Strategic Conformity in Affiliate Marketing

Itay Fainmesser* Xudong Zheng†

February 29, 2024

Abstract

Affiliate marketing is a prevalent performance-based marketing model. For the most part, retailers have formally transitioned from paying marketing affiliates pay-per-click to a pay-per-purchase compensation model. Yet, some pay-per-click incentives remain in many affiliate programs. Our results show that the remaining pay-per-click incentives give rise to a conflict of interest between affiliates and consumers. On the one hand, consumer surplus is maximized under affiliates' *recommendation conformity*, i.e., when affiliates converge on a single product in their recommendation profile. On the other hand, if affiliates are compensated significantly per click, they achieve the highest aggregate payoffs by minimizing recommendation conformity. Surprisingly, even if affiliates are compensated mostly per purchase and both consumers and affiliates would benefit from recommendation conformity, a *conformity equilibrium* may not exist. In contrast, a non-conformity equilibrium always exists. We characterize further how consumers' expected search length, affiliates' product information accuracy, and consumers' ability to learn directly about products' qualities affect market outcomes via its effect on the existence of conformity equilibria. Finally, our results shed light on how a retailer can craft affiliates' compensation structure to influence market outcomes.

Keywords: Affiliate marketing, Strategic conformity

*The Johns Hopkins Carey Business School and The Department of Economics, The Johns Hopkins University, Baltimore, MD 21202 (e-mail: itay_fainmesser@jhu.edu)

†The Department of Economics, The Johns Hopkins University, Baltimore, MD 21218 (e-mail: xzheng27@jhu.edu)

1 Introduction

Affiliates are websites (e.g., Wirecutter, CNET), bloggers, and social media influencers, who promote a retailer's product or service (Shopify 2024) using unique *affiliate links*—links that allow tracking by the retailer. An affiliate directs a consumer to a product's web page on the retailer's online platform, and is compensated based on the consumer's action. Common compensation models include pay-per-click (the retailer pays the affiliate for each click on an affiliate link by a consumer) and pay-per-purchase (the retailer pays the affiliate for each purchase by a consumer who is directed to the platform by the affiliate's link). Affiliate marketing is becoming increasingly important to consumers and retailers. Since its inception in the mid-90s, affiliate marketing has experienced a 10% year-on-year growth, becoming an industry worth over 15.7 billion U.S. dollars in 2023,[1] accounting for 16% of all internet orders globally (a staggering 53% of all sales on Amazon came from affiliates).[2] As a result, it is not surprising that over 80 percent of retailers consider affiliate marketing to be a key way to promote products and engage consumers (Jefferson 2016). Yet despite the prominence of affiliate marketing in practice, scant research has focused on affiliates' strategic product recommendations.

This paper develops a new model that captures three key features of affiliate marketing: (1) the majority of pay-per-purchase affiliate programs use a "cookie duration" to provide affiliates with pay-per-click-like incentives,[3] (2) consumers often contribute more clicks than purchases during their search processes, and (3) consumers and affiliates may be more likely to be aware of some prominent products ex ante (e.g., a running shoe from Nike) than of other less prominent products (e.g., a running shoe from Hoka)—we refer to such prominent product as *salient products*, and to affiliates' tendency to converge in their recommendations and recommend the same salient product

[1]Source: Printify, `https://printify.com/blog/affiliate-marketing-trends/`, accessed January 25, 2024.

[2]Source: Demandsage, `https://www.demandsage.com/affiliate-marketing-statistics/`, accessed January 25, 2024.

[3]That is, once a consumer clicks on an affiliate link, a cookie is set, and the affiliate receives a commission for the consumer's first purchase on website within a specific time frame, typically from twenty-four hours to several months. Notably, the affiliate receives a commission even if the consumer buys a different product than the one recommended.

as *recommendation conformity*.[4,5]

We answer the following questions:

1. How is recommendation conformity related to other market outcomes, such as affiliates' aggregate profits and consumer surplus?

2. How does the affiliate program's commission structure, in particular pay-per-purchase vs. pay-per-click, affect affiliates' recommendation conformity, and subsequently the related market outcomes?

3. Does increasing the information that affiliates have on products' qualities (e.g., by allowing affiliates to test products) make affiliates' recommendation conformity more or less likely?

4. How do features of consumers' search, such as search duration and the ability to observe a product's quality, affect affiliates' recommendation conformity?

Our analysis shows that two affiliates' recommendation profiles are especially important when it comes to welfare outcomes. In a *non-conformity strategy profile*, affiliates disregard the salient product, and each recommends a non-salient product for which he received a high quality signal. On the other extreme, in a *conformity strategy profile*, every affiliate who receives a high quality signal for the salient product, recommends the salient product. A *non-conformity equilibrium* and a *conformity equilibrium* are equilibria in which affiliates employ non-conformity and conformity strategy profiles respectively.

Our first result shows that a conformity strategy profile always maximizes the (expected) Consumer Surplus (CS). If affiliates' compensation has a significant per-purchase component, then a conformity strategy profile also maximizes affiliates' aggregate expected payoff, or Affiliate Sur-

[4]Our notion of conformity is distinct from the one studied in the (sequential) social learning literature on online reviews (Moe and Trusov 2011, Moe and Schweidel 2012, Sridhar and Srinivasan 2012). In that literature reviewers may observe earlier reviews and, through belief updating, have their opinions (and reviews) align with existing reviews. In contrast, affiliates in our framework provide their recommendations simultaneously. That is, their recommendation conformity is a strategic outcome based on affiliates' focus on maximizing their expected earnings given their compensation structure.

[5]In our model, there are many ex-ante substitutable products, each of either high or low quality. One of the products is "salient" in the following sense: Each affiliate observes signals that contains information on the quality of the salient product and the qualities of a randomly selected sample of the non-salient products. That is, all affiliates receive signals on the quality of the salient product, and it is highly unlikely that any two affiliates receive signals on the quality of the same non-salient product. As a result, recommendation conformity on a non-salient product is ruled-out.

plus (AS).[6] In that case, consumers and affiliates have aligned preferences for conformity. However, if affiliates' compensation has a significant per-click component, then AS is maximized by a non-conformity strategy profile.[7] In that case, consumers' and affiliates' preferences for payoff maximization are polar opposites. That is, we provide a bang-bang result: either AS and CS are maximized by the same recommendation strategy profile or by opposing ones.

To see why a conformity strategy profile always maximizes CS, note that a conformity strategy profile aggregates the most accurate information about the quality of the salient product. If the salient product receives many recommendations, it is very likely to be of high quality. The consumer's belief induces her to inspect it first and likely to make a purchase, thus ending her costly search after one click. The relationship between conformity and AS is more subtle. On the one hand, when affiliates are compensated heavily per purchase, they want to convert the consumer's clicks to purchases effectively, so that the consumer purchases with high probability before stopping to search. On the other hand, when affiliates are compensated heavily per click, they seek to maximize clicks, hence, they might prefer to risk the consumer stopping the search without purchasing, if that means the consumer clicks on more links in expectations. The affiliates' compensation threshold for preferring conformity over non-conformity depends on the consumer's expected search duration (in case of no purchase) and the accuracy of affiliates' information on product quality: (a) the shorter the expected search the better off affiliates are in focusing on a conformity strategy profile that maximizes the probability of purchase, and (b) the more accurate the affiliates' information, the more certain they are that the search will end with a purchase, so their focus shifts to maximizing clicks.

Our second and third results show that a non-conformity equilibrium always exists, whereas the existence of a conformity equilibrium subtly depends on affiliates' compensation. In fact, even when both the consumer and the affiliates would benefit from affiliates' conformity, a conformity equilibrium may not exist, leading to a sub-optimal welfare outcome. Yet, when affiliates' payoffs are maximized with a non-conformity strategy profile, a conformity equilibrium can exist. This suggests that conflict of interest built into the market structure may make it challenging for affiliates to coordinate on an equilibrium that is good for them and they may, instead, end with a market

[6] The same is true if the consumer is likely to abort the search quickly in the absence of a purchase, or if affiliates' information about products' qualities is likely inaccurate.

[7] The same is true if the consumer is likely to continue searching for longer, or if affiliates' information about products' qualities is sufficiently accurate.

outcome that is sub-optimal for them.

To see why a non-conformity equilibrium always exists, suppose that all but one affiliate recommend non-salient products. In that case, the remaining affiliate will have no incentive to recommend the salient product (regardless of his information about the salient product's quality). The reason is that salience leads the consumer to make quality inferences about the salient product based on the absence of recommendations—if some affiliates do not recommend it, they might have learned that the salient product has low quality. To the contrary, a conformity equilibrium may not exists if either the incentives for affiliates to conform are two weak (for example if compensation is heavily per click), or if their incentives to conform are too strong (for example if compensation is heavily per purchase). In the case in which incentives to conform are too strong, if all other affiliates employ conformity strategies, an affiliate would chose to recommend the salient product even if he received a low quality signal for it—leading to the collapse of the equilibrium and unraveling towards the non-conformity equilibrium.

An interesting implication of our results is that the retailer can affect when a conformity equilibrium exists through its control of affiliates' compensation, changing the weight between per-click and per-purchase compensation. However, the effect of shifting compensation weights between per-purchase and per-click is not straightforward and depends on the consumer's expected search duration in the following way: If a consumer is likely to abandon the search early, her first click-through is more critical to affiliates than the prospect of subsequent clicks. Hence, when affiliates are paid per-click, a conformity equilibrium exists for a larger set of parameters. In contrast, when the consumer's expected search is longer, affiliates are less sensitive to the first click. As a result, if consumer search is expected to be longer, a heavier pay-per-click component implies that a conformity equilibrium exists for a smaller set of parameters.

We generalize our model to accommodate an environment in which the consumer cannot perfectly observe the true quality of a product after clicking on an affiliate link, but instead she observes a noisy signal. We show that our results extend to this more general environment. The more general model also allows us to derive additional comparative statics. Interestingly, in the limit case of experience goods, where consumers do not observe any independent quality signal before purchase, affiliates as a group become indifferent between conformity and non-conformity strategy profiles, and a conformity equilibrium exists for a smaller range of parameters.

Our paper continues as follows. We discuss our contribution to the literature in the remaining of Section 1. Section 2 presents our baseline model. In Section 3, we present our results highlighting the conflict of interest between affiliates and consumers, and characterize the existence of conformity equilibria. Section 4 presents an extension generalizing the assumption on the consumer's learning and discusses implications for platform design. All proofs are deferred to the appendix.

1.1 Related Literature

Our paper contributes to the literature on social influencers. Katona (2020) analyzes how firms compete for influencers considering the overlap and degrees in the influencers' reach in a social network. Fainmesser and Galeotti (2021) study a market in which influencers maximize sponsorship revenues by choosing how much sponsored and organic content to post. Jain and Qian (2021) analyze how competition among online content producers, the size of the consumer base, and consumers' time constraints affect retailers' revenue-sharing strategies. Pei and Mayzlin (2022) study the optimal business relationship between a firm and influencers. Janssen and Williams (2021) study the role of influencers in affecting the order of consumer search. Berman et al. (2023) investigate how social influencers' creative contribution to a firm's marketing campaign affects the firm's profit by influencing the dispersion of the market demand for its product. Nistor and Selove (2022) focus on how an influencer's entertainment level affects the informativeness of their comment section and which products they endorse. This literature focuses on influencers who are paid per content endorsement (akin to how most advertisers are paid). Our main departure from this literature is that we study a performance-based marketing model, affiliate marketing, in which influencers are paid per performance (per click and/or per purchase).

There exists some empirical literature on affiliate marketing.[8] Papatla and Bhatnagar (2002) demonstrates the effect and advantage for an online retailer from having affiliates, even if such affiliates are also retailers of related products. Olbrich et al. (2019) find how Search Engine Optimization (SEO) can cannibalize the effect of affiliate marketing when merchants run multi-channel campaigns. Additional work focuses on various practical issues related to affiliate marketing, such as control mechanisms and contract designs (Gilliland and Rudd 2013), affiliate fraud (Edelman

[8]Early case studies include Hoffman and Novak (2000) on one of the earliest and most successful affiliate program, the BuyWeb Network of CD Now, a music retailer, and how it integrated affiliate marketing into its multifaceted customer acquisition strategy, and Duffy (2005), identifying a key to successful affiliate marketing, i.e., the construction of a win-win relationship between the advertiser and the affiliate.

and Brandi 2015), absence of endorsement disclosure (Mathur et al. 2018), deceptive and fabricated reviews (Karabas et al. 2021), and effects of language features on consumer engagement (Syrdal et al. 2023).

To our knowledge, there are two papers studying affiliate marketing theoretically. Libai et al. (2003) consider non-strategic affiliates and compares pay-per-purchase and pay-per-click from a retailer's perspective. Suryanarayana et al. (2019) study a model in which affiliates arrive sequentially, and analyze a retailer's optimal information disclosure to affiliates as they arrive. In this paper there is only one product and the retailer maximizes the revenue from selling that product.

2 Model

Consider a market with many products, one of which is more salient than the others. Affiliates (he/they) receive stochastic signals on the qualities of the products and each affiliate chooses one product to recommend online by providing an affiliate link. A consumer (she) observes the recommendations of the affiliates and decides whether to click on affiliate links as part of her costly search for a high quality product. We next describe each of the components of our model: products, the affiliates' information sets, the consumer, and the timeline of the model.

2.1 Products

A single retailer (e.g., Amazon) simultaneously offers a unit measure of distinct and substitutable products online. Ex-ante, products are divided into two types. There is one commonly known salient product, hereafter product 0.[9] The remaining products are non-salient. The formal difference between the salient and non-salient products will become clear in Subsection 2.2. We denote the set of non-salient products by NS.

Products are vertically differentiated. Each product $i \in \{0\} \cup NS$, regardless of its salience, has ex ante equal probability of being high or low quality, denoted by $q_i \in \{H, L\}$. That is, $Pr(q_i = H) = Pr(q_i = L) = 1/2$. The realizations of products' qualities are drawn independently. The prices of the products are all equal and normalized to 1.

[9] We leave unmodeled the reasons leading to product 0 being salient, such as brand market share, past marketing campaigns, or superior brand awareness.

2.2 Affiliates' Information

There are N affiliates, who actively post recommendations for the retailer's products (each on his channel, e.g., TikTok, YouTube, or a blog).[10] An affiliate j does not observe the qualities of the products but receives a signal $s_j(i) \in \{h, l\}$ for each product i from a set $\eta_j \subset \{0\} \cup NS$ that includes the salient product (product 0) and a countably infinite collection of products chosen uniformly at random (u.a.r.) from the set of non-salient products, NS.[11] The signal $s_j(i)$ is generated such that $Pr(s_j(i) = h | q_i = H) = Pr(s_j(i) = l | q_i = L) = p_0 > 1/2$ for all $j = 1, \ldots, N$ and $i \in \eta_j$. Signals are conditionally independent across affiliates and products. Taking into account his signals, each affiliate j recommends a single product online by posting an affiliate link. Such an affiliate link can direct the consumer to the web page of the recommended product on the retailer's online marketplace.

Each affiliate maximizes his monetary compensation, which has two components:

1. Per-click compensation: a commission of $\lambda \in (0, 1)$ if the consumer clicks on the link posted by the affiliate.

2. Per-purchase compensation: a commission of $1 - \lambda$ if the consumer clicks on the link posted by the affiliate *and* buys the recommended product.

That is, an affiliate's payoff is $\pi_j = \mathbb{1}_{click} \cdot \lambda + \mathbb{1}_{purchase} \cdot (1 - \lambda)$ where $\mathbb{1}_{click}$ and $\mathbb{1}_{purchase}$ are the indicator functions for whether the consumer clicks the affiliate's link and purchases the product recommended by the affiliate respectively.[12]

2.3 The consumer

There is one consumer with unit demand who is aware of all of the products. The consumer observes the product recommendations made by the affiliates, and engages in a sequential search: she can click on a link to a product (either an affiliate link provided by an affiliate, or a link she finds by searching a product online), evaluate and perfectly observe the product's quality, and make

[10]The set of N affiliates is presumably much smaller than the set of all affiliates active in the market and represents the affiliates the consumer considers. These affiliates could be the consumer's selection among the top results from her search for the products on a search engine.

[11]This approximates the case in which each affiliate receives signals for a large number of products, which is nevertheless small relative to the set of all products.

[12]The commission amount in practice is typically a fraction of the price of the product. The retailer only revises the fraction infrequently, if at all, and hence we keep the fraction given and fixed in the model.

a purchasing decision. If the consumer does not purchase the first product she evaluates, with (exogenous) probability $q \in [0, 1]$ she can click on a second link to another product and repeat the process.[13] For simplicity, we assume the consumer never clicks on a third link.[14] We abstract from product pricing and vertical differentiation and assume that the consumer has a net utility of 1 if she purchases a high quality product and -1 if she purchases a low quality product.

Evaluating a product's quality involves reviewing product specifications and reading other consumers' reviews. We capture the effort invested in learning about the product's quality by a cost c the consumer incurs each time she learns about a new product. Therefore, if the consumer searched $n \in \{1, 2\}$ products before finding a high quality product, her payoff is $1 - nc$. We focus attention on small enough c so that the consumer finds it profitable (in expectations) to click on a link to a product she believes to be of high quality with probability p_0 or higher.[15]

2.4 Timeline and equilibrium

To sum, the timeline of our model is as follows:

1. Nature determines products qualities.

2. Each affiliate j receives independent signals about the qualities of products in η_j and chooses one product to recommend. Affiliates make their recommendations simultaneously.

3. The consumer observes the product recommendations made by the affiliates and searches for a high quality product as described above.

A strategy for affiliate j is a map $\{h, l\}^{\eta_j} \to \Delta\eta_j$ from quality signals for products in the set η_j to a (potentially stochastic) recommendation. A strategy for the consumer is a map $(\{0\} \cup NS)^{\{1,2,\ldots,N\}} \to \Delta((\{0\} \cup NS \cup \varnothing) \times (\{0\} \cup NS \cup \varnothing))$ from the vector of the recommended products by all of the affiliates to a (potentially stochastic) choice of at most two products, one that she clicks on first, and another that she clicks on with probability q if she didn't purchase the product she clicked on first.

[13] There are many potential reasons for the consumer to stop her search: outside distractions, fatigue, or other time and mental constraints.

[14] The possibility of a second click is sufficient to capture the effect of consumer's search length on affiliates' recommendation decisions and conformity incentive, while maintaining the tractability of the model.

[15] In fact, if c is sufficiently large to deter the consumer from clicking on a link for a product that is of high quality with probability p_0, all equilibria of the game are such that no search takes place and no link is ever clicked.

A Perfect Bayesian Equilibrium is a collection of affiliates and consumer strategies and beliefs such that each affiliate's strategy maximizes his expected profit, the consumer's strategy maximizes her expected utility, and all beliefs are consistent with Bayes law whenever possible. In what follows, whenever we refer to an equilibrium, we mean a Perfect Bayesian Equilibrium.

2.5 Discussion of modeling assumptions

2.5.1 The commission model: pay-per-click vs. pay-per-purchase

As mentioned in the introduction, a minority of retailers and retail platforms still follow the outright pay-per-click affiliate commission model. However, most, if not all, retailers who follow a pay-per-purchase affiliate commission model pay an affiliate on any purchase made by consumers within a relatively large time frame (the "cookie duration"), regardless of whether the consumers purchased the products recommended by the affiliate. This practice provides affiliates with pay-per-click-like incentives. That is, an affiliate has at least a component of his expected payoff independent of whether the product he recommends is purchased.

In our baseline model, we abstract from the specific implementation of the compensation offered by the retailer and capture the pay-per-click-like component by λ. One can interpret a retailer's commission model with a high λ as the retailer offering its affiliates a longer cookie duration, providing affiliates with greater flexibility of receiving a commission beyond a sale of a linked product. In Section 4.2, we formulate a model that explicitly captures the pay-per-purchase-with-time-frame environment and shows how it is nested within our current model. We then apply this formalization to analyze the retailer's choice of affiliates' compensation model.

2.5.2 Single salient product

For tractability, we assume throughout that there is one salient product. While we expect many of our results to extend to an environment with multiple salient products, such an environment may also introduce the possibility of additional forms of conformity.[16] We leave the investigation of other conformity equilibria in a richer model for future research.

[16]For example, one strong form of conformity will be one in which affiliates' establish an ordering of the salient products, such that all affiliates who receive a high signal for product 'a' recommend it, all affiliates who receive a low signal for product 'a' and a high signal for product 'b' recommend product 'b', and so forth. The sustainability of some of these conformity strategy profiles in equilibrium might require the model to allow the consumer's longer search duration accordingly.

2.5.3 Informed consumer

We assume that once the consumer evaluates a product, she will perfectly learn the product's quality. We relax this assumption and explore the effect of the consumer's ability to self-evaluate products in Section 4.1. We show that our results extend qualitatively to this more general case and derive new insights from emerging quantitative changes.

3 Results: conflict of interest and affiliates' conformity

Due to the nature of the coordination game between affiliates and the asymmetric information between affiliates and the consumer, our model admits a large set of equilibria, including some less interesting (or realistic) ones in which some (or all) products are never recommended and never purchased with corresponding ad-hoc beliefs over zero probability events. Therefore, instead of characterizing the entire set of equilibria, we first show that two affiliates' strategy profiles stand out because, for any set of parameters, one of the two strategy profiles always maximizes affiliates' and/or the consumer's payoffs. These two strategy profiles are polar opposites, and their optimality highlights an inherent conflict of interest in this market. Then, we characterize the range of market parameters for which each of the two strategy profiles is an equilibrium, highlighting important trade-offs and inefficiencies.

3.1 Conflict of interest

Our first result characterizes the inherent conflict of interest between the affiliates as a group and the consumer with respect to recommendations of salient versus non-salient products and how it depends on the length of consumer search and on affiliates' compensation. Two affiliates' strategy profiles emerge as important in our analysis.

Definition 1. A *conformity strategy profile* is a strategy profile in which every affiliate who receives a high signal for the salient product recommends it and every affiliate who receives a low signal for the salient product recommends a non-salient product for which he received a high signal. A *conformity equilibrium* is an equilibrium in which affiliates employ a conformity strategy profile.

A *non-conformity strategy profile* is a strategy profile in which, regardless of their signals for the salient product, every affiliate recommends a non-salient product for which he received a high signal. A *non-conformity equilibrium* is an equilibrium in which affiliates employ a non-conformity strategy profile.

Proposition 1. *Suppose the consumer best responds given the affiliates' strategy profile (and some beliefs that are consistent with the Bayes's rule given the affiliates' strategy profile).*[17] *Then,*

(1) a conformity strategy profile maximizes the consumer's expected payoff among all affiliates' strategy profiles;

(2) If $\lambda \geq \frac{1-p_0 q}{1+q-p_0 q}$*, then a non-conformity strategy profile maximizes the affiliates' aggregate expected payoff among all affiliates' strategy profiles; otherwise, a conformity strategy profile maximizes the affiliates' aggregate expected payoff among all affiliates' strategy profiles.*

Proposition 1 delivers a bang-bang result: depending on the parameters of our model, the affiliates and the consumer may have completely aligned interests or completely opposite interests (but nothing in between) on whether affiliates recommend the salient product when they receive a positive signal on its quality.

Part (1) of Proposition 1 shows that regardless of the parameters of the model, the consumer benefits from affiliates' conformity. Intuitively, affiliates' conformity allows the consumer to have more accurate information on at least one product (the salient product). This increases the probability that the first product that the consumer clicks on is of high quality, leading to: (1) a shorter and less costly search, and (2) a higher probability of the search ending with a purchase of a high quality product.

Note that the condition $\lambda \geq \frac{1-p_0 q}{1+q-p_0 q}$ in part (2) of Proposition 1 can be equivalently expressed as $p_0 \geq \frac{1-\lambda-\lambda q}{q-\lambda q}$, or $q \geq \frac{1-\lambda}{\lambda+p_0-\lambda p_0}$. Therefore, affiliates' total expected payoffs are the highest under non-conformity, if: (1) a significant fraction of affiliates' compensation is per-click, (2) the consumer is likely to continue searching after a click that doesn't lead to a purchase, and (3) affiliates receive sufficiently accurate signals on products' qualities. In this case, the consumer and the affiliates have opposing interests.[18]

[17] This is akin to a hypothetical scenario in which affiliates' strategy profile can be exogenously determined and publicly announced and the consumer best responds to it.

[18] Note also that $\frac{1-\lambda-\lambda q}{q-\lambda q} > 1$ if and only if $1-\lambda > q$, or $\lambda + q < 1$. It follows that, if $\lambda + q < 1$, then $\frac{1-\lambda-\lambda q}{q-\lambda q} > p_0$ for all $p_0 \leq 1$, that is, the condition in part (2) of Proposition 1 cannot hold. Therefore, if a sufficiently large fraction of affiliates' compensation is per-purchase (low λ) or if the consumer's search is sufficiently short (low q). Affiliates (as a group) have aligned interests with the consumer regardless of the affiliates' information accuracy p_0.

Intuitively, when affiliates are compensated significantly per click, and when it is likely that the consumer will continue clicking after not purchasing, the average expected payoff for an affiliate increases in the expected length of the consumer's search. At the same time, as long as affiliates' compensation does not put significant weight on the purchase, affiliates' expected payoffs are not reduced much if the search is likely not to end in a purchase. Conversely, suppose affiliates are compensated mainly per purchase, and the consumer is expected to stop searching after an unsuccessful first click. In that case, affiliates' interests align with the consumer's preference for affiliates' conformity. Figure 1 captures these observations.

The effect of the accuracy of the affiliates' signals is more subtle. A high accuracy signal implies that the search will likely end in purchase even if affiliates do not conform to their recommendations, thus making the non-conformity outcome better for affiliates. This effect is more substantial when a larger fraction of affiliates' compensation is per purchase and the probability that the consumer continues searching after one low-quality product is high. We can see this effect by comparing the contour lines in Figure 1.

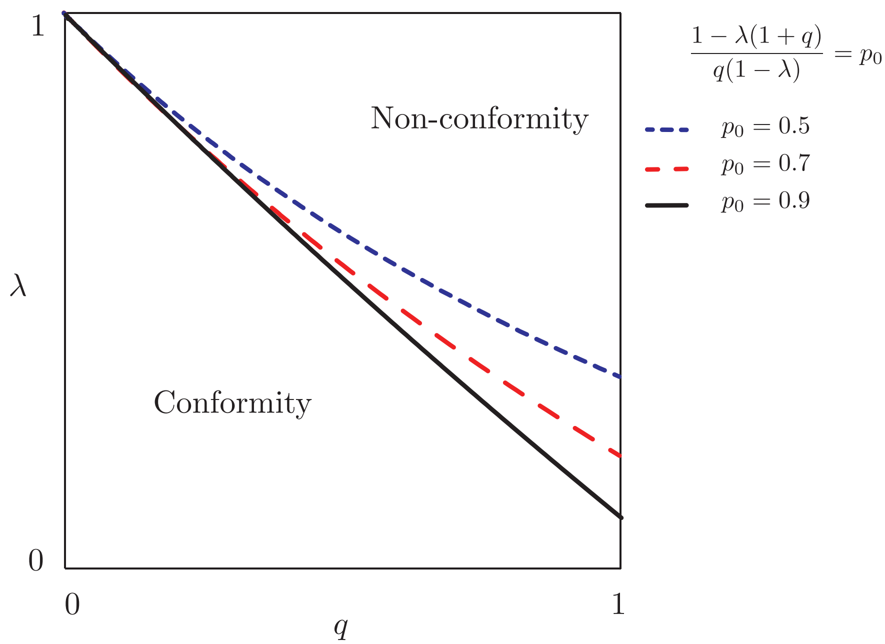

Figure 1: Affiliates' preference for conformity

For three values of p_0, the figure illustrates a partition of the q-λ space that distinguishes when affiliates prefer conformity (vs. non-conformity). Affiliates' aggregate payoff is highest under conformity when affiliates are compensated heavily per-purchase (small λ) and consumer search is short (small q). In contrast, affiliates' aggregate payoff is highest under non-conformity when affiliates are compensated heavily per-click (large λ) and consumer search is long (large q). If affiliates signals are more accurate (large p_0) the range of parameters for which affiliates' aggregate payoff is highest under conformity is smaller. (In the figure $N = 10$.)

3.2 Equilibrium and affiliates' conformity

We now analyze the conditions under which affiliates' conformity and non-conformity can be supported as equilibrium outcomes. We first show that regardless of the parameters of the model, there exists a non-conformity equilibrium, i.e., an equilibrium in which no affiliate ever recommends the salient product.

Proposition 2. *There always exists a non-conformity equilibrium.*

Proposition 2 provides bad news for the consumer, whose payoff is maximized under affiliates' conformity. Proposition 2 suggests that affiliates' non-conformity can emerge in equilibrium regardless of the parameters of the model. When combined with Proposition 1, Proposition 2 provides good news to the affiliates when affiliates' payoffs have a significant per-click component or consumers are more likely to continue their search beyond the first click.

To gain intuition for Proposition 2, it is useful to consider the case that the consumer follows exactly two affiliates. In this case, essentially all equilibria are non-conformity equilibria.[19] To gain intuition into why when $N = 2$ this equilibrium is unique, consider an equilibrium candidate in which any affiliate who receives a good signal for the salient product recommends the salient product with some positive probability. Assume by contradiction that this is an equilibrium. Then, if one affiliate recommends the salient product and one doesn't, this informs the consumer that one of them received a good signal for the salient product and the other, at least with some positive probability, received a negative signal for the salient product. As a result, the consumer's posterior puts a higher probability that the recommended non-salient product is of high quality than that the salient product is high quality. If such an equilibrium exists, the only situation in which the consumer will click the salient product is if both affiliates recommend the salient product. However, in that case, each of the affiliates can deviate to recommend a non-salient product and receives the consumer's click with probability 1.

When $N > 2$, additional equilibria may exist. Motivated by Proposition 1, which tells us that a conformity equilibrium maximizes the consumer's payoffs and sometimes maximizes also affiliates' payoffs (and thus aggregate social welfare), our next result characterizes the set of parameters for which affiliates' conformity is an equilibrium.

Proposition 3. *Fix any $N > 2$. For any q, λ there exist $\frac{1}{2} \leq \underline{p} \leq \bar{p} < 1$ such that:*[20]

(1) for every $p_0 > \bar{p}$, there exists a conformity equilibrium; and

(2) for every $p_0 < \underline{p}$, a conformity equilibrium does not exist.

The intuition for Proposition 3 is most straightforward when considering N large yet finite. That is, the consumer observes the recommendations of a large (finite) number of affiliates: Consider affiliate j, and suppose that the quality of the salient product is high and all other affiliates act according to the conformity strategy profile, that is, if they receive a positive signal on the salient

[19] There is only one additional equilibrium in which at most one affiliate recommends the salient product. This equilibrium is equivalent to a non-conformity equilibrium in the sense that no one product is ever recommended by two affiliates in equilibrium.

[20] Moreover, for any $0 \leq \lambda \leq 1$, there exists a unique $\hat{q} \in (0, 1)$ such that for any $q \neq \hat{q}$, the upper bound $\underline{p}(q, \lambda)$ for the range of p_0 on the non-existence of a conformity equilibrium is strictly larger than $\frac{1}{2}$. Therefore, generically, the proposition can be re-written with strict inequalities $\frac{1}{2} < \underline{p} \leq \bar{p} < 1$ and we can say that generically for any $0 \leq q \leq 1$ and $0 \leq \lambda \leq 1$ there exists ranges of p_0 in which a conformity equilibrium exists and ranges of p_0 in which a conformity equilibrium does not exist.

product they post a recommendation for it. Then, there is a considerable probability that the salient product gets most of the recommendations observed by the consumer. This probability increases in the accuracy of the affiliates' signals, p_0. As p_0 approaches 1, the probability that conditional on being high quality, the salient product gets an overwhelming majority of the recommendations is high. In that case, the consumer clicks first on a link to the salient product with high probability, and because the salient product is of high quality, the consumer will not search further. Therefore, to have a chance of having his link clicked on, affiliate j would need to recommend the salient product when the salient product is of high quality. It is just left to note that, for a high information accuracy p_0, the probability that the salient product is of high quality conditional on affiliate j receiving a positive signal is high.

When considered in tandem with Proposition 1, Proposition 3 carries two main welfare implications. First, even when both the consumer and the affiliates would benefit from the affiliates' conformity, a conformity equilibrium may not exist, leading to a sub-optimal equilibrium outcome. This sub-optimal outcome occurs, for example, when affiliates' signals are sufficiently inaccurate.

Second, even when affiliates as a group prefer non-conformity, a conformity equilibrium may exist. For example, part 1 of Proposition 3 tells us that a conformity equilibrium exists if affiliates receive sufficiently accurate product quality signals. However, Proposition 1 tells us that this is the case in which the affiliates can maximize their aggregate payoffs when they never recommend the salient product. Notably, a conformity equilibrium is always ideal for the consumer. We illustrate these two welfare implications in Figure 2.

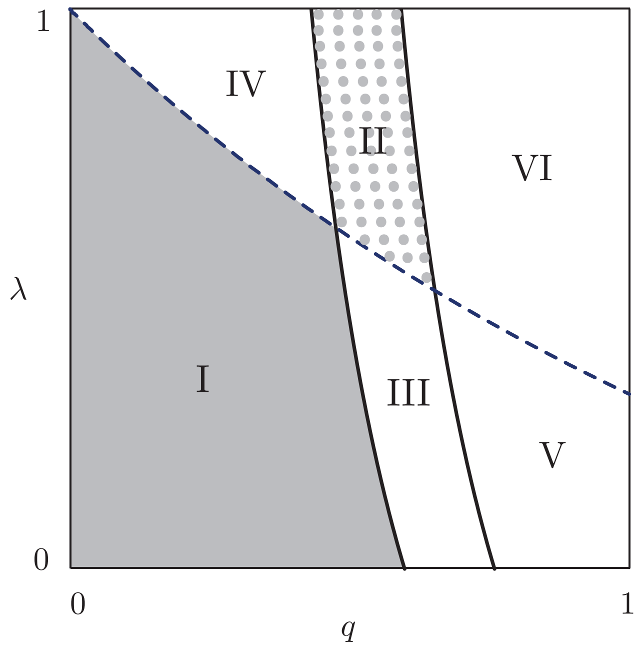

Figure 2: Conformity-driven inefficiency

Below the dark blue dashed line (regions I, III, and V), affiliates achieve the maximal aggregate payoff under affiliates' conformity, whereas above the dashed line (regions II, IV, and VI), affiliates achieve maximal aggregate payoff under their non-conformity. Between the two solid lines (regions II and III), a conformity equilibrium exists. In regions I and IV, a conformity equilibrium does not exist due to affiliates' too strong incentive to recommend the salient product even after they learn it is of low quality. In regions V and VI, a conformity equilibrium does not exist due to too-strong incentives to recommend the non-salient product. In region I (shaded with solid grey), a conformity equilibrium does not exist, but both the consumer and the affiliates achieve the maximal payoffs under affiliates' conformity. In region II (shaded with dots), a conformity equilibrium exists, but affiliates can maximize their aggregate payoff from non-conformity. (In the figure $N = 10$, $p_0 = 0.57$.)

We can say more.

Corollary 1. *Let $\bar{p}_0(q, \lambda)$ be the lowest signal accuracy above which there always exists a conformity equilibrium, i.e., $\bar{p}_0(q, \lambda) := \min\{\bar{p}(q, \lambda)\}$. Then,*

(1) there exist $0 < \underline{\underline{q}} \leq \overline{\overline{q}} < 1$ such that if $q < \underline{\underline{q}}$, then $\bar{p}_0(q, \lambda)$ is decreasing in λ, and if $q > \overline{\overline{q}}$, then $\bar{p}_0(q, \lambda)$ is increasing in λ;

(2) for any λ, there exist $0 < \underline{q} \leq \overline{q} < 1$ such that if $q < \underline{q}$, then $\bar{p}_0(q, \lambda)$ is decreasing in q, and if $q > \overline{q}$, then $\bar{p}_0(q, \lambda)$ is increasing in q.

Corollary 1 shows that the expected length of the consumer's search and the affiliates' compensation scheme have non-monotonic effects on the existence of a conformity equilibrium. As a result, a conformity equilibrium is most likely to exist if the affiliate compensation scheme is between per-click and per-purchase compensation and when consumers' search is of intermediate length. This assertion is interesting because, as mentioned in the introduction, in recent years, affiliate programs have converged on an affiliate compensation scheme that is effectively between per-click and per-purchase compensation.

To better understand the source of this non-monotonicity, assume that all affiliates employ conformity strategies and consider the following two conceivable unilateral deviations: (1) an affiliate who receives a high-quality signal for the salient product may, nevertheless, recommend a non-salient product, and (2) an affiliate who receives a low-quality signal for the salient product may, despite his signal, recommend the salient product.

Consider part (2) of Corollary 1. An increase in the expected duration q of the consumer's search makes deviation (1) more profitable and deviation (2) less profitable. So much so that when the search duration is short, there is no p_0 for which deviation (1) is viable. Hence, when the search duration is short, the range of p_0 for which any profitable deviation exists reduces in the search duration. Once the search duration is sufficiently high, deviation (2) is no longer viable, but for sufficiently low p_0, deviation (1) is. In this range of q, the range of p_0 for which any profitable deviation exists widens in the search duration. This effect is captured by part (2) of Corollary 1 and illustrated in Figure 3.

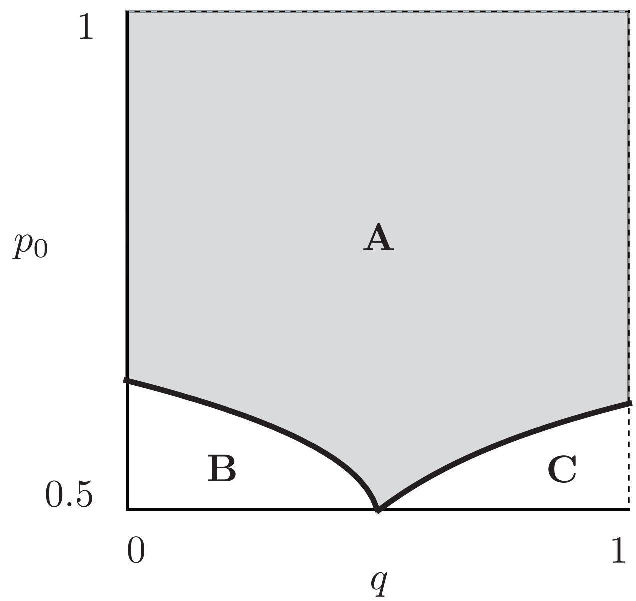

Figure 3: Existence of conformity equilibrium

In region A, a conformity equilibrium exists. In regions B and C a conformity equilibrium does not exist: in region B this is due to too-strong incentives to recommend the salient product (even when an affiliate receives a bad signal for it), whereas in region C this is due to too-strong incentives to recommend the non-salient product. (In the figure $\lambda = 0.5$, $N = 10$.)

The intuition for part (1) of the corollary follows a similar logic. An increase in the per-click compensation, λ, make deviation (1) more profitable because it reduces the importance of recommending a product that will be purchased (recalls that the salient product, if clicked on, is more likely to be purchased). For the same reason, an increase in λ also makes deviation (2) less profitable (more generally, it increases the expected profit from recommending a non-salient product vis-a-vis the salient product). Therefore, as part (1) of the corollary shows, an increase in λ has the opposite effect on p_0 when the conformity equilibrium is threatened by the possibility of deviation (1) versus when it is threatened by deviation (2).

4 Discussion: extensions and broader implications

In this section, we first show that our model is robust to the case in which the consumer does not observe the quality of a product directly after clicking on an affiliate link. We derive implications for the case of experience goods. We then consider the platform's problem: optimal design of affiliate incentives. We show that whether the platform prefers to put heavier weight on per-purchase vs.

per-click compensation depends on how good the platform is at converting clicks to purchases and on the consumer's expected search duration. We conclude by providing managerial implications.

4.1 Experience goods and the role of affiliates

In our baseline model we assumed that once a consumer clicks on a link to a product, she observes the product's quality accurately. As a result, affiliates had the role of influencing the consumer's search without having an effect on whether a consumer purchases a product after she viewed it. While this assumption provides a reasonable approximation for many product markets, some products require the consumer to purchase before learning the quality accurately. In such cases, the consumer may take into account the affiliate's recommendations in her purchase decision even after viewing the product page.

We now consider a more general model in which the consumer observes some information about the product in the form of a potentially noisy signal. Formally, suppose that after the consumer clicks an affiliate link for product $i \in \{0\} \cup NS$, she observes a signal $\tilde{s}(i) \in \{\tilde{h}, \tilde{l}\}$ such that $Pr(\tilde{s}(i) = \tilde{h}|q_i = H) = Pr(\tilde{s}(i) = \tilde{l}|q_i = L) = \pi_0 > 1/2$. We assume that the consumer's signal is conditionally independent from those of the affiliates. The affiliates' information, the payoffs for the affiliates and the consumer, and the timing of the game remain the same as in our baseline model (Section 2).

We find all our results for the baseline model hold.

Proposition 4. *In the model in which the consumer observes a signal from clicking on an affiliate's link:*

(1) (Proposition 1) A conformity strategy profile maximizes the consumer's expected payoff among all affiliates' strategy profiles.

(2) (Proposition 1) If $\lambda > \frac{1-p_0 q}{1+q-p_0 q}$, then a non-conformity strategy profile maximizes the affiliates' aggregate expected payoff among all affiliates' strategy profiles; otherwise, a conformity strategy profile maximizes the affiliates' aggregate expected payoff among all affiliates' strategy profiles.

(3) (Proposition 2) There always exists a non-conformity equilibrium.

(4) (Proposition 3) Fix any $N > 2$, there exist $\frac{1}{2} < p' \leq p'' < 1$ such that, for all q, λ:

(a) for every $p_0 > p''$, there exists a conformity equilibrium;

(b) for every $p_0 < p'$, a conformity equilibrium does not exist.

An interesting special case is the case of *experience goods*—goods on which the consumer obtains no quality signal before trying out the good.[21] This is captured by the limit case in which the consumer observes an uninformative signal, i.e., $\pi_0 = 1/2$. For this case we can say more.

Corollary 2. *In the model of experience goods:*

(1) both conformity and non-conformity strategy profiles maximize the affiliates' aggregate expected payoff among all strategy profile.

(2) affiliate compensation (λ) and the length of consumer search (q) do not affect the range of affiliate signal accuracy level (p_0) for which a conformity equilibrium exists.

An immediate broader implication of this observation is that in markets for experience goods, the retailer cannot affect equilibrium behavior via affiliate compensation.

The intuition behind Corollary 2 begins by noting that this scenario is mathematically equivalent to our baseline model when parameters' values are set to $\lambda = 1$ and $q = 0$. To see why, note that for every strategy profile pursued by the affiliates, the consumer's unique best response is to click on the link to a product that has the highest posterior probability of being high quality and to purchase that product. That is, any affiliate's expectation of being clicked on is identical to our baseline model in which the consumers clicks only once ($q = 0$). The affiliate then expects to be paid $\lambda + (1 - \lambda) = 1$ if his link is clicked on by the consumer, but the probability of purchase conditional on clicking is 1, and therefore identical to what he's paid in our baseline model when $\lambda = 1$.

4.2 Platform's design

So far we have treated the affiliate compensation scheme as exogenous. In reality, however, every online retailer has the ability to design the compensation scheme of its affiliate program. In the early days of e-commerce, formal pay-per-click compensations schemes were common. The rise of

[21] Many personal services fall into this category (e.g. restaurant, hairdresser, beauty salon, theme park, travel, holiday).

bots and click-fraud led many retailers to change their compensation model to pay-per-purchase, in order to guarantee that paid clicks arrive from real consumers.

However, as suggested in Section 2.5.1, a large component of the compensations scheme of many affiliate programs is still in expectations in the form of pay-per-click compensation. Affiliate programs achieve that by paying affiliates for *any* purchase of the consumer within a predetermined time-frame after the click. The longer the time-frame the larger the expected share of compensation coming from per-click-like payments (relative to per-purchase). Additional platform design consideration that increase the effective per-click compensation are larger product variety and better platform internal recommendation system, as they increase the probability that a consumer buys something on the platform.

We capture the platform's objective within our model as follows: Let r be the probability that, conditional on clicking, the consumer buys a different product on the platform, and let v be the expected sale price of that product. Normalize the sale price of the product recommended by the affiliate to 1. Finally, let ρ be the commission in terms of the fraction of the sale price paid to the affiliate.

The expected payoff for the affiliate, conditional on the consumer clicking on his recommendation, can be written as

$$\rho \cdot (\Pr(\text{consumer buys the affiliate-recommended product conditional on clicking}) + r \cdot v).$$

By varying the time frame for compensation and its recommendation algorithm the platform changes r (and maybe also v) and thus the relative weights given to pay-per-click vs. pay-per-purchase in the affiliate's compensation scheme. We note that this compensation scheme is captured by our modeling above when ρ is normalized so that $\rho + r \cdot v \cdot \rho = 1$.[22]

We denote by P_c the probability that the platform sells a product recommended by an affiliate in a conformity equilibrium and by P_{nc} the probability of such a sale in a non-conformity equilibrium.

[22] To see how, set $\lambda = r \cdot v \cdot \rho$ and note that by the normalization $1 - \lambda = \rho$. Moreover, because the normalization does not affect any of our results above, the platform can always change ρ without affecting the equilibrium—the only restriction is that ρ is sufficiently high so that, in equilibrium, affiliates' expected incomes are such that they stay active. Since the platform can modify ρ it will then pay affiliate the lowest total payoff to keep them active and can adjust the compensation structure to the equilibrium played so that it delivers to affiliates the same minimal expected payoff. Thus, without loss of generality, we can focus on the platform's choice of λ (through its choice of policies affecting r, such as the time frame considered for affiliate compensation), and assume that the platform chooses ρ to maximize its revenues.

We denote by τ the additional expected revenue to the platform from a click directed to it by an affiliate, and by x_c and x_{nc} the expected number of clicks in a conformity and non-conformity equilibrium respectively.

Let $E \in \{c, nc\}$. The platform's revenue is then captured by

$$\pi_E = P_E + x_E \cdot \tau \tag{1}$$

Next, note that our analysis above (Proposition 1) shows that $P_c \geq P_{nc}$. On the other hand, the same analysis shows that $x_c \leq x_{nc}$. Therefore, holding all else equal, a platform that converts a click into higher revenues (high τ) will prefer a non-conformity equilibrium, whereas a platform that does not converts clicks into high revenues will prefer a conformity equilibrium.

In addition, as we showed in Section 4.1, for a platform selling experience goods $x_c = x_{nc}$, and such a platform, therefore, will always prefer a conformity equilibrium.

The remaining question is then, how can a platform affect whether a conformity or non-conformity equilibrium is played. Due to the multiplicity of equilibria the platform may not be able to deterministically affect which equilibrium is played. However, it can affect whether a conformity equilibrium exists or not. Corollary 1 provides guidelines to a platform seeking to affect the existence of a conformity equilibrium. In particular,

1. If the consumer's search is expected to be short (low q), the platform can make conformity equilibrium exist (not exist) by increasing (decreasing) the share of compensation paid per-click (λ).

2. If the consumer's search is expected to be long (high q), the platform can make conformity equilibrium exist (not exist) by decreasing (increasing) the share of compensation paid per-click (λ).

4.3 Managerial implications

This paper provides managerial implications for affiliates, brands, and platforms alike.

Affiliates

If affiliates have accurate signals of products' qualities, they will benefit the most under a non-conformity equilibrium, which always exists. Therefore, affiliates can benefit from coordinating on a non-conformity equilibrium. However, this is the scenario in which a conformity equilibrium also exists, which makes affiliates' coordination on a non-conformity strategy profile tricky because consumers prefer a conformity equilibrium. A well-informed regulator might try to crack down on such attempts.

In contrast, if affiliates' signals are less accurate, they will benefit from conformity. Unfortunately, this is the scenario in which a conformity equilibrium may not exist, to the detriment of the affiliates and the consumers. All parties in this case (except non-salient brands) will benefit from increasing the availability of accurate information to all affiliates. Improving the accuracy of affiliates' product quality information can be done, for example, by a salient brand or platform providing samples to affiliates (more on that below) or by affiliates investing more intensively in gathering information.

In markets for experience goods, in which consumers do not receive quality signals independent of affiliates' recommendations, affiliates are indifferent between conformity and non-conformity equilibria. Because consumers prefer conformity equilibria, if affiliates engage in conformity equilibrium, more consumers will be attracted to the market, ultimately benefiting all parties.

Brands

When affiliates receive inaccurate signals about product quality, a conformity equilibrium does not exist. Therefore, the salient brand may want to make accurate quality information readily accessible as it will make it more likely that a conformity equilibrium exists. Interestingly, this is consistent with salient brands sending samples to influencers with no conditions on how favorable their reviews will be.

Platforms

Even while paying per-purchase to avoid click fraud, platforms can balance between per-purchase compensation and per-click compensation. Shortening the time window for purchase and introducing a policy that only the latest clicked-on affiliate is compensated for purchases (the "last click wins" model) reduce the "per-click" weight in affiliate compensation. Improvements in the matching and recommendations algorithms of a platform, increase in the scope of products carried by a platform, and general price competitiveness and consumer loyalty to the platform increase the "per-click" weight in the compensation.

Whether the platform prefers a conformity or non-conformity equilibrium depends on the platform's ability to extract additional unrelated purchases from a click. A platform with a good in-site recommendation system and a large variety of product categories, such as eBay or Amazon, may prefer to maximize clicks and, therefore, encourage a non-conformity equilibrium, perhaps by setting the compensation scheme in a way that will not admit a conformity equilibrium. On the other hand, a platform with a less established recommendation system or a more specialized platform with a narrow focus is likely to prefer a conformity equilibrium, which increases the probability of purchase.

If consumers are expected to conduct longer searches, a platform seeking to facilitate a conformity equilibrium will look for a per-purchase compensation scheme. The reverse is true if consumers conduct shorter searches on average. In contrast, when consumers cannot learn about the product quality independently of the affiliates' recommendations, the platform's affiliate compensation scheme does not affect equilibrium behavior and whether salient or non-salient products are sold in equilibrium.

References

Berman, R., A. Öry, and X. Zheng (2023). Influence or advertise: The role of social learning in influencer marketing. Available at SSRN 4324888.

Duffy, D. L. (2005). Affiliate marketing and its impact on e-commerce. *Journal of Consumer MarketingJournal of Consumer Marketing 22*(3), 161–163.

Edelman, B. and W. Brandi (2015). Risk, information, and incentives in online affiliate marketing. *Journal of Marketing Research 52*(1), 1–12.

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

Gilliland, D. I. and J. M. Rudd (2013). Control of electronic channel affiliates: An exploratory study and research propositions. *Journal of Business Research 66*(12), 2650–2656.

Hoffman, D. L. and T. P. Novak (2000). How to acquire customers on the web? *Harvard Business Review 78*(3), 179–188.

Jain, S. and K. Qian (2021). Compensating online content producers: A theoretical analysis. *Management Science 67*(11), 7075–7090.

Janssen, M. and C. Williams (2021). Influencing search. CEPR Discussion Paper 15811.

Jefferson, A. (2016). Networks help drive affiliate marketing into the mainstream: Advertisers and publishers evolve as the industry shifts. *A Forrester Consulting Thought Leadership Paper Commissioned By Rakuten Affiliate Network*.

Karabas, I., I. Kareklas, T. J. Weber, and D. D. Muehling (2021). The impact of review valence and awareness of deceptive practices on consumers' responses to online product ratings and reviews. *Journal of Marketing Communications 27*(7), 685–715.

Katona, Z. (2020). Competing for influencers in a social network. NET Institute Working Paper No. 13-06, Available at SSRN 2335679.

Libai, B., E. Biyalogorsky, and E. Gerstner (2003). Setting referral fees in affiliate marketing. *Journal of Service Research 5*(4), 303–315.

Mathur, A., A. Naryanan, and M. Chetty (2018). Endorsements on social media: An empirical study of affiliate marketing disclosures on youtube and pinterest. *Proceedings of the ACM on Human-Computer Interaction 2*(CSCW), 1–26.

Moe, W. W. and D. A. Schweidel (2012). Online product opinions: Incidence, evaluation, and evolution. *Marketing Science 31*(3), 372–386.

Moe, W. W. and M. Trusov (2011). The value of social dynamics in online product ratings forums. *Journal of Marketing Research 48*(3), 444–456.

Nistor, C. and M. Selove (2022). Influencers: The power of comments. Available at SSRN 4118010.

Olbrich, R., P. M. Bormann, and M. Hundt (2019). Analyzing the click path of affiliate-marketing campaigns. *Journal of Advertising Research 59*(3), 342–356.

Papatla, P. and A. Bhatnagar (2002). Choosing the right mix of on-line affiliates: How do you select the best? *Journal of Advertising 31*(3), 69–81.

Pei, A. and D. Mayzlin (2022). Influencing social media influencers through affiliation. *Marketing Science 41*(3), 593–615.

Shopify (2024). What is affiliate marketing? everything you need to know in 2024. `https://www.shopify.com/blog/affiliate-marketing`. Accessed: 2024-01-22.

Sridhar, S. and R. Srinivasan (2012). Influence effects in online product ratings. *Journal of Marketing 76*(5), 70–88.

Suryanarayana, S. A., D. Sarne, and S. Kraus (2019). Information disclosure and partner management in affiliate marketing. *DAI '19: Proceedings of the First International Conference on Distributed Artificial Intelligence* (3), 1–8.

Syrdal, H. A., S. Myers, S. Sen, P. J. Woodroof, and W. C. McDowell (2023). Influencer marketing and the growth of affiliates: The effects of language features on engagement behavior. *Journal of Business Research 163*, 113875.

Appendix: Proofs of Sections 3 and 4

Proof of Proposition 1. We first use the ex ante homogeneity across the non-salient products to simplify the consideration of affiliates' strategy spaces. Because any two non-salient products are ex ante identical, if an affiliate j receives high signals for two non-salient products, then any probability distribution of recommendations between the two induces the same outcome. That is, for any $i_1, i_2 \in \eta_j$, with $i_1, i_2 \neq 0$ and $s_j(i_1) = s_j(i_2) = h$, all of the strategies for affiliate j holding $Pr(\text{recommend } i_1) + Pr(\text{recommend } i_2)$ constant induce the same outcome. Therefore, for the consideration of payoffs, it suffices to pick a representative from each set of strategies with the same total probability of recommending non-salient products—one that allocates that probability entirely to one non-salient product. As a result, we are reduced to consider the following set of strategies for each affiliate j: $(\alpha_j, \beta_j) \in [0,1]$, where α_j and β_j represent the probabilities affiliate j will recommend product 0 contingent on his signal about it. That is, with probability α_j, affiliate j will recommend product 0 after a high signal about product 0, and with probability $(1-\alpha_j)$, he will recommend a non-salient product for which he receives a high signal; with probability β_j, affiliate j will recommend product 0 after a low signal about product 0, and with probability $(1-\beta_j)$, he will recommend a non-salient product for which he receives a high signal.

Denote by $Pr\left(H|0^{(n)}\right)$ the consumer's posterior belief of product 0 being of high quality when she sees n recommendations for product 0. Being a risk-neutral expected utility maximizer, the consumer's best response in her search to a strategy profile for the affiliates is: (1) when $Pr\left(H|0^{(n)}\right) \geq p_0$, search product 0 first, if it is of high quality, stop the search and make a purchase of one unit of product 0, otherwise (with probability q) search a non-salient product among the recommended ones; when $Pr\left(H|0^{(n)}\right) < p_0$, search any non-salient product that is recommended first, if it is of high quality, stop and make a purchase of the searched product, otherwise, (with probability) search another recommended non-salient product. Therefore, the consumer's expected utility is

$$E[U(q)|n] = \begin{cases} (1-c)p_0 + (1-2c)(1-p_0)p_0 q & \text{if } Pr\left(H|0^{(n)}\right) < p_0 \\ (1-c)Pr\left(H|0^{(n)}\right) + (1-2c)\left(1 - Pr\left(H|0^{(n)}\right)\right)p_0 q & \text{otherwise} \end{cases} \tag{2}$$

Collecting $Pr\left(H|0^{(n)}\right)$ in the later case, we have

$$(1-c)Pr\left(H|0^{(n)}\right) + (1-2c)\left(1 - Pr\left(H|0^{(n)}\right)\right)p_0 q = [1 - p_0 q - c(1-2p_0 q)]Pr\left(H|0^{(n)}\right) + (1-2c)p_0 q.$$

The expression above is linear in $Pr\left(H|0^{(n)}\right)$, whose coefficient $1 - p_0 q - c(1-2p_0 q)$ must be positive. This is because when $1 - 2p_0 q < 0$, then clearly it is positive; and when $1 - 2p_0 q > 0$, $1 - p_0 q > 1 - 2p_0 q > c(1-2p_0 q)$, so it is positive as well. Therefore, consumer's expected utility is bounded from below by that when she only considers non-salient products and attains its maximum when her posterior about product 0 is at the highest, conditional on the number of recommendations for product 0 is high enough for her to consider product 0 first.

To prove part (1) of the proposition, it is sufficient to show that the conformity strategy profile maximizes the consumer's posterior belief about product 0's quality conditional on the number of recommendations it receives.

Suppose that the affiliates have chosen the strategy profile represented by $\{(\alpha_j, \beta_j)\}_{j=1,\dots,N}$. Let $n \leq N$ be the number of recommendations for product 0 the consumer observes. To streamline notations, we define the following partitions on the set of affiliates. Let $H(n)$ ($L(n)$, resp.) be the set of affiliates who receive high signals (low signals, resp.) about product 0, such that $|H(n)| = N_1 \leq N$ and $|L(n)| = N - N_1$. Furthermore, let $H^0(n) \subset H(n)$ ($L^0(n) \subset L(n)$, resp.) denote the set of affiliates who choose to recommend product 0. Then by definition, $|H^0(n)| + |L^0(n)| = n$.

First, the probability that product 0 receives n recommendations conditional on that N_1 affiliates received high signals about it is

$$Pr\left(0^{(n)}|(h^{N_1}, l^{N-N_1})\right) = \sum_{m=1}^{n} \binom{N_1}{m}\binom{N-N_1}{n-m} \prod_{\substack{j \in H^0(n) \\ |H^0(n)|=m}} \alpha_j \prod_{g \in H(n)\setminus H^0(n)} (1-\alpha_g) \prod_{k \in L^0(n)} \beta_k \prod_{l \in L(n)\setminus L^0(n)} (1-\beta_l). \tag{3}$$

Next, as the number N_1 of high signals for product 0 ranges from n to N, the probability that product 0 will receive n recommendations conditional on its true quality is

$$Pr\left(0^{(n)}|H\right) = \sum_{N_1=n}^{N} Pr\left((h^{N_1}, l^{N-N_1})|H\right) Pr\left(0^{(n)}|(h^{N_1}, l^{N-N_1})\right), \tag{4}$$

$$Pr\left(0^{(n)}|L\right) = \sum_{N_1=n}^{N} Pr\left((h^{N_1}, l^{N-N_1})|L\right) Pr\left(0^{(n)}|(h^{N_1}, l^{N-N_1})\right). \tag{5}$$

where the probabilities that product 0 receives N_1 high signals and $N - N_1$ low signals conditional on its true quality are

$$Pr\left((h^{N_1}, l^{N-N_1})|H\right) = \binom{N}{N_1} p_0^{N_1}(1-p_0)^{N-N_1}, \tag{6}$$

$$Pr\left((h^{N_1}, l^{N-N_1})|L\right) = \binom{N}{N_1} (1-p_0)^{N_1} p_0^{N-N_1}. \tag{7}$$

Given that the prior quality for product 0 being high or low is equally likely, by Bayes's Rule, the consumer's posterior belief that product 0 is of high quality is

$$Pr\left(H|0^{(n)}\right) = \frac{Pr\left(0^{(n)}|H\right)}{Pr\left(0^{(n)}|H\right) + Pr\left(0^{(n)}|L\right)}. \tag{8}$$

By eqs. (6) and (7), the distribution of signals about product 0 received by the affiliates is independent of the affiliates' strategies, and hence, we are reduced to maximize the conditional probabilities $Pr\left(0^{(n)}|(h^{N_1}, l^{N-N_1})\right)$ given by eq. (3). By the symmetry between the α_j's and β_j's, and because $p_0 > 1/2$, it is without loss of generality to assume $\alpha_j \geq \beta_j$ for $j = 1, \ldots, N$. Note that when $0 \leq \beta_j \leq \alpha_j \leq 1$, the expression

$$\prod_{\substack{j \in H^0(n) \\ |H^0(n)|=m}} \alpha_j \prod_{g \in H(n) \setminus H^0(n)} (1 - \alpha_g) \prod_{k \in L^0(n)} \beta_k \prod_{l \in L(n) \setminus L^0(n)} (1 - \beta_l) \leq 1,$$

and it attains maximum value of 1 when either

(1) $(\alpha_j, \beta_j) = (1, 1)$ for $j = 1, \ldots, N$, requiring $H^0(n) = H(n)$ and $L^0(n) = L(n)$; or

(2) $(\alpha_j, \beta_j) = (1, 0)$ for $j = 1, \ldots, N$.

However, the first situation above is only possible when $N_1 = N$, i.e., the entire group of affiliates all receive high signals about product 0. It is hence non-generic. When $(\alpha_j, \beta_j) = (1, 0)$ for $j = 1, \ldots, N$, necessarily, $m = n = N_1$, implying that the number of recommendations for product 0 that the consumer sees coincides with the number of affiliates who received high signals about product 0. Therefore, we have shown the strategy profile specified as $(\alpha_j, \beta_j) = (1, 0)$ for $j = 1, \ldots, N$ maximizes the consumer's expected utility from best responding to this strategy profile. This completes the proof of part (1).

To prove part (2), suppose that the affiliates act according to a strategy profile characterized by $\{(\alpha_j, \beta_j)\}$ for $j = 1, \ldots, N$, yielding n recommendations for product 0. Then the affiliates' total expected payoff is

$$\begin{aligned} EU(\alpha, \beta) = \lambda &\left[1 + q\left(1 - \max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\}\right)\right] \\ &+ (1 - \lambda)\left[\max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\} + p_0 q\left(1 - \max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\}\right)\right]. \end{aligned} \tag{9}$$

Collecting $\max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\}$, we can express eq. (9) as

$$EU(\alpha, \beta) = [1 - p_0 q - (1 + q - p_0 q)\lambda] \cdot \max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\} + p_0 q + (1 + q - p_0 q)\lambda. \tag{10}$$

Note that the affiliates' total expected payoff is linear in $\max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\}$. Therefore, $EU(\alpha, \beta)$ attains its maximum when $\max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\} = p_0$ if $1 - p_0 q - (1 + q - p_0 q)\lambda < 0$, and when

$\max\left\{p_0, Pr\left(H|0^{(n)}\right)\right\} = Pr\left(H|0^{(n)}\right)$ is at its highest if $1 - p_0q - (1 + q - p_0q)\lambda \geq 0$. The sign of the expression $1 - p_0q - (1 + q - p_0q)\lambda$ induces the affiliates' preference dichotomy. When $1 - p_0q - (1 + q - p_0q)\lambda \leq 0$, the affiliates want to maximize the probability that $Pr\left(H|0^{(n)}\right) \leq p_0$ across all possible realized recommendation profiles, which will be fulfilled by a non-conformity profile. To the contrary, when $1 - p_0q - (1 + q - p_0q)\lambda \geq 0$, the affiliates want to maximize the probability that $Pr\left(H|0^{(n)}\right) \geq p_0$, which will be achieved by a conformity profile. This completes the proof of part (2). □

Proof of Proposition 2. Consider the following strategies: all affiliates recommend non-salient products regardless of their signals for product 0. The consumer follows only recommendations for no-salient products. It is immediate that the affiliates' strategies are best responses to the consumer's strategy. Given that the probability that an affiliate recommends the salient product is zero, the consumer's strategy can be supported as best response to the affiliates' strategies with the following beliefs: if an affiliate recommends product zero it is a mistake, which is equally likely when the affiliate receives a high or low signal for the salient product. □

Proof of Proposition 3. (1) We verify that a conformity strategy profile constitutes a pure-strategy equilibrium in the extreme scenario when $p_0 = 1$ by checking the non-profitability for any affiliate from a unilateral deviation. Then part (1) of Proposition 3 follows by the continuity of an affiliate j's expected payoff in p_0.

Case 1. Affiliate j observes a high signal $s_j(0) = h$ for product 0.

When $p_0 = 1$, $Pr\left((h^k, l^{N-1-k})|s_j(0) = h\right) = 1$ if $k = N - 1$ and 0 otherwise. That is, when the signal about product quality is perfectly informative, one affiliate's high signal about product 0 implies all the other affiliates also received high signal about product 0. When all other affiliates use a conformity strategy, they will all elect to recommend product 0. Then,

$$\mathbb{E}[U_j(0|s_j(0) = h), p_0 = 1] = \frac{1}{N} > \mathbb{E}[U_j(j|s_j(0) = h, s_j(i) = h), p_0 = 1] = 0. \tag{11}$$

Case 2. When affiliate j observes a low signal $s_j(0) = l$. When $p_0 = 1$, recommending product 0 yields affiliate j a lower expected payoff than recommending a non-salient product i:

$$\mathbb{E}[U_j(0|s_j(0) = l, p_0 = 1)] = 0 < \mathbb{E}[U_j(j|s_j(0) = l, s_j(i) = h, p_0 = 1)] = \frac{1}{N}. \tag{12}$$

Note that affiliate j's expected utilities following any strategy when all the other affiliates use the strategy $(\alpha_l, \beta_l) = (1, 0)$ for any $l \neq j$ are continuous in p_0 for $\frac{1}{2} < p_0 \leq 1$. Then, there exists some $\frac{1}{2} \leq \bar{p}(q, \lambda) < 1$ such that affiliate j finds no profitable deviation from $(\alpha_j, \beta_j) = (1, 0)$ when $\bar{p}(q, \lambda) < p_0 \leq 1$. This completes the proof of part (1).

(2) We show that $\bar{p}_0(q, \lambda) > \frac{1}{2}$ except at one point. We consider the limits of affiliate j's expected utilities as p_0 approaches $\frac{1}{2}$. When $p_0 > \frac{1}{2}$, product 0 needs $\frac{N}{2}$ recommendations for the consumer to consider first when N is even, and $\frac{N+1}{2}$ recommendations when N is odd.

Since affiliate j's expected utility from recommending a non-salient product is strictly increasing in q for any $\frac{1}{2} < p_0 < 1$ when $s_j(0) = h$, there exists $\hat{q}$ such that

$$\lim_{p_0 \to \left(\frac{1}{2}\right)^+} \mathbb{E}[U_j(0|s_j(0) = h)|q > \hat{q}] < \lim_{p_0 \to \left(\frac{1}{2}\right)^+} \mathbb{E}[U_j(j|s_j(0) = h, s_j(i) = h)|q > \hat{q}], \tag{13}$$

$$\lim_{p_0 \to \left(\frac{1}{2}\right)^+} \mathbb{E}[U_j(0|s_j(0) = l)|q < \hat{q}] > \lim_{p_0 \to \left(\frac{1}{2}\right)^+} \mathbb{E}[U_j(j|s_j(0) = l, s_j(i) = h)|q < \hat{q}]. \tag{14}$$

The reason that the same $\hat{q}$ can serve as the cutoff for both inequalities (13) and (14) is that as $p_0 \to \left(\frac{1}{2}\right)^+$, affiliate j's expected utility is independent from his signal about product 0. Therefore, when the probability of a second click by the consumer becomes sufficiently high, and when the signal informativeness becomes sufficiently low, and when all the other affiliates follow the strategy $(\alpha_k, \beta_k) = (1, 0)$, affiliate j finds it more profitable to recommend a non-salient product than product 0 after receiving a high signal about product 0. Therefore, for any given q (keeping λ and N fixed throughout) and any p_0, either $q > \hat{q}$ or $q \leq \hat{q}$ must hold. Then affiliate j finds either recommending product 0 more profitable regardless of his signal about product 0 or recommending a non-salient product more profitable regardless of his signal about product 0. Therefore, $(\alpha, \beta) = (1, 0)$ cannot be an equilibrium profile as $p_0 \to \left(\frac{1}{2}\right)^+$. This proves that $\bar{p}_0(q, \lambda) > \frac{1}{2}$ unless $q = \hat{q}$. It also follows that for any $q \neq \hat{q}$, there exists an interval $[\frac{1}{2}, \underline{p}_0(q, \lambda)]$ over which the profile $(\alpha, \beta) = (1, 0)$ does not constitute an equilibrium. This completes the proof of part (2). □

Proof of Corollary 1. (1) By the proof of Proposition 3, part (1), for any λ, as $p_0 \to \left(\frac{1}{2}\right)^+$, it is strictly more profitable for an affiliate to recommend product 0 when $q = 0$, and strictly more profitable to recommend a non-salient product when $q = 1$. By the compactness of the unit interval $[0, 1]$, the function $\lambda \mapsto \hat{q}(\lambda)$ is strictly bounded between 0 and 1. That is, there exist $\underline{q}$ and $\bar{q}$ with $0 < \underline{q} \leq \bar{q} < 1$ so that $\underline{q} \leq \hat{q}(\lambda) \leq \bar{q}$ for any $\lambda \in [0, 1]$.

When $q < \underline{q}$, $\bar{p}_0(q, \lambda)$ is determined by the indifference for an affiliate between recommending any non-salient product for which he received a high signal and product 0 when he received a low signal for it, whereas it is always more profitable to recommend product 0 after receiving a high signal for it for any $p_0 > \frac{1}{2}$. As λ increases, say, from λ_1 to $\lambda_1 + \epsilon$, recommending a non-salient product becomes more profitable: for any sufficiently small $\epsilon > 0$,

$$\begin{aligned}&\mathbb{E}[U_j(i|s_j(0) = l, s_j(i) = h, p_0 = \bar{p}_0(q, \lambda_1), \lambda = \lambda_1 + \epsilon)]\\ &> \mathbb{E}[U_j(0|s_j(0) = l, s_j(i) = h, p_0 = \bar{p}_0(q, \lambda_1), \lambda = \lambda_1 + \epsilon)].\end{aligned}$$

It follows that $\bar{p}_0(q, \lambda_1 + \epsilon) < \bar{p}_0(q, \lambda_1)$.

When $q > \bar{q}$, $\bar{p}_0(q, \lambda)$ is determined by the indifference for an affiliate between recommending any non-salient product for which he received a high signal and product 0 when he received a high signal for it, whereas it is always more profitable to recommend a non-salient after receiving a high signal for it for any $p_0 > \frac{1}{2}$. As λ increases, say, from λ_1 to $\lambda_1 + \epsilon$, recommending a non-salient product becomes more profitable: for any sufficiently small $\epsilon > 0$,

$$\begin{aligned}&\mathbb{E}[U_j(j|s_j(0) = h, s_j(i) = h, p_0 = \bar{p}_0(q, \lambda_1), \lambda = \lambda_1 + \epsilon)]\\ &> \mathbb{E}[U_j(0|s_j(0) = h, s_j(i) = h, p_0 = \bar{p}_0(q, \lambda_1), \lambda = \lambda_1 + \epsilon)].\end{aligned}$$

It follows that $\bar{p}_0(q, \lambda_1 + \epsilon) > \bar{p}_0(q, \lambda_1)$.

(2) The existence of $\underline{q}(\lambda)$ and $\bar{q}(\lambda)$ follows from the same compactness argument as in the proof of part (1). The expected utility from recommending a non-salient product is increasing with q, whereas the expected utility from recommending the salient product remains constant as q increases. When $q < \underline{q}(\lambda)$ and $p_0 = \bar{p}_0(q, \lambda)$, making a recommendation after receiving a low signal about the salient product is the binding condition for an affiliate. Then, $\bar{p}_0(q, \lambda) > \bar{p}_0(q + \epsilon, \lambda)$ for any sufficiently small $\epsilon > 0$. When $q > \bar{q}(\lambda)$, making a recommendation after receiving a high signal about the salient product is binding for an affiliate at $p_0 = \bar{p}_0(q, \lambda)$. Again, increasing q will only increase the expected utility for an affiliate from recommending a non-salient product. So as q increases to $q + \epsilon$ for some sufficiently small $\epsilon > 0$, $\bar{p}_0(q, \lambda)$ is weakly increasing. □

Proof of Proposition 4. (1) Let n_0 and n_1 be the unique numbers of recommendations the salient product receives in a recommendation profile so that

$$
\begin{aligned}
n_0 &:= \left\{1 \leq n \leq N | Pr(H|0^{(n)}) > p_0 \geq Pr(H|\tilde{l}, 0^{(n)})\right\} \\
n_1 &:= \min_{1 \leq n \leq N} \left\{Pr(H|\tilde{l}, 0^{(n)}) > p_0\right\}
\end{aligned}
$$

That is, n_0 is the threshold number of recommendations that will incentivize the consumer to first consider the salient product but will only purchase it when her signal is high, and n_1 is the threshold number of recommendations so that the consumer will purchase the salient product regardless of her signal. By definition, $n_0 < n_1$. Depending on the number of recommendations for the salient product the consumer sees, we consider the following three cases regarding the consumer's best response.

Case 1. $n \geq n_1$. This case leads the consumer to purchase the salient product regardless of her signal about it. The consumer's expected payoff is

$$(1-c)Pr(H|0^{(n)}). \tag{15}$$

Case 2. $n_0 \leq n < n_1$. The consumer will first consider the salient product. She will purchase a salient product only when her signal about it is high. Otherwise, she will continue to consider a non-salient product. The consumer's expected payoff is

$$
\begin{aligned}
&(1-c)Pr(\tilde{h}|0^{(n)}) \cdot Pr(H|\tilde{h}, 0^{(n)}) + q(1-2c)\left(1 - Pr(\tilde{h}|0^{(n)}) \cdot Pr(H|\tilde{h}, 0^{(n)})\right) \cdot Pr(\tilde{h}|h) \cdot Pr(H|\tilde{h}, h). \\
&= (1-c)\pi_0 \cdot Pr\left(H|0^{(n)}\right) + q(1-2c)\pi_0 p_0 \left(1 - \pi_0 \cdot Pr\left(H|0^{(n)}\right)\right) \\
&= [1 - p_0 q - c(1-2p_0 q)]\pi_0 Pr\left(H|0^{(n)}\right) + q(1-2c)p_0\pi_0.
\end{aligned}
$$

Case 3. $n < n_0$. In this case, the consumer will only consider non-salient products. She will only purchase it after seeing a high signal about the focal non-salient product. This is because when she sees a low signal about it, her posterior belief about the salient product is higher than that about the non-salient product under her consideration. After the consumer clicks on a link to a recommended non-salient product i, her belief about the clicked non-salient product's quality is updated as $Pr(H|\tilde{h}, h) = \frac{p_0\pi_0}{p_0\pi_0 + (1-p_0)(1-\pi_0)}$ or $Pr(H|\tilde{l}, h) = \frac{p_0(1-\pi_0)}{p_0(1-\pi_0)+(1-p_0)\pi_0}$. The consumer's expected payoff is:

$$
\begin{aligned}
&(1-c)\cdot Pr(\tilde{h}|h)\cdot Pr(H|\tilde{h},h) + q\cdot(1-2c)\cdot(1-Pr(\tilde{h}|h)\cdot Pr(H|\tilde{h},h))\cdot Pr(\tilde{h}|h)\cdot Pr(H|\tilde{h},h),\\
&=(1-c)p_0\pi_0 + q(1-2c)p_0\pi_0(1-p_0\pi_0). \qquad (16)
\end{aligned}
$$

The equality in (16) follows because

$$Pr(\tilde{h}|h) = Pr(H|h)Pr(\tilde{h}|H) + Pr(L|h)Pr(\tilde{h}|L) = p_0\pi_0 + (1-p_0)(1-\pi_0). \qquad (17)$$

Similar to the baseline model, the consumer's expected payoff attains maximum when her posterior belief about the salient product is the most accurate, which must be induced by a conformity profile.

(2) Denote by $p_0^+ := Pr(H|\tilde{h},h)$ and $p_n^+ := Pr(H|\tilde{h},0^{(n)})$ the consumer's posterior belief about a recommended non-salient product and the salient product after she clicks on a link and receives a high signal, where n is the number of recommendations she sees about product 0.

When $n \geq n_1$, the consumer will purchase the salient product after one click. The affiliates' aggregate expected payoff is

$$\lambda + (1-\lambda)\cdot Pr\left(H|0^{(n)}\right). \qquad (18)$$

When $n < n_1$, the affiliates' aggregate expected payoff is

$$\lambda\left[1+q(1-\max\{p_n^+,p_0^+\})\right] + (1-\lambda)\left[\max\{p_n^+,p_0^+\} + p_0q(1-\max\{p_n^+,p_0^+\})\right]. \qquad (19)$$

Comparing (18) and (19), we find that a conformity profile maximizes the affiliates' aggregate payoff if and only if either

(1) $1 - p_0q - (1+q-p_0q)\lambda > 0$; or

(2) $Pr\left(H|0^{(n)}\right) > \frac{\lambda}{1-\lambda}q(1-p_0) + p_0 + p_0q(1-p_0)$.

The first scenario is identical to the baseline model, yielding a lower bound on p_0: $p_0 < \frac{1-\lambda-q\lambda}{q(1-\lambda)}$. The second scenario leads to an upper bound on p_0. Note that the condition in the second scenario is not vacuous if and only if $\frac{\lambda}{1-\lambda}q(1-p_0)+p_0+p_0q(1-p_0)<1$, which gives rise to

$$\lambda < \frac{1-p_0q}{1+q-p_0q}. \tag{20}$$

Note that the condition above is the identical condition for scenario 1. Therefore, we identify the same condition that distinguishes the affiliates' aggregate preference for conformity as in the baseline model. In addition, within the case when the affiliates' aggregate payoff is maximized under conformity, we identify a further reason for their preference for conformity, which is when the consumer's posterior belief about product 0 is sufficiently high, the incentive for the affiliates to induce the consumer to purchase product 0 after one click leads to the greatest aggregate payoff for the affiliates.

For parts (3) and (4), the proofs in the baseline model on the existence of a non-conformity equilibrium and conditions for the existence of a conformity equilibrium still work. The reason is that, for checking the profitability of unilateral deviations, we only need to compare affiliates' aggregate payoffs for special sets of parameter values (e.g., $p_0 = 1$ in the case to characterize the condition for the existence of a conformity equilibrium). And those comparisons would not change under the current more general setting. □

Proof of Corollary 2. (1) In the model of experience goods, for every strategy profile for the affiliates, the consumer's best response is to click on a link to a product that has the highest posterior probability of being high quality and to purchase that product. (Because the consumer's posterior beliefs are identical for all of the recommended non-salient products, it is therefore unique up to identification of non-salient products that are recommended.) Therefore, the consumer's expected search length is 1, reducing $q = 0$. Also, the affiliates expect to be paid without distinguishing being clicked only or inducing a purchase. The affiliates' aggregate payoff is identical to the special case of the baseline model with $\lambda = 1$.

Therefore, to analyze the model of experience goods, it suffices to look at the special case with $q = 0$ and $\lambda = 1$. In this scenario, the condition identified in Proposition 1 becomes vacuous, yielding affiliates as a group indifferent between a conformity profile and non-conformity profile.

(2) Part (2) follows immediately as we have shown in part (1) the equivalence of the model of experience goods with the special model with $q = 0$ and $\lambda = 1$. □

Online appendix I: Preliminary analysis

In this section we cover a few observations that highlight the inner workings of the model. First, because each affiliate observes only a small fraction of all products, the probability that any two affiliates observe (and recommend) the same non-salient product is 0 and the only product that has positive probability of receiving multiple recommendations is the salient product (product 0). Therefore, affiliates' recommendations segment products into three categories:[23] (1) the salient product—a recognisable product that can possibly receive more than one recommendation; (2) non-salient products recommended by one affiliate; and (3) the remaining non-recommended non-salient products.

Second, because only a fraction of all products is observed by any affiliate, if a non-salient product does not receive a recommendation, the consumer assigns probability 1 for the product not being observed by any affiliate. As a result, the posterior belief of the consumer is that a non-salient product that is not recommended (category 3) is a high quality product is probability 1/2. Note that because all affiliates observe the salient product, the same is not true for an unrecommended salient product.

Third, for any $\lambda < 1$, an affiliate's expected payoffs increase in the probability that, conditional on clicking on the product he recommends, the consumer purchases the product. Because the consumer is able to evaluate the product before purchasing, the probability that a consumer purchases the product following a click is increasing in the quality of the product. As a result, as long as an affiliate receives a high signal for at least one non-salient product (which happens with probability 1), it is a weakly dominated strategy for him to recommend a non-salient product for which he received a low signal. For the same reason, in any equilibrium with weakly undominated strategies the posterior belief of the consumer is that a non-salient product that is recommended (category 2) is of high quality with probability p_0.

Fourth, the only undominated strategy for the consumer is to click on links to products in the order of their posterior probabilities to be of high quality (as long as the expected utility is higher than c), and buy a product she clicked on immediately if it is revealed to be of high quality. Moreover, for any $c > 0$ the consumer will never click on a link to a product that is not

[23] We abstract from the zero probability event in which a non-salient product receives more than one recommendation. Accommodating that possibility adds some special cases (all zero probability) to the analysis but doesn't introduce any new insights.

recommended by any affiliate.

Online appendix II: Further characterizations of equilibria

First, it is straightforward to verify that non-conformity always leads to an equilibrium.

Proposition OA.1. There always exists a pure-strategy equilibrium in which none of the affiliates ever recommend the salient product regardless of their signals about it and the consumer proceeds by only searching among non-salient products that are recommended.

Next, let $N > 2$ be the number of affiliates whose recommendations the consumer will observe. For any $q \in [0,1]$ and $\lambda \in [0,1]$, let $\overline{p}_0(q,\lambda)$ be identified by Proposition 3. That is, $\overline{p}_0(q,\lambda)$ is the smallest value so that the strategy profile with affiliates recommending according to $(\alpha,\beta) = (1,0)$ and the consumer searching in a posterior-belief-descending order can be sustained for any $p_0 \in (\overline{p}_0(q,\lambda), 1]$.

Proposition OA.2.

(1) For any $p_0 > \overline{p}_0(q,\lambda)$, the following profile is supported as a mixed-strategy equilibrium: the affiliates play according to $(\alpha(p_0), 0)$ for some uniquely determined $\alpha(p_0) \in (0,1)$, and the consumer proceeds her search in a posterior-belief-descending order.

(2) For any $p_0 > \overline{p}_0(q,\lambda)$, the following profile is supported as a mixed-strategy equilibrium: the affiliates play according to $(1, \beta(p_0))$ for some uniquely determined $\beta(p_0) \in (0,1)$, and the consumer proceeds her search in a posterior-belief-descending order.

Proof. (1) We first consider possible mix-strategy equilibria of the form $(\alpha, 0)$ with $0 < \alpha < 1$—every affiliate recommends the salient product with probability α after receiving a high signal about it and recommends a non-salient product after receiving a low signal about product 0, and the consumer proceeds with her search in a posterior-belief-decreasing order (among the recommended non-salient products, since the consumer's posterior belief about them are equal, when she decides to consider non-salient products, she will randomly pick one with equal probability).

Suppose that affiliate i receives a high signal about product 0: $s_i(0) = h$. His estimate of the probability that among the other $N-1$ affiliates s of them also receive high signals about product

0 is

$$Pr\left((h^s, l^{N-1-s})|s_i(0) = h\right) = \binom{N-1}{s}\left[p_0^{s+1}(1-p_0)^{N-1-s} + (1-p_0)^{s+1}p_0^{N-1-k}\right]. \tag{21}$$

Hence, after a high signal about product 0, affiliate i's expectation that product 0 will receive k recommendations from other affiliates is

$$\mathbb{E}_h\left[0^{(k)}\right] = \sum_{s=k}^{N-1}\binom{s}{k} Pr\left((h^s, l^{N-1-s})|s_i(0) = h\right)\alpha^k(1-\alpha)^{s-k}. \tag{22}$$

For consumer's posterior belief about product 0, conditional on product 0 having high quality, the probability that the consumer will see k recommendations of product 0 is

$$Pr\left(0^{(k)}|H\right) = \alpha^k\left[\sum_{i=0}^{N-k}\binom{N}{N-i}\binom{N-i}{k}p_0^{k+i}(1-p_0)^{N-k-i}(1-\alpha)^i\right]. \tag{23}$$

We can also derive the probability $Pr(0^{(k)}|L)$ that the consumer will see k recommendations of product 0 conditional on product 0 having low quality. Then the consumer's posterior about product 0 after seeing k recommendations is

$$\begin{aligned} Pr\left(H|0^{(k)}\right) &= \frac{Pr\left(0^{(k)}|H\right)}{Pr\left(0^{(k)}|H\right) + Pr\left(0^{(k)}|L\right)} \\ &= \frac{\sum_{i=0}^{N-k}\binom{N}{N-i}\binom{N-i}{k}p_0^{k+i}(1-p_0)^{N-k-i}(1-\alpha)^i}{\sum_{i=0}^{N-k}\binom{N}{N-i}\binom{N-i}{k}[p_0^{k+i}(1-p_0)^{N-k-i} + (1-p_0)^{k+i}p_0^{N-k-i}](1-\alpha)^i}. \end{aligned} \tag{24}$$

Similar to the pure-strategy case where $(\alpha, \beta) = (1, 0)$, we can find a unique $k_0(\alpha)$ such that product 0 needs to receive at least $k_0(\alpha)$ recommendations for the consumer to consider first. Since the consumer's search lasts at most two steps, she will not consider product 0 when more than one non-salient products receive recommendations and she did not start the search with product 0.

In the proposed mixed-strategy equilibrium with $0 < \alpha < 1$, after seeing a high signal about product 0, affiliate i must be indifferent between recommending product 0 and non-salient product

j for which he receives a high signal. The indifference condition is

$$\mathbb{E}[U_i(0|s_i(0) = h, p_0, \alpha)] = \mathbb{E}[U_i(j|s_i(0) = h, s_i(j) = h, p_0, \alpha)] \tag{25}$$

$$\Leftrightarrow \sum_{k=k_0(\alpha)-1}^{N-1} \frac{\lambda + (1-\lambda)Pr\left(H|0^{(k+1)}\right)}{k+1}\mathbb{E}_h\left[0^{(k)}\right] = \sum_{k=0}^{k_0(\alpha)-2} \Gamma(\lambda, q)\mathbb{E}_h\left[0^{(k)}\right] \tag{26}$$

where as before

$$\Gamma(\lambda, q) = \lambda\left(\frac{1}{N-k} + \frac{(1-p_0)q}{N-k-1}\right) + (1-\lambda)\left(\frac{p_0}{N-k} + \frac{(1-p_0)p_0 q}{N-k-1}\right). \tag{27}$$

To confirm the existence of a mixed-strategy equilibrium of the form $(\alpha, 0)$, it suffices to show that Eq. (26) has a solution in $\alpha \in (0,1)$ for any given set of parameters λ, q and N, and when p_0 is sufficiently high.

When $p_0 > \hat{p}_0(q, \lambda)$ and when all other affiliates deploy the strategy $(\alpha, \beta) = (1, 0)$, it is strictly more profitable for affiliate i to also play according to $(\alpha, \beta) = (1, 0)$, based on the previous analysis on the pure-strategy equilibrium $(\alpha, \beta) = (1, 0)$. Similarly, when $p_0 > \hat{p}_0(q, \lambda)$ and when all other affiliates deploy the strategy $(\alpha, \beta) = (0, 0)$, it is strictly more profitable for affiliate i to also play according to $(\alpha, \beta) = (0, 0)$ from the previous analysis on the pure-strategy equilibrium $(\alpha, \beta) = (0, 0)$. Therefore, there exists some $\alpha(p_0, q, \lambda) \in (0, 1)$ that equalizes an affiliate's expected utility from recommending product 0 and a non-salient product after receiving high signals about both when all the other affiliates play according to the mixed strategy $(\alpha(p_0, q, \lambda), 0)$.

Lastly, when affiliate i receives a low signal about product 0, and when all the other affiliates play according to the mixed strategy $(\alpha(p_0, q, \lambda), 0)$, it is strictly more profitable for affiliate i to recommend a non-salient product at $\alpha = \alpha(p_0, q, \lambda)$ when $p_0 > \hat{p}_0(q, \lambda)$. Therefore, the mixed-strategy profile of the form $(\alpha(p_0, q, \lambda), 0)$ gives rise to an equilibrium when the consumer searches in a posterior-belief-descending order.

(2) Consider the mix-strategy profile of the form $(1, \beta)$ with $0 < \beta < 1$. Note that after a high signal about product 0, affiliate i's expectation that product 0 will receive k recommendations from other affiliates is

$$\mathbb{E}_h\left[0^{(k)}\right] = \sum_{s=0}^{k} \binom{N-1-s}{k-s} Pr\left((h^s, l^{N-1-s})|s_i(0) = h\right) \beta^{k-s}(1-\beta)^{N-1-k}. \tag{28}$$

Conditional on product 0 having high quality, the probability that the consumer will see k recommendations of product 0 is

$$Pr\left(0^{(k)}|H\right) = (1-\beta)^{N-k}\left[\sum_{i=0}^{k}\binom{N}{i}\binom{N-i}{k-i}p_0^i(1-p_0)^{N-i}\beta^{k-i}\right]. \tag{29}$$

We can similarly derive the probability $Pr\left(0^{(k)}|L\right)$ that the consumer will see k recommendations of product 0 conditional on product 0 having low quality. Then the consumer's posterior about product 0 after seeing k recommendations is

$$\begin{aligned} Pr\left(H|0^{(k)}\right) &= \frac{Pr\left(0^{(k)}|H\right)}{Pr\left(0^{(k)}|H\right)+Pr\left(0^{(k)}|L\right)} \\ &= \frac{\sum_{i=0}^{k}\binom{N}{i}\binom{N-i}{k-i}p_0^i(1-p_0)^{N-i}\beta^{k-i}}{\sum_{i=0}^{k}\binom{N}{i}\binom{N-i}{k-i}[p_0^i(1-p_0)^{N-i}+(1-p_0)^i p_0^{N-i}]\beta^{k-i}}. \end{aligned} \tag{30}$$

We can find unique $k_0(\beta) < k_0'(\beta)$ such that product 0 needs to receive at least $k_0(\beta)$ and at most $k_0'(\beta)$ recommendations for the consumer to consider first.

After seeing a low signal about product 0, affiliate i is necessarily indifferent between recommending product 0 and a non-salient product j for which he receives a high signal. The indifference condition is

$$\mathbb{E}[U_i(0|s_i(0)=l,p_0,\beta)] = \mathbb{E}[U_i(j|s_i(0)=l,s_i(j)=h,p_0,\beta)] \tag{31}$$

$$\Leftrightarrow \sum_{k=k_0(\beta)-1}^{k_0'(\beta)-1}\frac{\lambda+(1-\lambda)Pr\left(H|0^{(k+1)}\right)}{k+1}\mathbb{E}_h\left[0^{(k)}\right] = \sum_{k<k_0(\beta)-1 \text{ or } k\geq k_0'(\beta)}\Gamma(\lambda,q)\mathbb{E}_h\left[0^{(k)}\right] \tag{32}$$

where as before

$$\Gamma(\lambda,q) = \lambda\left(\frac{1}{N-k}+\frac{(1-p_0)q}{N-k-1}\right)+(1-\lambda)\left(\frac{p_0}{N-k}+\frac{(1-p_0)p_0q}{N-k-1}\right). \tag{33}$$

To confirm the existence of a mixed-strategy equilibrium of the form $(1,\beta)$, it suffices to show that Eq. (32) has a solution in $\beta \in (0,1)$ for any given set of parameters λ and q, and when p_0 is sufficiently high.

When all the other affiliates play according to the strategy $(\alpha,\beta)=(1,1)$, when p_0 is sufficiently high, it is strictly more profitable for affiliate i to also play according to $(\alpha,\beta)=(1,1)$. Then there exists some $\beta(p_0, q, \lambda) \in (0, 1)$ that equalizes affiliate i's expected utility from recommending product 0 after a low signal about it and a non-salient product for which he receives a high signal, when all other affiliates also play according to $(\alpha, \beta) = (1, \beta(p_0, q, \lambda))$:

$$\mathbb{E}[U_i(0|s_i(0) = l, \beta = \beta(p_0, q, \lambda))] = \mathbb{E}[U_i(j|s_i(0) = l, s_i(j) = h, \beta = \beta(p_0, q, \lambda))]. \tag{34}$$

This also shows that when affiliate i receives a high signal about product 0, it is strictly more profitable for him to recommend product 0. This completes the proof of the existence of a mixed-strategy equilibrium in which affiliates play according to $(\alpha, \beta) = (1, \beta(p_0, q, \lambda))$ and the consumer searches in a posterior-belief-descending order. $\square$