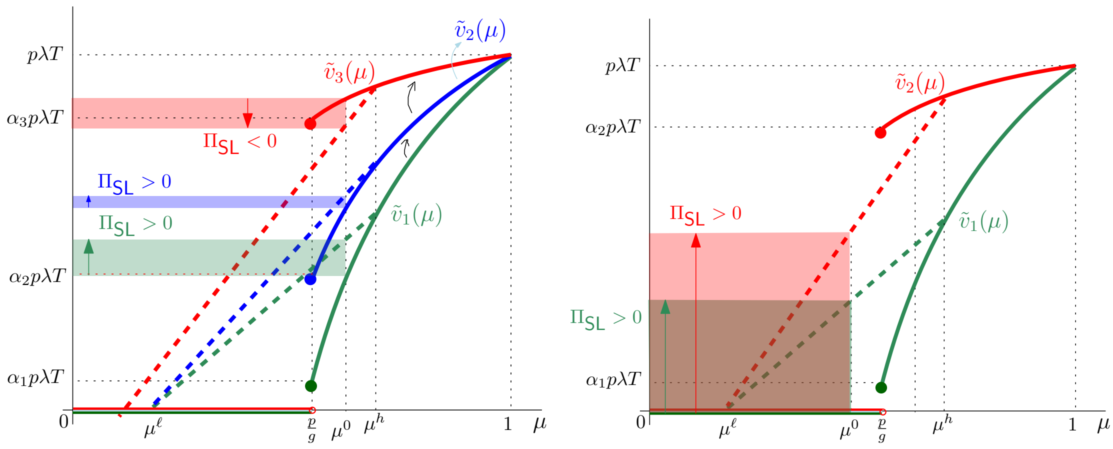

(i) Π_{SL} decreases in α in case (b) of Figure 4. (ii) Π_{SL} increases in α in case (c) of Figure 4

Figure 5: An illustration of how influencer value can either increase or decrease as charisma increases depending on the relative position of μ^0 with respect to p/g. Three curves corresponding to increasing levels of charisma: $\alpha_1 < \alpha_2 < \alpha_3$

by introducing competing organic signals that can override the influencer's message. Conversely, when customers already intend to purchase (case b, where $\mu^0 > p/g$), social learning *increases* influencer value since an endorsement strengthens skeptics' starting belief, protecting the belief from drifting downward when organic signals are negative.

Figure 5 further reveals how the proportion of devotees (α) interacts with these effects. The figure shows three curves corresponding to increasing levels of charisma: $\alpha_1 < \alpha_2 < \alpha_3$. In case (c), panel (ii), influencer value increases with fewer skeptics that rely on competing organic signals. In case (b), panel (i), influencer value decreases with fewer skeptics whose belief is strengthened by a positive signal.

Figure 6 summarizes these findings by mapping when influencers are valuable as a function of the decision threshold (horizontal axis) and charisma (vertical axis). Note that the informativeness π changes the position of μ^ℓ and μ^h on the horizontal axis via (4). Several insights emerge: (i) When priors are low ($\mu^0 < p/g$), influencers are always (weakly) valuable regardless of charisma. (ii) When priors are high ($p/g < \mu^0$) and a "no endorsement" signal is credible ($\pi_{\ell L}$ is sufficiently high so $\mu^\ell < p/g$), the only influencers that create value are those with a sufficient proportion of skeptical followers. (iii) As the no-endorsement signal becomes more credible (as $\pi_{\ell L}$ increases and μ^ℓ decreases), the influencer with social learning becomes more valuable.[5] In the extreme case, an influencer with perfectly credible negative signals (inducing $\mu^\ell = 0$) always has non-negative value as long there are enough skeptics.

Our main analytical result formalizes these observations and its proof is in Appendix E.1.

Theorem 1. For a given prior μ^0 and an influencer with a given information structure π, if an influencer is valuable (i.e., $\Pi_{\mathsf{NSL}} \geq 0$) without social learning, then that same influencer is also valuable when there is social learning (i.e., $\Pi_{\mathsf{SL}} \geq 0$). However, the converse is not always true.

[5]This insight is further illustrated in Figure 9 in Appendix A.

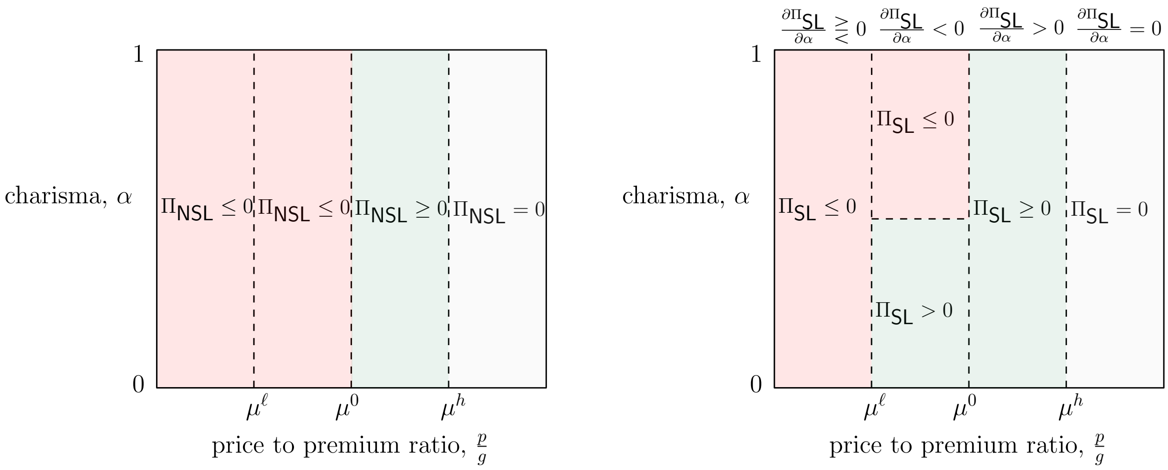

(i) No social learning (ii) Social learning

Figure 6: A visual summary of observations and intuitions arising in Figures 4 and 5.

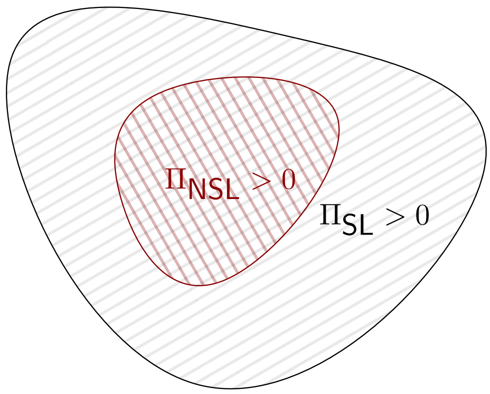

Figure 7: Parameter values $(\alpha, \mu^0, p/g)$ that make the π-influencer valuable

Moreover,

(i) if $\Pi_{NSL} > 0$, then $\Pi_{NSL} \geq \Pi_{SL} > 0$;
(ii) if $\Pi_{NSL} = 0$, then $\Pi_{NSL} \geq \Pi_{SL}$.
(iii) if $\Pi_{NSL} < 0$, then $\Pi_{NSL} < \Pi_{SL}$;

Theorem 1 reveals a surprising asymmetry: the set of instances where influencers create value is strictly larger with social learning than without, as illustrated in Figure 7. One might expect social learning to always weaken influencer value by introducing competing information sources that dilute the influencer's message. Though this is true, those influencers still create positive (though weakened) value. The expansion of the set valuable influencers occurs because social learning can transform value-destroying influencers into value-creating ones.

6 Optimal information structure

We now consider a firm's decision when choosing among multiple available influencers. Which influencer should the firm select? This question is central to Bayesian persuasion (Kamenica and Gentzkow, 2011), where a sender (the firm) chooses how to reveal information to a receiver (customers) to maximize its objective. In our setting, the firm selects an influencer characterized by informativeness π, effectively choosing the information structure that followers observe.

We assume the firm has commitment power: it can commit to working with a particular influencer whose information structure π is publicly known. Following standard assumptions in Bayesian persuasion, we adopt a perfect Bayesian equilibrium (PBE) where both the firm and followers behave optimally. The firm's problem is to choose π from the feasible set $[0,1]^2 \cap \{\pi_{\ell L} + \pi_{hH} \leq 1\}$ to maximize its expected revenue.

Formally, the *optimal influencer problem* for given parameters μ^0, λ, α, and T is:

$$\begin{aligned} \underset{\pi}{\text{maximize}} \quad & \Pi(\pi, \mu^0) \\ \text{subject to} \quad & \mu^s = \frac{\mu^0 \pi_{sH}}{\mu^0 \pi_{sH} + (1-\mu^0)\pi_{sL}}, \quad s \in \{h, \ell\} \end{aligned} \tag{26}$$

where $\Pi(\pi, \mu^0)$ is a π-influencer's value defined in (8) and the constraints ensure consistency with Bayesian updating from (3).

6.1 An LP reformulation

Using the transformation from Lemma 1 and standard techniques in Bayesian persuasion (Kamenica and Gentzkow, 2011), we can reformulate (26) in terms of the distribution over posterior distributions τ defined in (9). This reformulation reveals the problem's structure more clearly and the proof is in Appendix F.1.

Proposition 4. The optimal influencer problem (26) is equivalent to:

$$\underset{\tau}{\text{maximize}} \quad \mathbb{E}_\tau [v(z)] - v(\mu^0) = \int_0^1 \tau(z) v(z) \mathrm{d}z - v(\mu^0) \tag{27a}$$

$$\text{subject to} \quad \mathbb{E}_\tau (z) = \int_0^1 \tau(z) z \mathrm{d}z = \mu^0 \tag{27b}$$

$$\int_0^1 \tau(z) \mathrm{d}z = 1 \tag{27c}$$

$$\tau(z) \geq 0 \text{ for all } z \in [0,1] \tag{27d}$$

where $\mathbb{E}_\tau$ is the expectation operator with respect to distribution τ of post-signal beliefs z for $z \in [0,1]$.

The reformulation (27) has an important structural property: constraint (27b) requires that the expected posterior equals the prior (Bayes plausibility), while constraints (27c)–(27d) ensure τ is a valid probability distribution. Since there is only one moment constraint, the dual Carathéodory theorem implies that the optimal τ has support on at most two points (see Barvinok 2002 and Lemma A.3). In other words, the optimal influencer sends signals that induce exactly two possible posterior beliefs.

This insight allows us to characterize the solution of (27) using the dual problem:

$$\underset{\lambda_1,\lambda_2}{\text{minimize}}\ \lambda_1 + \mu^0\lambda_2 - v(\mu^0) \tag{28a}$$

$$\text{subject to } \lambda_1 + z\lambda_2 \geq v(z) \text{ for all } z \in [0,1]. \tag{28b}$$

The dual problem has an intuitive interpretation: find the linear function $\lambda_1 + z\lambda_2$ that upper-bounds $v(z)$ everywhere and is tightest at the prior μ^0. Hence, $\lambda_1 + z\lambda_2$ can be thought of as a linear approximation of the expected revenue $v(z)$. The dual has only two variables, so at most two constraints are active at optimality. We denote these active constraints by their indices z_1 and z_2. By complementary slackness, $\tau(z) = 0$ for all $z \neq z_1, z_2$. Hence, the optimal influencer of the primal (27) has informativeness π^* with post-signal beliefs z_1 and z_2.

The following theorem characterizes the optimal influencer's value based on the solution to the dual problem. Its proof is in Appendix F.3.

Lemma 7 (Optimal influencer from z_1, z_2)**.** Suppose $v(\cdot)$ is differentiable almost everywhere. Let z_1 and z_2 (where $z_1 \leq z_2$) denote the indices of active constraints to the dual problem (28) under the optimal solution.

(i) If $z_1 < z_2$, the optimal influencer's *a priori* probabilities of post-signal beliefs are

$$\tau(z_1) = \frac{z_2 - \mu^0}{z_2 - z_1}, \quad \tau(z_2) = \frac{\mu^0 - z_1}{z_2 - z_1}, \tag{29}$$

and the optimal influencer's value is

$$\Pi^* = v(z_1) + (\mu^0 - z_1)\frac{v(z_2) - v(z_1)}{z_2 - z_1} - v(\mu^0). \tag{30}$$

(ii) If $z_1 = z_2$, then $\Pi^* = 0$ (no influencer provides value).

Corollary 1. Suppose $v(\cdot)$ is differentiable almost everywhere. Let z_1 and z_2 (where $z_1 \leq z_2$) denote the indices of active constraints to the dual problem (28) under the optimal solution. Sponsoring any influencer is beneficial to the firm if and only if $v(z_1) + (\mu^0 - z_1)\frac{v(z_2)-v(z_1)}{z_2-z_1} - v(\mu^0)$ is nonnegative.

In Figure 10 in Appendix A, we visualize the optimal influencer value in the case of $z_1 < z_2$. Geometrically, (30) represents the vertical distance between $v(\mu^0)$ and the chord connecting $(z_1, v(z_1))$ and $(z_2, v(z_2))$. Motivated by this, we provide a general Algorithm 1 for computing z_1 and z_2 (and hence Π^*) in Appendix J. Algorithm 1 implements a secant-slope search to identify the active constraints (z_1, z_2) in the dual formulation. Starting from the left, the algorithm repeatedly finds the point that maximizes the secant slope from the current position, effectively constructing the concave hull of $v(\cdot)$ segment by segment. The search terminates when it identifies the secant line segment whose interval $[z_1, z_2]$ contains the prior belief μ_0; these endpoints are the two active constraints that determine the optimal information structure. Algorithm 1 has complexity $\mathcal{O}(n^2)$ if the grid precision of $[0,1]$ is n.

We next apply Lemma 7 to characterize the optimal influencer to both our baseline model (Section 4.2) and social learning model (Section 4.3). To do so, we must verify that the respective

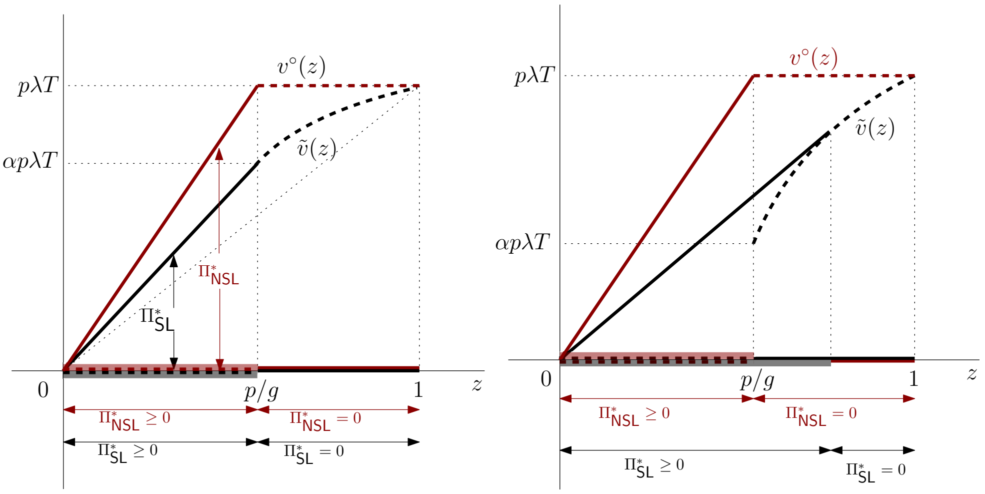

(a) High charisma
($\tilde{v}'(p/g^+) \leq \alpha\lambda T g$)

(b) Low charisma
($\tilde{v}'(p/g^+) > \alpha\lambda T g$)

Figure 8: Optimal influencer value with social learning (Π^*_{SL}, solid black) and without (Π^*_{NSL}, solid red) as a function of the prior μ^0. Dashed curves show the underlying revenue functions $\tilde{v}(\cdot)$ (black) and $v^\circ(\cdot)$ (red). Shaded regions on the horizontal axis indicate ranges of μ^0 where optimal influencers create positive value with social learning (black shade) and without (red shade). We use $\tilde{v}'(x^+)$ to denote the right differential of $\tilde{v}$ at x.

value functions are monotone, which ensures they are almost everywhere differentiable (see, e.g., page 112 of Royden and Fitzpatrick 1988). The proof is in Appendix F.4.

Lemma 8. The value functions $v^o(z)$ (baseline model) and $\tilde{v}(z)$ (social learning model) are both weakly increasing in $z \in [0, 1]$.

This monotonicity result is crucial: it allows us to apply Lemma 7 to characterize optimal influencers in both settings. For the baseline model without social learning, we can derive a closed-form solution.

Proposition 5 (Optimal influencer without social learning)**.** In the baseline model:

(i) If $\mu^0 < p/g$, the optimal influencer's value is $\Pi^* = \mu^0 p\lambda T/(p/g)$. The optimal information structure is $\pi_{hH} = 1$ and $\pi_{\ell L} = (p/g - \mu^0)/[(1-\mu^0)(p/g)]$ which induces post-signal beliefs $\mu^\ell = 0$ and $\mu^h = p/g$ with *a priori* probabilities $\tau(\mu^\ell) = 1 - \mu^0/(p/g)$ and $\tau(\mu^h) = \mu^0/(p/g)$.

(ii) If $\mu^0 \geq p/g$, optimal influencer's value is $\Pi^* = 0$.

The optimal influencer without social learning has a simple structure: it always sends a positive signal in the high state ($\pi_{hH} = 1$) and calibrates the true negative rate ($\pi_{\ell L}$) to place μ^h exactly at the purchase threshold p/g. Intuitively, the firm wants followers reach the purchase threshold after a high signal, but gains nothing from pushing beliefs higher.

With social learning, closed-form solutions are difficult to obtain. However, we can compute optimal values numerically using the algorithm in Appendix J. Figure 8 illustrates how optimal influencer value differs with and without social learning across different prior beliefs μ^0 and charisma levels α.

When the prior is low ($\mu^0 < p/g$), the figure shows that the optimal influencer without social learning creates more value than with social learning, regardless of charisma level. This aligns with the intuition from case (c) in Section 5: when customers need persuasion to purchase, social learning dilutes the influencer's impact by introducing competing organic signals.

However, when the prior is high ($\mu^0 \geq p/g$) and charisma is low (panel b), a striking reversal occurs: the optimal influencer *with* social learning can be more valuable than without. Without social learning, no influencer creates value when $\mu^0 \geq p/g$ since customers already purchase. With social learning, however, the optimal influencer raises skeptics' starting belief from μ^0 to μ^h, reducing the probability their beliefs drift below p/g over time. This reversal only occurs when charisma is low (many skeptics benefit from the raised starting point). Furthermore, the effect is weakest when μ^0 is just above p/g, since beliefs starting just above the threshold face high drift risk even after being raised to μ^h, limiting the protection the influencer can provide.

More broadly, Figure 8 shows that the set of instances where the optimal influencer with social learning creates positive value subsumes the set where the optimal influencer without social learning creates positive value. This pattern mirrors the observation from Theorem 1, but now applies to optimally designed influencers rather than arbitrary ones. This relationship is further illustrated in Figure 11 in Appendix A.

Beyond these comparisons of optimal influencer value, social learning fundamentally changes the structure of the optimal information policy. To formalize this, we introduce a standard definition from information economics.

Definition 1 (Adapted from Blackwell 1953)**.** Let τ and τ' be the distributions over posterior beliefs induced by the information structure π and π' under the same prior μ^0. We say that π *Blackwell dominates* π' if τ is a mean-preserving spread of τ'.[6]

Blackwell dominance provides a formal way to compare information structures. When π Blackwell dominates π', the posteriors induced by π have more spread while maintaining the same mean as those induced by π'. Greater spread means the signals are more informative: they distinguish more sharply between states. Our main result shows that the optimal information structure with social learning Blackwell dominates the optimal structure without social learning—in other words, social learning increases optimal informativeness.

Theorem 2. When the prior satisfies $\mu^0 \leq p/g$:

(i) The optimal influencer with and without social learning induces posteriors $\mu^\ell = 0$ and $\mu^h \in (0, 1)$, corresponding to information structure $\pi_{hH} = 1$ and $\pi_{\ell L} \in (0, 1)$.

(ii) Let π^{SL} and π^{NSL} denote optimal information structures with and without social learning. Then $\pi^{\mathsf{SL}}_{\ell L} \geq \pi^{\mathsf{NSL}}_{\ell L} > 0$, meaning π^{SL} Blackwell dominates π^{NSL}.

We call the information structure in part (i) h-biased because it biases followers toward observing positive signals: the influencer always sends signal h in a high state, but sends h with

[6] A distribution τ is a mean-preserving spread of τ' if there exists random variables $X \sim \tau$ and $X' \sim \tau'$ satisfying $\mathbf{E}(X|X') = X'$. Intuitively, each realization of X' is spread out into a random X with the same conditional mean.

positive probability even in the low state (since $\pi_{hL} = 1 - \pi_{\ell L} > 0$). Part (ii) reveals the key insight: with social learning, the optimal h-bias is weaker (higher $\pi_{\ell L}$ means lower false-positive rate π_{hL}), making the optimal influencer more informative.

This finding is counterintuitive. One might expect that abundant organic information would make precise influencer signals less important, leading firms to prefer less informative influencers. The opposite holds: social learning increases the value of informativeness. To see why, recall that higher informativeness creates more spread between μ^h and μ^ℓ. When combined with the concavity of $\tilde{v}(\cdot)$ and the diffusion process, this spread generates more value by strengthening initial beliefs so that they are more robust against drift.

Concretely, suppose an influencer with a 10% false-positive error rate is optimal with social learning. Then an influencer with a 20% false-positive error rate might be optimal without social learning. The firm should seek more accurate influencers when social learning is present.

Thus, social learning has two effects. First, it expands the set of instances where influencers create value—some influencers that would be rejected without social learning become worth sponsoring. Second, among settings where influencers are used, social learning shifts the optimal information structure toward greater informativeness. Both effects arise from the same mechanism: the interaction between the influencer's signal and the subsequent belief diffusion process that social learning enables.

Remark 3 (The impact of influencer's cost.)**.** When the influencer's cost is a function of its signaling structure, i.e., a function of π, then the optimal influencer problem is related to another literature: Bayesian persuasion with costs of acquiring information. Such costs are commonly modeled as (i) prior dependent cost, or (ii) prior independent cost. When the cost depends on prior belief, it is often interpreted as the cost of absorbing the information content relative to what the signal receiver already knows, so it is a cost related to customers, not the influencer. When the cost is independent of the prior, it is commonly modeled as a function of the noise. Pomatto et al. (2023) propose axioms that such cost functions should satisfy and characterize the cost functions over Blackwell experiments under their axioms. However, modeling the general cost of producing information has remained an unsolved problem. It is an interesting future direction to explore the costly information acquisition under social learning.

7 Privately informed firms

We now consider a setting where the firm has private information about the state $w \in \{H, L\}$, for example, from extensive test-market studies or prior experience in similar markets. After observing a private signal, the firm knows the true state. We refer to a firm that knows the state is w as the *type-w firm*. The type-w firm then chooses an influencer with information structure $\pi^w : S \times \Omega \mapsto [0, 1]$.[7]

Followers are rational and understand the firm is privately informed. In this setting, followers learn information through two channels: (1) from observing which influencer the firm selects,

[7] In this section, we focus on the optimal influencer only because we want to analyze how the initial private signal might affect the equilibrium choice of signaling strategy π.

and (2) from observing the signal that influencer subsequently sends.

The firm's choice of influencer potentially reveals information about its type. Let $\beta(\pi)$ denote followers' interim belief that the state is H after observing the firm's choice π, but before seeing the influencer's signal. This belief depends on which types would choose π in equilibrium:

$$\beta(\pi) = \begin{cases} 0 & \text{if only } L\text{-type chooses } \pi \\ 1 & \text{if only } H\text{-type chooses } \pi \\ \mu^0 & \text{if both types choose } \pi \end{cases} \tag{31}$$

If different types choose different information structures ($\pi^L \neq \pi^H$), we call this a *separating equilibrium* because the firm's choice perfectly reveals its type and the true state. If both types choose the same information structure ($\pi^L = \pi^H$), we call this a *pooling equilibrium* because the firm's choice reveals no information, leaving followers' interim belief at the prior μ^0.

After updating the prior to $\beta(\pi^w)$, followers observe the influencer's signal s generated according to π^w, and the update beliefs dynamically following Section 4.3. The only difference from the common prior setting is that the firm now has an extreme prior (0 or 1 since it knows the true state) while followers have interim belief $\beta(\pi^w)$. Given the firm's type w, the expected profit from choosing influencer π is:

$$\Pi\left(\pi, \beta(\pi), w\right) \triangleq \sum_{s \in S} \left(\mathbf{1}_{w=H}\pi(s|H) + \mathbf{1}_{w=L}\pi(s|L)\right) v(\mu^s) - v(\mu^0) \tag{32}$$

where $\mu^s = \beta(\pi)\pi(s|H) / \left(\beta(\pi)\pi(s|H) + (1 - \beta(\pi))\pi(s|L)\right)$. All the proofs of this section are in Appendix G.

This two-stage updating process means followers can learn the true state even in a pooling equilibrium if the firm pools on a fully revealing information structure. We distinguish two extremes: A *full disclosure signal* satisfies $\pi_{hH} = \pi_{\ell L} = 1$, meaning the influencer's signal perfectly reveals the true state. A *no disclosure signal* leaves the posterior belief unchanged from the interim belief $\beta(\pi)$.

Our main result characterizes equilibrium behavior when the firm is informed.

Theorem 3. When the firm is privately informed about the state:
(i) All equilibrium outcomes can be supported by a pooling equilibrium where $\pi^H = \pi^L$.
(ii) In any equilibrium, we have $\pi^H = \pi^L = \pi^{\mathsf{FD}}$, where π^{FD}. In other words, followers learn the true state.

Theorem 3 has a striking implication: when the firm knows the true state, it always uses a full disclosure influencer in equilibrium. The reasoning is straightforward. The high-type firm prefers full disclosure because $v(1) \geq v(z)$ for any $z < 1$. The low-type firm cannot benefit from choosing any other structure because followers would correctly infer the firm is low type. Therefore, both types pool on the fully revealing structure, and followers learn the true state—not from the choice of influencer (since both types make the same choice), but from the influencer's subsequent signals. This result holds regardless of whether social learning is present, as followers learn the state immediately from the influencer's signal.

This changes the influencer's role. Strategic information design (i.e., choosing how much to reveal) becomes redundant because incentive compatibility forces complete revelation. However, the influencer is essential for credible communication: the firm cannot simply announce "our product is high quality," as followers would rationally discount such claims. In settings where firms have verified quality information, influencer marketing serves primarily to credibly transmit this known information rather than to strategically manage uncertainty.

Connecting back to our main model reveals a striking implication: private information can harm the firm. An uninformed firm can implement sophisticated partial disclosure strategies through Bayesian persuasion (e.g., choosing an influencer with $\pi(h|H) = 0.7$, $\pi(h|L) = 0.3$). In contrast, an informed firm must fully disclose despite having access to the same commitment technology.

Why does commitment fail to enable partial disclosure when the firm is informed? This connects to the cheap talk literature (Crawford and Sobel, 1982) where a firm can send costless messages without commitment. Our firm possesses more power than in cheap talk: it can commit *ex-ante* to any information structure. Yet Theorem 3 shows this commitment becomes irrelevant once the firm has private information. The reason is that the firm's *choice* of information structure is itself observable. Despite having commitment power unavailable in cheap talk, the informed firm cannot avoid full separation.

The economic consequence is substantial. An informed firm's *ex-ante* revenue (before learning the state) is $\mu_0 v(1) + (1 - \mu_0)v(0)$, lying on the line segment connecting $v(0)$ and $v(1)$. As shown in Figure 8, this is always weakly below the concave curve achieved by an uninformed firm using optimal Bayesian persuasion. The firm would prefer to remain ignorant since commitment to partial disclosure (available when uninformed) is more valuable than conditioning on the realized state (available when informed). This "ignorance is bliss" result highlights why the uninformed firm setting in Sections 5 and 6 (where firms genuinely do not know product quality and followers understand this uncertainty) may better reflect modern marketing environments. It is precisely in this setting where the interaction between influencer signals and social learning creates the richest strategic considerations.

8 Conclusion

This paper analyzes how influencer marketing interacts with social learning in product markets. We develop a model where firms can sponsor influencers to send signals about product quality, and customers may learn both from these sponsored signals and from organic information generated by other customers' purchases. Our analysis characterizes when influencers create value and how firms should optimally design information structures, yielding several key insights.

First, social learning affects when and by how much influencers are valuable to the firm. When customers need persuasion to purchase (low priors), the influencer can be beneficial with or without social learning, but its value under social learning is always smaller, as social learning introduces competing organic signals. However, when customers already intend to purchase (high priors), the influencer can be even more effective under social learning than without social learning. This occurs when influencers raise skeptics' starting beliefs enough so that the customers' beliefs become robust to organic negative signals that can be generated from social learning. As a result, the influencer may become more likely to be valuable with social learning. The set of parameter values where influencers create value is strictly larger with social learning than without, a contradiction to the intuition that abundant organic information should diminish influencers' impact.

Second, when firms can choose among influencers, social learning increases the optimal informativeness. Firms should seek more accurate influencers (lower false-positive rates) when social learning is present compared to when it is absent. This finding is counterintuitive: one might expect abundant organic information to make precise influencer signals less important. Instead, higher informativeness creates a greater spread in posterior beliefs, which strengthens skeptics' prior beliefs to be robust against alternative information.

Third, when firms have private information about the state, these strategic considerations disappear. Incentive compatibility forces both high- and low-quality firms to pool on fully revealing information structures, making followers learn the true state immediately. This "ignorance is bliss" result shows that private information can harm firms: commitment to information structures only creates value when genuine uncertainty exists.

The traditional distinction between paid advertising and organic word-of-mouth may be too stark. Influencer marketing occupies a middle ground: it is paid persuasion that operates through trusted intermediaries, whose signals interact dynamically with organic information flows. Understanding these interactions becomes increasingly important as social learning mechanisms proliferate across digital platforms, i.e. from product reviews to public health messaging. The analytical tools we develop here, particularly for handling dynamic belief diffusion in networks, provide a foundation for analyzing how actors can strategically navigate complex information environments where centralized control gives way to distributed learning.

Several extensions could enrich our framework and broaden its applicability. First, one could examine how signal format choices (text, photo, video) as costly signals affect influence value. In public health contexts, for instance, video testimonials from healthcare workers may be more credible than text posts, but also more expensive to produce and verify. Second, the timing of information release deserves attention. In the video game space, for example, it has traditionally been a sign of poor quality that a game is released to reviewers close to the launch date. However, recent trends have pushed the release of preview copies later.[8] Public health authorities must decide when to release information during evolving crises: releasing too early risks misinformation to take root, while releasing too late allows competing narratives to take hold. Lastly, dynamic disclosure policies merit investigation, where firms sequentially sponsor different influencers. However, one needs to be careful about the applicability of the model assumptions. For example, one would need to assume that influencers have similar groups of followers.

[8] For a discussion, see the article: https://arstechnica.com/gaming/2016/10/why-early-reviews-of-video-games-are-getting-rarer-and-rarer/

References

Saed Alizamir, Francis de Véricourt, and Shouqiang Wang. Warning against recurring risks: An information design approach. *Management Science*, 66(10):4612–4629, 2020.

Ricardo Alonso and Odilon Câmara. Persuading voters. *American Economic Review*, 106(11):3590–3605, 2016a.

Ricardo Alonso and Odilon Câmara. Political disagreement and information in elections. *Games and Economic Behavior*, 100:390–412, 2016b.

Itai Arieli and Yakov Babichenko. Private Bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.

Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.

Alexander Barvinok. *A Course in Convexity*, volume 54. American Mathematical Soc., 2002.

Nicole Bäuerle. Optimal control of queueing networks: An approach via fluid models. *Advances in Applied Probability*, 34(2):313–328, 2002.

Richard Bennett. *Top 10 Richest YouTubers in 2024*, 2024. URL https://filmora.wondershare.com/vlogger/top-wealthiest-youtubers.html. Accessed on 8 June 2024.

Omar Besbes and Marco Scarsini. On information distortions in online ratings. *Operations Research*, 66(3):597–610, 2018.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.

David A Blackwell and Meyer A Girshick. *Theory of Games and Statistical Decisions*. Courier Corporation, 1979.

Tamer Boyaci, Soudipta Chakraborty, and Huseyin Gurkan. Persuading skeptics and fans in the presence of additional information. *Production and Operations Management*, pages 1142–1154, 2022.

Ozan Candogan. Persuasion in networks: Public signals and k-cores. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 133–134, 2019.

Ozan Candogan and Kimon Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research*, 2020.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

Laurens Debo and Nicola Secomandi. Revenue management for products with unknown quality and observable price and sales history. *Carnegie Mellon University. Journal contribution*, 2018. URL https://doi.org/10.1184/R1/6707804.v1.

Chrysanthos Dellarocas. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management Science*, 52(10):1577–1593, 2006.

Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, 1986. ISBN 0-471-08186-8. doi: 10.1002/9780470316658.

Stewart N Ethier and Thomas G Kurtz. *Markov Processes: Characterization and Convergence*. John Wiley & Sons, 2009.

Itay P. Fainmesser and Andrea Galeotti. The market for online influence. *American Economic Journal: Microeconomics*, 13(4):332–72, November 2021.

Paul D Feigin. Stable convergence of semimartingales. *Stochastic Processes and their Applications*, 19(1):125–134, 1985.

Avner Friedman. *Partial Differential Equations of Parabolic Type*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1964.

Joe Gagliese. *The Rise Of The Influencer: Predictions For Ways They'll Change The World*, 2022. URL https://www.forbes.com/sites/theyec/2022/07/08/the-rise-of-the-influencer-predictions-for-ways-theyll-change-the-world/. Accessed on 8 June 2024.

Werner Geyser. *https://influencermarketinghub.com/influencer-marketing-benchmark-report/*, 2024. URL https://influencermarketinghub.com/influencer-marketing-benchmark-report/. Accessed on 8 June 2024.

Peter W Glynn. Diffusion approximations. *Handbooks in Operations Research and Management Science*, 2: 145–198, 1990.

Tom Hamami. Network effects, bargaining power, and product review bias: Theory and evidence. *The Journal of Industrial Economics*, 67(2):372–407, 2019.

Joe Hitchcock. Influencer pricing: The cost of influencers in 2025, January 2025. URL https://www.shopify.com/uk/blog/influencer-pricing. Accessed: 2025-04-28.

Bar Ifrach, Costis Maglaras, Marco Scarsini, and Anna Zseleva. Bayesian social learning from consumer reviews. *Operations Research*, 67(5):1209–1221, 2019.

InfluencerMarketingHub. *The State of Influencer Marketing 2024: Benchmark Report*, 2024. URL https://influencermarketinghub.com/influencer-marketing-benchmark-report/.

Maya Jalloul and Vasiliki Kostami. Advertising by recruiting influencers. *HEC Paris Research Paper No. MOSI-2022-1454*, 2022.

Björn Johansson. On the central limit theorem for point process martingales. *Statistics & Probability Letters*, 20 (2):125–130, 1994.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Robert V Kozinets, Kristine De Valck, Andrea C Wojnicki, and Sarah JS Wilner. Networked narratives: Understanding word-of-mouth marketing in online communities. *Journal of Marketing*, 74(2):71–89, 2010.

Can Küçükgül, Özalp Özer, and Shouqiang Wang. Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Management Science*, 68(7):4899–4918, 2022.

Xiao Lin and Ce Liu. Credible persuasion. *Journal of Political Economy*, 132(7):000–000, 2024.

Costis Maglaras, Marco Scarsini, Dongwook Shin, and Stefano Vaccari. Product ranking in the presence of social learning. *Operations Research*, 71:1136–1153, 2022.

Dina Mayzlin. Promotional chat on the Internet. *Marketing Science*, 25(2):155–163, 2006.

Mohamed Mostagir and James Siderius. Strategic reviews. *Management Science*, 69(2):904–921, 2023.

Yiangos Papanastasiou and Nicos Savva. Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*, 63(4):919–939, 2016.

Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information: The case of constant marginal costs. *American Economic Review*, 113(5):1360–1393, 2023. doi: 10.1257/aer.20190185.

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.

Halsey Lawrence Royden and Patrick Fitzpatrick. *Real Analysis*. Macmillan New York, 1988.

Alexander Shapiro. On duality theory of conic linear problems. In *Semi-Infinite Programming*, pages 135–165. Springer, 2001.

Steven E Shreve. *Stochastic Calculus for Finance II*. Springer Science & Business Media, 2004.

M Zhao and M Xie. On maximum likelihood estimation for a general non-homogeneous Poisson process. *Scandinavian Journal of Statistics*, 23(4):597–607, 1996.

Noa Zychlinski. Applications of fluid models in service operations management. *Queueing Systems*, 103(1): 161–185, 2023.

Online Appendices to "Valuing Influence with Social Learning"

Hyun-Soo Ahn[1], **Christopher Thomas Ryan**[2], **Joline Uichanco**[3], **Mengzhenyu Zhang**[4]
[1]University of Michigan, Ross School of Business, `hsahn@umich.edu`
[2]University of British Columbia, Sauder School of Business, `chris.ryan@sauder.ubc.ca`
[3]New York University, Tandon School of Engineering, `ju441@nyu.edu`
[4]University College London, School of Management, `zhenyu.zhangmeng@ucl.ac.uk`

Appendix A Additional figures

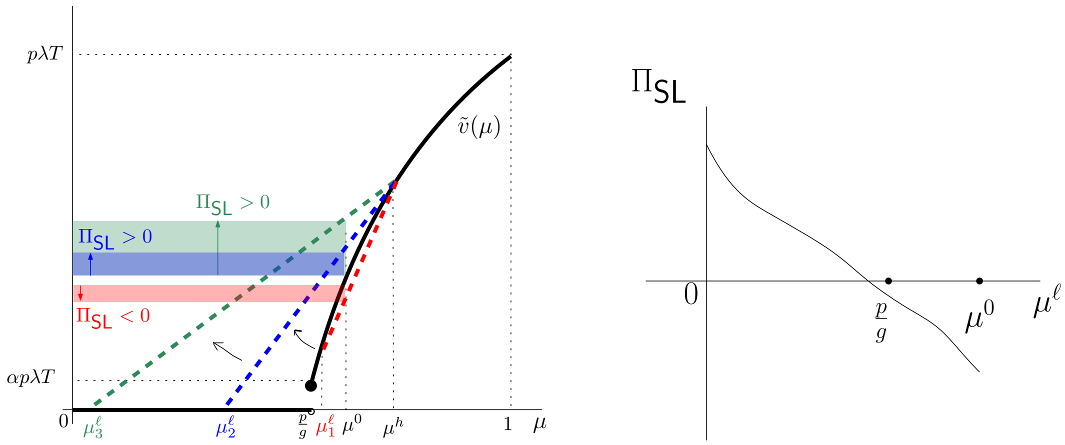

Figure 9: Π_{SL} increases as μ^{ℓ} decreases in case (b) of Figure 4.

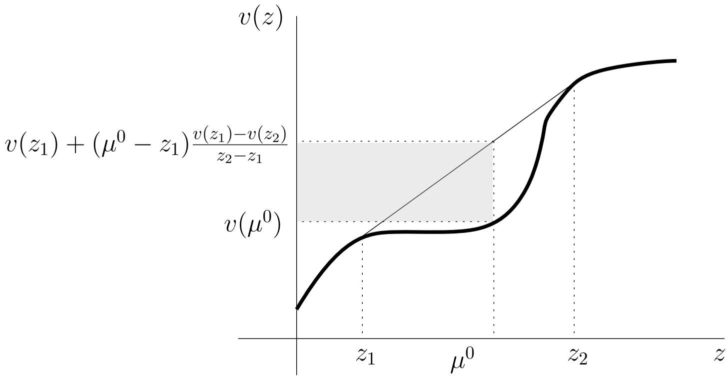

Figure 10: Visualization of the optimal value in Lemma 7. The solid line connecting $(z_1, v(z_1))$ and $(z_2, v(z_2))$ is $\lambda_1^* + z\lambda_2^*$.

Appendix B Visualizing the value of the influencer

Example A.1. Suppose the common prior belief is that the low state and the high state are equally likely to occur; i.e., $\mu^0(H) = \mu^0(L) = 1/2$. The influencer's informativeness π has

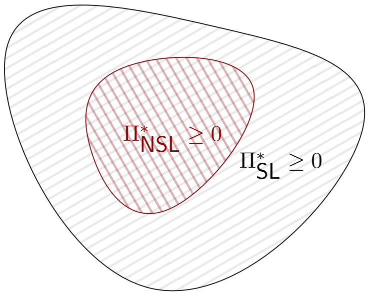

Figure 11: Parameter values $(\alpha, p/g, \mu^0)$ that make the optimal influencer valuable

$\pi_{hH} = 0.8$, $\pi_{\ell H} = 0.2$, $\pi_{hL} = 0.4$, and $\pi_{\ell L} = 0.6$.

The post-signal priors μ^h and μ^ℓ are computed *a posteriori* (after observing the signal) using Bayes' rule according to (4). However, before the signal is realized, we can compute the *a priori* probability of each post-signal belief using (9): $\tau(\mu^h)$ represents the probability of observing μ^h and $\tau(\mu^\ell)$ represents the probability of observing μ^ℓ.

Specifically, we have post-signal priors

$$\mu^h = \frac{0.5 \times 0.8}{0.5 \times 0.8 + 0.5 \times 0.4} = \frac{2}{3}$$
$$\mu^\ell = \frac{0.5 \times 0.2}{0.5 \times 0.2 + 0.5 \times 0.6} = \frac{1}{4}$$

with *a priori* probabilities

$$\tau(2/3) = 0.5 \times 0.8 + 0.5 \times 0.4 = 0.6,$$
$$\tau(1/4) = 1 - \tau(2/3) = 0.4.$$

From (10), the influencer value Π can be expressed in the following way

$$\Pi = \tau\left(\tfrac{2}{3}\right) v\left(\tfrac{2}{3}\right) + \tau\left(\tfrac{1}{4}\right) v\left(\tfrac{1}{4}\right) - v\Bigg(\underbrace{\tau\left(\tfrac{2}{3}\right)\tfrac{2}{3} + \tau\left(\tfrac{1}{4}\right)\tfrac{1}{4}}_{=\mu_0(H)=\frac{1}{2}}\Bigg). \tag{A.1}$$

Hence, we can visualize Π geometrically as the height of the gray rectangle in Figure 12. In particular, it is the difference evaluated at the point $z = 1/2$ between the curve $v(\cdot)$ and the secant line that passes through $(1/4, v(1,4))$ and $(2/3, v(2/3))$.

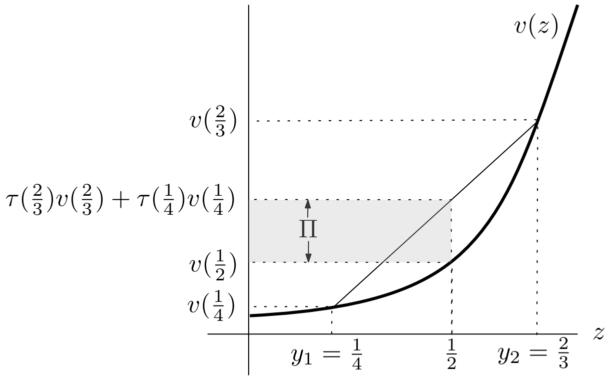

Figure 12: Visualization of Π in Example A.1 as the vertical height of the gray-colored rectangle. By (10), the horizontal axis is the post-signal prior z.

Appendix C When demand information is revealed, skeptics and devotees have the same belief over time

Lemma A.1. For a given arrival rate λ and a prior belief in the high state μ (i.e., μ^0 or μ^s), if the only information revealed at time t is the demand history $\{D_u, u \leq t\}$ then the demand process is a Poisson process with some constant rate

$$\Lambda(\mu) = \begin{cases} \lambda & \text{if } \mu > p/g, \\ 0 & \text{otherwise.} \end{cases} \tag{A.2}$$

In particular, charisma α does not affect the value of the influencer.

Proof. We know that the arrival rate λ is not state-dependent. Also, a follower's purchase decision is not state-dependent. Thus, for a given state $j \in \{H, L\}$, the likelihood of $\{D_u, u \leq t\}$ is

$$\mathbb{P}(\{D_u, u \leq t\} \mid j) = \prod_{k=1}^{D_t} f(\alpha, \mu_0, \mu_{t_k}) \cdot \lambda^{D_t} \exp\left(-\lambda \int_0^t f(\alpha, \mu_0, \mu_u)\right) du.$$

Then, by Bayes' rule, we have

$$\mu_t = \frac{\mathbb{P}\left(\{D_u, u \leq t\} \mid H\right) \mu_0}{\mathbb{P}\left(\{D_u, u \leq t\} \mid H\right) \mu_0 + \mathbb{P}\left(\{D_u, u \leq t\} \mid L\right) (1 - \mu_0)} = \mu_0$$

where μ_0 equals to μ^0 or μ^s. This means that follower beliefs are always equal to the prior. This concludes the proof. □

Appendix D Proofs of Section 4

D.1 Proof of Lemma 1

Proof. Note that (10) is a direct reformulation of (8) following the definition of $\tau(\cdot)$. Because the post-signal prior is computed from Bayes' rule (4), we can verify that $\mathbb{E}[\mu^s] = \mu^0$, where the left-hand side is the expected value of the belief for the high state after the influencer's signal. Note that this equation is equivalent to the following: $\tau(\mu^\ell)\mu^\ell + \tau(\mu^h)\mu^h = \mu^0$. □

D.2 Proof of Proposition 2

The following is a technical lemma derived from Proposition 1 to characterize the value of the influencer for any fixed ratio $\frac{\pi_{\ell L}}{\pi_{\ell H}}$, which is indicative of the accuracy of the influencer because it is the ratio of a true negative to a false negative. Note that (5) implies $\pi_{\ell L} + \pi_{hH} \geq 1$ and $\frac{\pi_{\ell L}}{\pi_{\ell H}} \geq 1$.

Lemma A.2. Keeping $\beta \triangleq \frac{\pi_{\ell L}}{\pi_{\ell H}}$ fixed, if $\frac{\beta-1}{\beta} \leq \pi_{hH} \leq \frac{\beta-1}{\beta - \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}}$ and $\mu^0 < p/g$, we have $\Pi > 0$; otherwise, we have $\Pi \leq 0$.

Proof of Lemma A.2. Let $\pi_1 \triangleq \pi_{hH}, \pi_2 \triangleq \pi_{\ell L}$. Also, in this proof, we use μ_0 to denote $\mu^0(H)$ for simpler expressions. Then, by definition, we have

$$y_1 = \frac{\mu_0(1-\pi_1)}{\mu_0(1-\pi_1) + (1-\mu_0)\pi_2} = \frac{\mu_0}{\mu_0 + (1-\mu_0)\frac{\pi_2}{1-\pi_1}} = \frac{\mu_0}{\mu_0 + (1-\mu_0)\beta},$$

$$y_2 = \frac{\mu_0\pi_1}{\mu_0\pi_1 + (1-\mu_0)(1-\pi_2)} = \frac{\mu_0}{\mu_0 + (1-\mu_0)\frac{1-\pi_2}{\pi_1}} = \frac{\mu_0}{\mu_0 + (1-\mu_0)\frac{1-\beta+\beta\pi_1}{\pi_1}}.$$

Thus, by fixing β, we fix y_1. Also, by assuming $\beta \geq 1$, we have $y_1 \leq y_2$.

Letting $y_2 \geq p/g$, we have

$$\frac{\mu_0}{\mu_0 + (1-\mu_0)\frac{1-\beta+\beta\pi_1}{\pi_1}} \geq p/g,$$

which gives us

$$\pi_1 \leq \frac{\beta-1}{\beta - \frac{\mu_0}{1-\mu_0}\frac{1-p/g}{p/g}}.$$

Moreover, $y_2 \leq 1$, which yields

$$\pi_1 \geq \frac{\beta-1}{\beta}.$$

Then, according to Proposition 1, we prove Lemma A.2. □

Proof of Proposition 2. In Lemma A.2, we derive the cases that give us $\Pi > 0$ and $\Pi \leq 0$ for fixed β. Also, we know $\beta \geq 1$. Therefore, we can identify the boundary by solving the following curve parameterized by β with $\beta \geq 1$:

$$\pi_{hH} = \frac{\beta-1}{\beta - \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}}$$

$$\pi_{\ell L} = -\beta\frac{\beta-1}{\beta - \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}} + \beta.$$

This gives us

$$\pi_{\ell L} = 1 - \frac{\mu^0}{1-\mu^0}\frac{1-p/g}{p/g}\pi_{hH}.$$

□

D.3 Proof of Lemma 2

Proof. According to Bayes rule, we have

$$
\begin{aligned}
\mu_t &= \frac{\mathbb{P}(\mathcal{H}_t \mid H)\,\mu_0}{\mathbb{P}(\mathcal{H}_t \mid H)\,\mu_0 + \mathbb{P}(\mathcal{H}_t \mid L)\,(1-\mu_0)} \\
&= \frac{\mu_0(\pi^R_{hH})^{R_{ht}}(\pi^R_{\ell H})^{R_{\ell t}}}{\mu_0(\pi^R_{hH})^{R_{ht}}(\pi^R_{\ell H})^{R_{\ell t}} + (1-\mu_0)(\pi^R_{hL})^{R_{ht}}(\pi^R_{\ell L})^{R_{\ell t}}} \\
&= \frac{\mu_0}{\mu_0 + (1-\mu_0)\left(\frac{\pi^R_{hL}}{\pi^R_{hH}}\right)^{R_{ht}}\left(\frac{\pi^R_{\ell L}}{\pi^R_{\ell H}}\right)^{R_{\ell t}}}.
\end{aligned}
$$

□

D.4 Proof of Lemma 3

Proof. To prove the proposition, we need Itô's lemma for a jump process (see Shreve (2004) Section 11.5.1). Specifically, suppose that $Z_t = X_t + J_t$ is a stochastic process, where J_t is a pure jump process and X_t is a continuous-path process with differential form:

$$X_t = X_0 + \int_0^t \Gamma_s ds + \int_0^t \Theta_s ds,$$

where t is a standard Wiener process (Brownian motion) and Γ_s and Θ_s are adapted processes. Then the following theorem (Theorem 11.5.1 in Shreve 2004) provides the expression for the dynamics of a function of Z_t.

Theorem 4 (Itô-Doeblin formula for a jump process)**.** Let Z_t be a jump process and $h(x)$ a function for which $h'(x)$ and $h''(x)$ are defined and continuous. Then,

$$h(Z_t) = h(Z_0) + \int_0^t h'(Z_s)dX_s + \frac{1}{2}\int_0^t h''(Z_s)dX_s dX_s + \sum_{0<s\le t}\left[h(Z_s) - h(Z_{s-})\right].$$

Armed with Theorem 4, we are now ready to prove the proposition.

Define $Z_t \triangleq R_{ht}\ln\frac{\pi^R_{hL}}{\pi^R_{hH}} + R_{\ell t}\ln\frac{\pi^R_{\ell L}}{\pi^R_{\ell H}}$. Then we have $\mu_{Ht} = h(Z_t)$ where

$$h(x) \triangleq \frac{\mu_0}{\mu_0 + (1-\mu_0)e^x}.$$

According to Theorem 4 in its differential form, the following holds:

$$
\begin{aligned}
d\mu_t &= \left[h\left(Z_{t-} + \ln\frac{\pi^R_{hL}}{\pi^R_{hH}}\right) - h(Z_{t-})\right] dR_{ht} + \left[h\left(Z_{t-} + \ln\frac{\pi^R_{\ell L}}{\pi^R_{\ell H}}\right) - h(Z_{t-})\right] dR_{\ell t} \\
&= \frac{\mu_{t-}(1-\mu_{t-})\left(\pi^R_{hH} - \pi^R_{hL}\right)}{\pi^R_{hH}\mu_{t-} + \pi^R_{hL}(1-\mu_{t-})} dR_{ht} \\
&\quad + \frac{\mu_{t-}(1-\mu_{t-})\left(\pi^R_{\ell H} - \pi^R_{\ell L}\right)}{\pi^R_{\ell H}\mu_{t-} + \pi^R_{\ell L}(1-\mu_{t-})} dR_{\ell t}.
\end{aligned}
$$

□

D.5 Proof of Lemma 5

Proof. We only need to check the conditions of the martingale central limit theorem (Feigin, 1985; Johansson, 1994; Ethier and Kurtz, 1986). First, we have $0 < \mu_t^{(n)} < 1$ almost surely. Second, we have $|\mu_t^{(n)} - \mu_s^{(n')}| \le 2$ almost surely for any n and n' so the jumps vanishes as $n \to \infty$. Third, we know $\{\mu_t^{(n)}, t \geq 0\}$ is square integrable and $\mathbf{E}(\langle \mu^{(n)} \rangle_t)$ converges as $n \to \infty$ because R_{ht} and $R_{\ell t}$ are nonhomogeneous poison processes. Lastly, the limits 0 and 1 are absorbing.

□

D.6 Proof of Proposition 3

Proof. When $z \geq p/g$, following the definition diffusion approximation (for example, Ethier and Kurtz (2009) Theorem 4.1), the approximated expected revenue-to-go function $\tilde{V}(y, \bar{t}; z)$ for $y > p/g$ is governed by the following partial differential equation

$$
\begin{aligned}
\frac{\partial \tilde{V}(y,\bar{t};z)}{\partial \bar{t}} &= \lambda p + \Big[\left(\pi_{hH}^R y + \pi_{hL}^R (1-y)\right) \frac{\partial \tilde{V}(y,\bar{t};z)}{\partial y} J_{hy} \\
&+ \left(\pi_{\ell H}^R y + \pi_{\ell L}^R (1-y)\right) \frac{\partial \tilde{V}(y,\bar{t};z)}{\partial y} J_{\ell y} \Big] \cdot \lambda\gamma \\
&+ \frac{1}{2} \Big[\left(\pi_{hH}^R y + \pi_{hL}^R (1-y)\right) \frac{\partial \tilde{V}^2(y,\bar{t};z)}{\partial y^2} J_{hy}^2 \\
&+ \left(\pi_{\ell H}^R y + \pi_{\ell L}^R (1-y)\right) \frac{\partial \tilde{V}^2(y,\bar{t};z)}{\partial y^2} J_{\ell y}^2 \Big] \cdot \lambda\gamma \\
&= \lambda p + \frac{1}{2} \frac{\lambda\gamma y^2 (1-y)^2 \left(\pi_{hH}^R - \pi_{hL}^R\right)^2}{\left(\pi_{hH}^R y + \pi_{hL}^R (1-y)\right)\left(\pi_{\ell H}^R y + \pi_{\ell L}^R (1-y)\right)} \frac{\partial^2 \tilde{V}(y,\bar{t};z)}{\partial y^2}.
\end{aligned}
\tag{A.3}
$$

The boundary conditions are

$$
\begin{aligned}
\tilde{V}(y,0;z) &= 0, && \text{for any } y \in [p/g, 1], \\
\tilde{V}(1,\bar{t};z) &= \lambda p t, && \text{for any } \bar{t} \in [0,T], \text{ and} \\
\tilde{V}(p/g,\bar{t};z) &= \alpha\lambda p \bar{t}, && \text{for any } \bar{t} \in [0,T].
\end{aligned}
$$

□

D.7 Proof of Lemma 6

Proof. When $\mu \in [0, p/g)$, we know that $f(\alpha, \mu, \mu) = 0$. Thus, we have $\tilde{v}(\mu) = \tilde{V}(\mu, T; \mu) = V(\mu, T; \mu) = 0$. Next consider any given $\mu \in [p/g, 1]$. We know $\frac{\partial \tilde{V}(y,\bar{t};z)}{\partial \bar{t}} \leq \lambda p$ because λp is the maximum revenue rate we can have no matter the belief of skeptics. Moreover, we have

$$
\frac{\lambda\gamma y^2 (1-y)^2 \left(\pi_{hH}^R - \pi_{hL}^R\right)^2}{\left(\pi_{hH}^R y + \pi_{hL}^R (1-y)\right)\left(\pi_{\ell H}^R y + \pi_{\ell L}^R (1-y)\right)} \geq 0,
$$

then according to (A.3), we have $\left.\frac{\partial^2 \tilde{V}(y,\bar{t};z)}{\partial y^2}\right|_{y=\mu} \leq 0$. This concludes the proof. □

Appendix E Proofs of Section 5

E.1 Proof of Theorem 1

Proof. $\Pi_{\mathsf{NSL}} > 0$ if and only if $\mu^0 \leq p/g \leq \mu^h$. From Figure 4c and Lemma 6, we know $\Pi_{\mathsf{NSL}} \geq \Pi_{\mathsf{SL}} > 0$.

$\Pi_{\mathsf{NSL}} = 0$ if and only if $p/g \leq \mu^\ell \leq \mu^0 \leq \mu^h$ or $\mu^\ell \leq \mu^0 \leq \mu^h \leq p/g$. From Figure 4a, Figure 4d and Lemma 6, we know $\Pi_{\mathsf{SL}} \leq \Pi_{\mathsf{NSL}} = 0$.

$\Pi_{\mathsf{NSL}} < 0$ if and only if $\mu^\ell \le p/g \le \mu^0 \le \mu^h$. From Figure 4b, Figure 5i and Lemma 6, we know $\Pi_{\mathsf{SL}} > \Pi_{\mathsf{NSL}}$. □

Appendix F Proofs of Section 6

F.1 Proof of Proposition 4

Proof. We first show that for any feasible solution π to (26), we can construct a feasible solution τ to (27) and they yield the same value of the objective functions. For any $s \in S$, we simply let $\tau(z) \triangleq \int_{s:\mu(H|s)=z} \sum_{w\in\Omega} \mu^0(w)\pi(s \mid w)\mathrm{d}s$ and $\mu(w \mid s) \triangleq \frac{\mu^0(w)\pi(s|w)}{\sum_{w'\in\Omega}\mu^0(w')\pi(s|w')}$.

This solution τ is feasible to (27) because it satisfies constraint (27b):

$$\mathbb{E}_\tau(z) = \int_0^1 z \int_{s:\mu(H|s)=z} \sum_{w\in\Omega} \mu^0(w)\pi(s \mid w)\mathrm{d}s\mathrm{d}z \tag{A.4}$$

$$= \int_0^1 \sum_{w\in\Omega} \mu^0(w)\pi(s \mid w)\tfrac{\mu^0(H)\pi(s|H)}{\sum_{w\in\Omega}\mu^0(w)\pi(s|w)}\mathrm{d}s \tag{A.5}$$

$$= \int_0^1 \mu^0(H)\pi(s|H)\mathrm{d}s = \mu^0(H).$$

Here, (A.4) comes from the construction of $\tau(z)$ and (A.5) comes from the construction of $\mu(w \mid s)$. It is easy to check that τ also satisfies the remaining constraints. Further, the objective value of τ in (27a) is the same as the objective value of π in (26) since:

$$\mathbb{E}_\tau[v(z)] = \int_0^1 \left[\int_{s:\mu(H|s)=z} \sum_{w\in\Omega} \mu^0(w)\pi(s \mid w)\right] v(z)\mathrm{d}s\mathrm{d}z = \sum_{w\in\Omega} \mu^0(w) \int_S \pi(s \mid w) v\left(\mu(H|s)\right)\mathrm{d}s \tag{A.6}$$

Next, we show the opposite direction. For the optimal solution τ to (27), we construct a feasible solution π, and show they have the same objective values. We first show that the optimal solution to the linear program (27) is a distribution with at most two points in its support. Because the optimal solution is an extremal distribution for problem (27). It suffices to show that the extremal distributions are discrete with support on at most two points. Proof of Lemma A.3 can be found in Appendix F.2.

Lemma A.3 (Extremal distributions)**.** The extremal distributions for problem (27) are discrete with support on at most two points.

Consider $\tau(\cdot)$ to be a solution to (27) with a two-point support z^ℓ and z^h, where $z^\ell \le z^h$. Note that:

$$\mathbb{E}_\tau[v(z)] = \tau(z^\ell)v(z^\ell) + \tau(z^h)v(z^h)$$

Let $\pi(s|w)$ be a solution to the following system equations:

$$\tau(z^\ell) = (1-\mu^0)\pi(\ell|L) + \mu^0\pi(\ell|H)$$
$$\tau(z^h) = (1-\mu^0)\pi(h|L) + \mu^0\pi(h|H)$$
$$1 = \pi(h|L) + \pi(\ell|L)$$
$$1 = \pi(h|H) + \pi(\ell|H)$$
$$\mu^0 = z^\ell\tau(z^\ell) + z^h\tau(z^h)$$

We next check that this constructed solution has the same objective value as τ:

$$\begin{aligned}
&\sum_{w\in\Omega}\mu^0(w)\sum_{s\in S}\pi(s|w)v(\mu^s)\\
&= \mu^0\left(\pi(h|H)v(z^h) + \pi(\ell|H)v(z^\ell)\right) + (1-\mu^0)\left(\pi(h|L)v(z^h) + \pi(\ell|L)v(z^\ell)\right)\\
&= v(z^\ell)\left[\mu^0\pi(\ell|H) + (1-\mu^0)\pi(\ell|L)\right] + v(z^h)\left[\mu^0\pi(h|H) + (1-\mu^0)\pi(h|L)\right]\\
&= v(z^\ell)\tau(z^\ell) + v(z^h)\tau(z^h) = \mathbb{E}_\tau[v(z)]
\end{aligned}$$

Thus, completing the proof. □

F.2 Proof of Lemma A.3

Proof. For problem (27), there is only one moment constraint (27b). According to Lemma 3.1 in Shapiro (2001), the support of the extremal distributions contains at most 2 points. □

F.3 Proof of Lemma 7

Proof. For case (i), λ_1^* and λ_2^* are solved from the following system of linear equations:

$$\lambda_1^* + y_1\lambda_2^* = v(z_1), \lambda_1^* + z_2\lambda_2^* = v(z_2).$$

With the primal solution defined in (29), we can easily check it generates the same objective value as $(\lambda_1^*, \lambda_2^*)$, where

$$\lambda_1^* = v(z_1) - z_1\frac{v(z_2)-v(z_1)}{z_2-z_1},$$
$$\lambda_2^* = \frac{v(z_2)-v(z_1)}{z_2-z_1}.$$

For case (ii), it is the case where there is only one active constraint at the optimal solution. The only way we can have this case is when the objective function is tangent to or coincide with the boundary of the feasible set generated by the infinitely many linear constraints (28b).

If part of the boundary of the feasible set is $\lambda_1 = -\mu^0\lambda_2 + v(\mu^0)$ for some range of λ_2, then the optimal value is $v(\mu^0)$. This is the only case where the objective function coincides with the part of the boundary, and there are infinitely many optimal solutions generating the objective value of $v(\mu^0)$.

For the other case that the objective function is tangent to the boundary at a single point that gives us the optimal solution, we must have that the boundary is continuous, differentiable and strictly convex around the tangent point (otherwise, we have two active constraints or part of the boundary is $\lambda_1 = -\mu^0\lambda_2 + v(\mu^0)$). We define $F(z,\lambda_1,\lambda_2) \triangleq \lambda_1 + z\lambda_2 - v(z)$. We must have the following relationship satisfied by (λ_1, λ_2) to describe this part of the boundary

$$\lim_{\epsilon \to 0} \frac{\lambda_1 + (z+\epsilon)\lambda_2 - v(z+\epsilon) - [\lambda_1 + z\lambda_2 - v(z)]}{\epsilon} = 0. \tag{A.7}$$

This is because any point (λ_1, λ_2) on the boundary should satisfy $F(z+\epsilon, \lambda_1, \lambda_2) = 0$ for any ϵ.

We define the subdifferential of $v(\cdot)$ as $v'(z) \triangleq \lim_{\epsilon \to 0} \frac{v(z+\epsilon)-v(z)}{\epsilon}$. Therefore, according to (A.7), we can characterize the boundary in the following parametric form (parametrized by z)

$$\lambda_2 = v'(z),$$
$$\lambda_1 = v(z) - zv'(z).$$

Moreover, we have $\frac{\mathrm{d}\lambda_1}{\mathrm{d}\lambda_2} = -z$. Let z^* denote the parameter value at the tangent point. Then, it satisfies $-z^* = -\mu^0$, which gives us the optimal solution to be $\lambda_1^* = v(\mu^0) - \mu^0 v'(\mu^0), \lambda_2^* = v'(\mu^0)$, and the optimal objective value is $v(\mu^0)$. □

F.4 Proof of Lemma 8

Proof. According to its definition in (13), $v^o(z)$ is monotonically increasing in z.

According to Lemma 6, we know $\tilde{v}(z) = \tilde{V}(z, T; z)$ solved from (25) is concave in z. Since $\tilde{v}(p/g) = \alpha\lambda pT \leq \tilde{v}(1) = \lambda pT$, we know $\tilde{v}(z)$ is monotone increasing in z. □

F.5 Proof of Proposition 5

Proof. The proof directly follows from Lemma 7 and (13). □

F.6 Proof of Theorem 2

Proof. Following from Lemma 6, all the possible cases of $\tilde{v}(\cdot)$ are illustrated in Figure 4. Following from Lemma 7, the concave closure of $\tilde{v}(\cdot)$ is the optimal value and can be depicted in Figure 8. Because $\tilde{v}(p/g) \leq v^\circ(p/g)$, we know the optimal μ^h in the social learning model is greater or equal to p/g. By Bayes' rule, we have $\pi^{\mathsf{SL}}_{\ell L} \geq \pi^{\mathsf{NSL}}_{\ell L}$, which means π^{SL} Blackwell dominates π^{NSL} by definition (see, for example, Blackwell and Girshick (1979) Chapter 12). □

Appendix G Proofs of Section 7

G.1 Proof of Theorem 3

Proof. (i) Suppose there exists a fully separating equilibrium in which the L type chooses π^L and the H type chooses $\pi^H \neq \pi^L$. Then, we have $\beta(\pi^H) = 1$ and $\beta(\pi^L) = 0$. To make the L type have no incentive to deviate to π^H, we must set $\pi^H = \pi^{\mathsf{FD}}$, where π^{FD} is a full disclosure information structure. However, the same outcome can be obtained by a pooling equilibrium by setting $\pi^L = \pi^H = \pi^{\mathsf{FD}}$.

(ii) The H type firm prefers to hire an influencer having a π^{FD} information structure because $v(1) \geq v(z)$ for any $z < 1$ in both no-social-learning and social-learning settings. In other words, the H type firm is always weakly better off by deviating to selecting a π^{FD} influencer. Thus, along with Theorem 3 (i), we conclude the outcome that followers know the true state is supported as an equilibrium. □

Appendix H Why use a dynamic social learning model?

One might be interested in the question of whether a simpler setting without dynamic beliefs updates (20) could have already conveyed the same messages captured in Figures 4 and 5 and Theorem 1. Indeed, suppose the additional information to followers arrives in a "single shot" instead of growing dynamically with staggered follower purchases. To this end, we present two static models such that only one piece of additional information is involved. As before, only skeptics learn from this additional information.

(a) **A static model with one external signal (Static Model 1).** In the first static model with additional information, we consider the setting that skeptics receive one additional signal generated from the distribution $\{\pi^E(s|w), s \in \{h, \ell\}, w \in \Omega\}$ and this distribution π^E is optimally chosen in the sender-preferred way. Following from the Bayesian persuasion literature (Kamenica and Gentzkow, 2011) (also proved in Lemma 7 in the next section), for a given prior z, the expected revenue $v(z)$ is

$$v(z) = \begin{cases} (1-\alpha)\frac{p\lambda T}{p/g} z & \text{for } z < p/g, \\ p\lambda T & \text{for } z \geq p/g. \end{cases} \tag{A.8}$$

We can visualize (A.8) together with that of no additional information model (13) in Figure 13 in Appendix H. The first part of (A.8) ($z < p/g$) comes from the fact that skeptics' belief will be changed by the additional signal generated from the optimally designed π^E, but devotees will not purchase because $z < p/g$. Therefore, the expected revenue will be positive. The second part of (A.8) ($z \geq p/g$) is because everyone will purchase when $z > p/g$. Thus, it is obvious that the interesting cases (b) (c) in Figure 4 will be lost, as well as (iii) in Theorem 1. In the Static Model 1, the influencer will never have any value if the prior belief is high (i.e., $\mu^0 > p/g$). The reason is that the market demand is fully deterministic when the prior belief is high, and there is no dynamic learning among followers. However, in our dynamic social learning model, the randomness in demand will not disappear even when the prior belief is high. This is because the noisy $\{R_{ht}, R_{\ell t}, t \geq 0\}$ realizations generate stochastic skeptics' beliefs, which again will feed back into market demand to generate future organic information.

(b) **A static model with one additional "organic" signal (Static Model 2).** In the second static model with additional information, instead of an external additional signal, we consider this additional signal to be generated "organically", meaning that this additional signal is generated from someone who purchased the product. Only skeptics will learn from this organic piece of information.

When $z < p/g$, no one will purchase, and no organic information can be released, so $v(z) = 0$ for $z < p/g$. When $z > p/g$, everyone will purchase, so the optimally designed organic information will choose not to change the beliefs of skeptics, which gives us $v(z) = p\lambda T$ for $z > p/g$. Therefore, this case collapses to our basic mode in Section 3.

Thus, from the analysis of both Static Models 1 and 2, it is the nonlinear structure of the expected revenue in post-signal prior that makes the problem interesting. The nonlinear part of

$v(z)$ comes precisely from the dynamic social learning aspect of our model: skeptics dynamically learn from the accumulated market organic information over time, which is triggered by the starting signal posted by the influencer and will dynamically feed back into demand generation, thus generates more organic information.

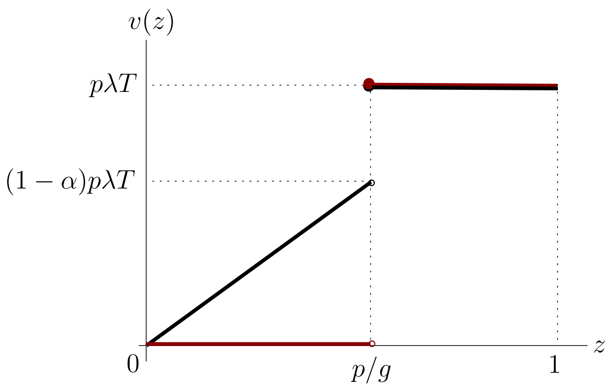

Figure 13: Black curve: $v(z)$ for the Static Model 1; Red curve: $v(z)$ for the basic model defined in Section 4.2.

Appendix I The micro-foundation of model for organic signals

In our model, we assume that buyers generate organic signals (reviews) at rate $\lambda\gamma f(\alpha, \mu_0, \mu_t)$. We assume that the informativeness of organic reviews, characterized by $\pi^R(h|H)$ and $\pi^R(\ell|H)$, depends only on the true state. This section provides a behavioral foundation for this specification and discusses why a seemingly natural alternatives are analytically intractable.

Micro-foundation for the review model. Our assumption that π^R is exogenous follows standard practice in the information economics literature (Banerjee, 1992; Bikhchandani et al., 1992) and admits a natural behavioral foundation. Suppose that review decisions are driven by experienced quality rather than purchase value, and that the experience utility is independent of the purchase utility defined in Table 1. Specifically, for a purchaser i, the experience utility is:

$$U_i^{\text{experience}} = w + \epsilon_i$$

where $w \in \{H, L\}$ is the true product quality and $\epsilon_i \sim F(\cdot)$ is i.i.d. across consumers and independent of their *ex-ante* purchase decisions. A purchaser generates a positive review h if $U_i^{\text{experience}} \geq \theta$ and a negative review ℓ otherwise. This yields:

$$\pi^R(h|H) \triangleq \mathbb{P}(\epsilon_i \geq \theta - H) = 1 - F(\theta - H),$$
$$\pi^R(h|L) \triangleq \mathbb{P}(\epsilon_i \geq \theta - L) = 1 - F(\theta - L),$$

with $\pi^R(\ell|H) = F(\theta - H)$ and $\pi^R(\ell|L) = F(\theta - L)$.

This two-component utility structure is behaviorally plausible: consumers decide whether to purchase based on their expected value (incorporating prior beliefs and personal preferences), but they decide whether to post a positive or negative review based on their actual experience with the product. The independence of ϵ_i from the purchase decision ensures that review informativeness is determined solely by product quality. Importantly, this specification captures the essential economic feature that high-quality products generate more positive reviews ($\pi^R(h|H) > \pi^R(h|L)$), while maintaining analytical tractability.

We next discuss alternative specifications and discuss their tractability.

What if γ depends on state (selection bias)? A natural extension would link the probability of generating organic reviews directly to followers' realized purchase utilities, creating state-dependent review informativeness through the purchase decision itself. Specifically, for a follower i who purchases the product, one might assume she generates a positive review h if her realized utility U_i (as defined in Table 1) exceeds a threshold θ, and a negative review ℓ otherwise. We now show why this specification, while intuitive, is analytically intractable.

Under this alternative, conditional on state H and given devotees' belief z and skeptics' belief μ_t, the probabilities that follower i releases a positive signal, a negative signal, or does not purchase are denoted by $\gamma_t(h|H)$, $\gamma_t(\ell|H)$, and $\gamma_t(N|H)$, respectively:

$$\begin{aligned}
\gamma_t(h|H) &\triangleq \mathbb{P}\left(U_i \geq \theta, i \text{ makes a purchase}|H\right)\\
&= \alpha \mathbf{1}_{g-p\geq\theta,-p+zg\geq 0} + (1-\alpha)\mathbf{1}_{g-p\geq\theta,-p+\mu_t g\geq 0}\\
\gamma_t(\ell|H) &\triangleq \mathbb{P}\left(U_i < \theta, i \text{ makes a purchase}|H\right)\\
&= \alpha \mathbf{1}_{g-p<\theta,-p+zg\geq 0} + (1-\alpha)\mathbf{1}_{g-p<\theta,-p+\mu_t g\geq 0}\\
\gamma_t(N|H) &\triangleq \mathbb{P}\left(i \text{ does not make a purchase}|H\right)\\
&= \alpha \mathbf{1}_{-p+zg<0} + (1-\alpha)\mathbf{1}_{-p+\mu_t g<0}.
\end{aligned}$$

Similarly, conditional on state L, the probabilities are:

$$\begin{aligned}
\gamma_t(h|L) &\triangleq \mathbb{P}\left(U_i \geq \theta, i \text{ makes a purchase}|L\right)\\
&= \alpha \mathbf{1}_{-p\geq\theta,-p+zg\geq 0} + (1-\alpha)\mathbf{1}_{-p\geq\theta,-p+\mu_t g\geq 0} = \mathbf{1}_{-p\geq\theta},\\
\gamma_t(\ell|L) &\triangleq \mathbb{P}\left(U_i < \theta, i \text{ makes a purchase}|L\right)\\
&= \alpha \mathbf{1}_{-p<\theta,-p+zg\geq 0} + (1-\alpha)\mathbf{1}_{-p<\theta,-p+\mu_t g\geq 0}\\
\gamma_t(N|L) &\triangleq \mathbb{P}\left(i \text{ does not make a purchase}|L\right)\\
&= \alpha \mathbf{1}_{-p+zg<0} + (1-\alpha)\mathbf{1}_{-p+\mu_t g<0}.
\end{aligned}$$

Under this specification, the organic signal processes $\{R_{ht}, t \geq 0\}$ and $\{R_{\ell t}, t \geq 0\}$ are nonhomogeneous Poisson processes whose rates at time t equal $\lambda\gamma_t(h|j)$ and $\lambda\gamma_t(\ell|j)$ for state $j \in \{H, L\}$. The likelihood of the history $\mathcal{H}_t := \{R_{hu}, R_{\ell u}, u \leq t\}$ can be expressed in terms of the inter-arrival times $\{t_{H1}, t_{H2}, \ldots, t_{HR_{ht}}\}$ and $\{t_{L1}, t_{L2}, \ldots, t_{LR_{\ell t}}\}$, where t_{Hk} denotes the time of the k^{th} positive signal and t_{Lk} denotes the time of the k^{th} negative signal. Given state $j \in \{H, L\}$, the likelihood at time t is:

$$\mathbb{P}\left(\mathcal{H}_t \mid j\right) = \exp\left(-\int_0^t \lambda f(\alpha, \mu_0, \mu_u) du\right) \cdot \prod_{k=1}^{R_{ht}} \left(\lambda\gamma_{t_{Hk}}(h|j)\right) \prod_{k=1}^{R_{\ell t}} \left(\lambda\gamma_{t_{Lk}}(\ell|j)\right). \tag{A.9}$$

Bayes' rule can be applied to update beliefs. Specifically, immediately after a signal s arrives at time t, we have the following update:

$$\mu_{t+} = \frac{\mu_t\gamma_t(s|H)}{\mu_t\gamma_t(s|H) + (1-\mu_t)\gamma_t(s|L)},$$

We can understand the analytical challenge of the alternative specification by comparing the Bayesian update under the original model:

$$\mu_{t+} = \frac{\mu_t \pi^R(s|H)}{\mu_t \pi^R(s|H) + (1-\mu_t)\pi^R(s|L)},$$

The challenge arises when deriving a stochastic differentiation equation for $d\mu_t$ using Itô's lemma. The functions $\gamma_t(s|H)$ and $\gamma_t(s|L)$ are discontinuous in the posterior belief μ_t due to the indicator functions. When μ_t crosses the purchase threshold p/g, skeptics begin purchasing, which discretely changes the informativeness of reviews. Specifically, the likelihood ratio:

$$\frac{\gamma_t(s|H)}{\gamma_t(s|L)}$$

depends on μ_t itself, creating a complex endogenous feedback: the informativeness of signals depends on beliefs, which in turn are updated based on those signals.

This differs fundamentally from our original model, where the likelihood ratio $\frac{\pi^R(s|H)}{\pi^R(s|L)}$ is constant. The original model maintains constant signal informativeness, which enables the application of Itô's lemma on a jump process (see the proof of Lemma 3). In contrast, the discontinuous dependence of $\frac{\gamma_t(s|H)}{\gamma_t(s|L)}$ on μ_t prevents the application of Itô's lemma, which makes the analysis intractable.

What if non-buyers can post reviews? Following current practice on most shopping platforms where only verified buyers can post reviews, we assume only buyers can post. However, we now examine how our results extend if non-buyers can also post negative reviews.

Suppose non-buyers post negative reviews at the same propensity as buyers, characterized by the same state-dependent probability $\pi^R(\ell|w)$. At time t in state w, the arrival rates are:

- Positive reviews (only from buyers): $\lambda\gamma f(\alpha, \mu_0, \mu_t)\pi^R(h|w)$
- Negative reviews (from both buyers and non-buyers):

$$\lambda\gamma f(\alpha, \mu_0, \mu_t)\pi^R(\ell|w) + \lambda(1 - f(\alpha, \mu_0, \mu_t))\pi^R(\ell|w) = \lambda\pi^R(\ell|w)$$

Note that the negative review rate becomes independent of $f(\alpha, \mu_0, \mu_t)$ when both buyers and non-buyers post at the same propensity. The likelihood of the review history $\mathcal{H}_t := \{R_{hu}, R_{\ell u}, u \leq t\}$ in state w is:

$$\mathbb{P}(\mathcal{H}_t \mid w) = \exp\left(-\int_0^t \left[\lambda\gamma f(\alpha, \mu_0, \mu_u)\pi^R(h|w) + \lambda\pi^R(\ell|w)\right] du\right)$$
$$\times \prod_{i=1}^{R_{ht}} \left[\lambda\gamma f(\alpha, \mu_0, \mu_{t_{hi}})\pi^R(h|w)\right] \prod_{j=1}^{R_{\ell t}} \left[\lambda\pi^R(\ell|w)\right]$$

where t_{hi} and $t_{\ell j}$ denote the arrival times of the i-th positive and j-th negative review, respectively.

When computing the likelihood ratio for Bayesian updating, observe that:

$$\begin{aligned}\frac{\mathbb{P}(\mathcal{H}_t \mid H)}{\mathbb{P}(\mathcal{H}_t \mid L)} &= \exp\left(-\int_0^t \lambda\gamma f(\alpha,\mu_0,\mu_u)\left[\pi^R(h|H)-\pi^R(h|L)\right]du\right)\\ &\times \exp\left(-\int_0^t \lambda\left[\pi^R(\ell|H)-\pi^R(\ell|L)\right]du\right)\\ &\times \left(\frac{\pi^R(h|H)}{\pi^R(h|L)}\right)^{R_{ht}}\left(\frac{\pi^R(\ell|H)}{\pi^R(\ell|L)}\right)^{R_{\ell t}}\end{aligned}$$

The key observation is that the exponential term depends on the cumulative purchase behavior over time but not on the specific timing of individual reviews. Since negative reviews arrive at a constant rate independent of purchases, they provide the same belief update as in the original model. Positive reviews, which only come from buyers, also maintain their original informativeness. Therefore, the belief dynamics characterized in Lemmas 2 and 3 remain unchanged, and all our main results continue to hold.

When tractability breaks down. If buyers and non-buyers have different propensities to post negative reviews, say $\pi^R_{\text{buyer}}(\ell|w) \neq \pi^R_{\text{non-buyer}}(\ell|w)$, then the negative review rate becomes:

$$\lambda\gamma f(\alpha,\mu_0,\mu_t)\pi^R_{\text{buyer}}(\ell|w) + \lambda(1-f(\alpha,\mu_0,\mu_t))\pi^R_{\text{non-buyer}}(\ell|w)$$

In this case, the state-dependent terms cannot be factored out, and the entire timing history $\{t_{\ell j}, j = 1,\ldots,R_{\ell t}\}$ matters for belief updating. This path-dependence makes the problem analytically intractable, similar to the selection bias issue discussed in the previous discussion.

What if λ is state-dependent? We now consider the case where the arrival intensity λ depends on the true state. Let λ_w denote the arrival rate in state $w \in \{H, L\}$, where we naturally assume $\lambda_H > \lambda_L$ (high-quality products attract more customer traffic). In this setting, purchase counts themselves become informative about product quality.

Let $\{D_u, u \le t\}$ denote the cumulative purchase count process. Given state w, the number of purchases D_t by time t follows a Poisson distribution with mean $\int_0^t \lambda_w f(\alpha,\mu_0,\mu_u)\mathrm{d}u$, where $f(\alpha,\mu_0,\mu_t)$ is the purchase probability. The likelihood is:

$$\mathbb{P}\left(D_t \mid w\right) \propto \lambda_w^{D_t}\exp\left(-\lambda_w\int_0^t f(\alpha,\mu_0,\mu_u)\mathrm{d}u\right) \quad \text{for } w\in\{H,L\}.$$

Applying Bayes' rule, the posterior belief of skeptics at time t is:

$$\mu_t = \frac{\mu_0}{\mu_0 + (1-\mu_0)\left(\frac{\lambda_L}{\lambda_H}\right)^{D_t}\exp\left((\lambda_H-\lambda_L)\int_0^t f(\alpha,\mu_0,\mu_u)\mathrm{d}u\right)}. \tag{A.10}$$

This formula is intuitive: more purchases D_t increase belief in state H (since $\lambda_H < \lambda_L$). Conversely, if time passes without purchases, the exponential term grows, diminishing belief in state H. Hence, observing no arrivals suggests the product may not be a good match.

Applying Itô's lemma to (A.10), we derive the stochastic differential equation for posterior beliefs:

$$\mathrm{d}\mu_t = -\mu_{t-}(1-\mu_{t-})(\lambda_H-\lambda_L)f(\alpha,\mu_0,\mu_{t-})\mathrm{d}t + \frac{\mu_{t-}(1-\mu_{t-})(\lambda_H-\lambda_L)}{\lambda_H\mu_{t-}+\lambda_L(1-\mu_{t-})}\mathrm{d}D_t. \tag{A.11}$$

Equation (A.11) differs from our original model (Lemma 3) in two key ways:

1. **Drift term**: The term $-\mu_{t-}(1-\mu_{t-})(\lambda_H-\lambda_L)f(\alpha,\mu_0,\mu_{t-})\mathrm{d}t$ represents belief degradation over time when no purchases occur. This serves a similar function to negative reviews in our baseline model. It pulls beliefs downward.
2. **Information structure**: Purchase events themselves (rather than reviews) carry information. The jump size $\frac{\mu_{t-}(1-\mu_{t-})(\lambda_H-\lambda_L)}{\lambda_H\mu_{t-}+\lambda_L(1-\mu_{t-})}$ is always positive, meaning every purchase increases belief in state H.

Despite these differences, our main economic insights remain intact. The belief process (20) is still a martingale (the drift and diffusion terms balance in expectation), so the variance of belief evolution captured in Lemma 5 continues to determine how social learning affects the firm's revenue. Hence, the firm faces a similar tradeoff as before.

The presence of the drift term introduces an additional consideration: even without negative reviews, beliefs naturally decay when purchases are sparse. However, the fundamental optimization problem remains tractable due to the Markovian structure of the belief process in (20). A complete characterization would require resolving the optimization problem with these modified belief dynamics.

Appendix J An algorithm to find the optimal influencer

According to Lemma 7, we provide the following algorithm to solve the problem (27). In the algorithm, we find the indices z_1 and z_2 of the two active constraints in (28) that characterize the optimal solution. To find the indices of these active constraints, we do a simple search under the assumption that the indices lie within a given grid of precision parameterized by n.

Recall that the optimal value of (27) is the height of the shaded rectangle in Figure 10. The following simple algorithm can efficiently find the optimal solution when restricted to a chosen grid. The input to the algorithm is the value of v at discretized points in the post-signal belief space. That is, we discretize $[0,1]$ to $z_0=0, z_1=1/n, z_2=2/n,\ldots,z_n=1$ and calculate $v(z_k)$ for all $k=0,1,\ldots,n$.

Algorithm 1 Finding the indices of active constraints in Lemma 7.

procedure OptimalActiveIndices($\bar{v}(z_k), k=0,1,\ldots,n$)
  $i_0 \leftarrow 0, z_1 \leftarrow z_{i_0}, z_2 \leftarrow z_{i_0}$ ▷ initialize z_1, z_2
  while $z_{i_0} < \mu_0$ **do**
    $i \leftarrow \arg\max_{j\geq i_0} \frac{v(z_j)-v(z_{i_0})}{z_j-z_{i_0}}$
    if $i == i_0+1$ **then**
      $z_1 \leftarrow z_i, z_2 \leftarrow z_i$
      $i_0 \leftarrow i$
    else
      if $z_i \geq \mu_0$ **then**
        $z_2 \leftarrow z_i$ **break**
      else
        $z_1 \leftarrow z_i, z_2 \leftarrow z_i$
        $i_0 \leftarrow i$

return z_1, z_2

Proposition A.1. Under the assumption that the indices of the active constraints lie in a grid of precision n, Algorithm 1 gives the indices of the active constraints at the optimal solution to problem (28).

Proof. According to Lemma 7, the optimal value of the objective function is either (30) or $\bar{v}(\mu^0)$. Thus, for a given z_1, (30) is maximized at $z_2^* = \arg\max_{z_2 \geq z_1} \frac{\bar{v}(z_2)-\bar{v}(z_1)}{z_2-z_1}$. Moreover, because $\bar{v}(\cdot)$ is a monotone increasing function, for any $z_1' \geq z_1$, we have

$$\arg\max_{z_2 \geq z_1} \frac{\bar{v}(z_2)-\bar{v}(z_1)}{z_2-z_1} \geq \arg\max_{z_2 \geq z_1'} \frac{\bar{v}(z_2)-\bar{v}(z_1')}{z_2-z_1'}.$$

Therefore, Algorithm 1 starts from $z_1 = 0$, and search for z_2 to maximize $\frac{\bar{v}(z_2)-\bar{v}(z_1)}{z_2-z_1}$. Gradually increasing z_1, once it finds a pair of (z_1, z_2) that satisfies $z_1 \leq \mu^0(H) \leq z_2$, this pair is the optimal indices (because increasing z_1 decreases $\max_{z_2 \geq z_1} \frac{\bar{v}(z_2)-\bar{v}(z_1)}{z_2-z_1}$ and decreases $\arg\max_{z_2 \geq z_1} \frac{\bar{v}(z_2)-\bar{v}(z_1)}{z_2-z_1}$, as a result, decreases the value (30)).

If the algorithm cannot find such a pair, it returns the index at μ^0. □

Appendix K Monte Carlo validation of diffusion approximation

Our key approximation is to approximate the belief jump process by its diffusion limit. In the simulation, we compare the original belief jump process μ_t with its diffusion limit $\tilde{\mu}_t$ and show the following figures. To illustrate the approximation (Lemma 5), Figures 14 and 16 is an example of the sample belief path, and Figures 15 and 17 are the normalized error over time ($\frac{|\mu_t-\tilde{\mu}_t|}{\sqrt{t}}$), and Figure 18 is the normalized total error ($\frac{\sum_{t=0}^{T} |\mu_t^{(n)}-\tilde{\mu}_t|}{\sqrt{n}}$) when λ scales up at the rate of n. All expectations are calculated based on 200 sample paths.

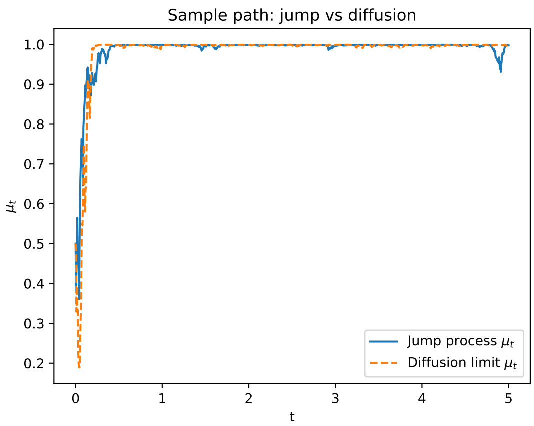

Figure 14: Sample path ($\lambda = 500$)

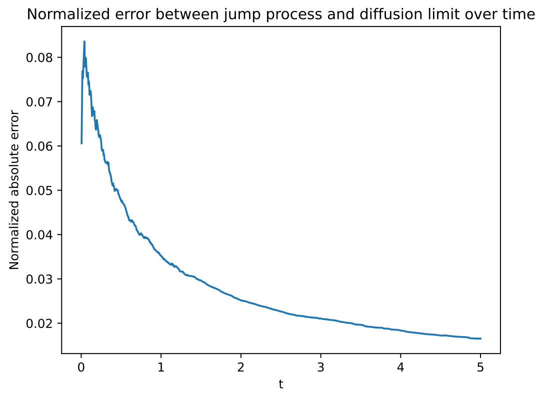

Figure 15: Normalized error (based on 200 sample paths, $\lambda = 500$)

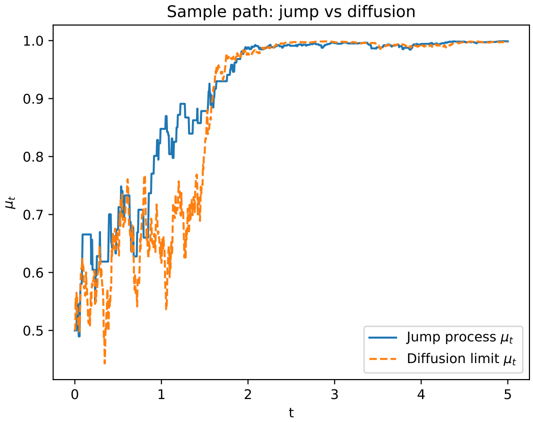

Figure 16: Sample path ($\lambda = 50$)

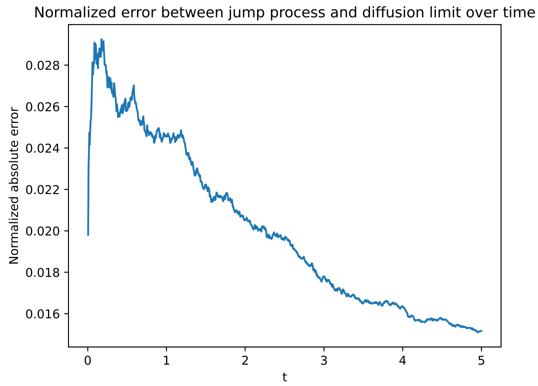

Figure 17: Normalized error (based on 200 sample paths, $\lambda = 50$)

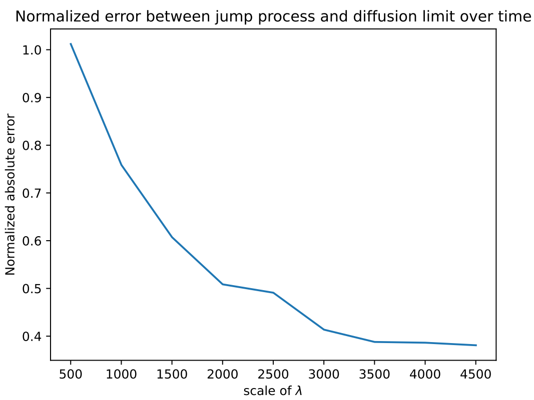

Figure 18: Normalized total absolute error during $[0, T]$