Mega or Micro?
Influencer Selection Using Follower Elasticity

Zijun Tian[*,†], Ryan Dew[†], Raghuram Iyengar
University of Pennsylvania

September 21, 2023

Abstract

Influencer marketing, in which companies sponsor social media personalities to promote their brands, has exploded in popularity in recent years. One common criterion for selecting an influencer partner is popularity. While some firms collaborate with "mega" influencers with millions of followers, other firms partner with "micro" influencers with only several thousand followers, but who also cost less to sponsor. To quantify this trade-off between popularity and cost, we develop a framework for estimating the *follower elasticity of impressions*, or FEI, which measures a video's percentage gain in impressions (i.e., views) corresponding to a percentage increase in the number of followers of its creator. Computing FEI involves estimating the causal effect of an influencer's popularity on the view counts of their videos, which we achieve through a combination of: (1) a unique dataset collected from TikTok, (2) a representation learning model for quantifying video content, and (3) a machine learning-based causal inference method. We find that FEI is always positive, averaging 0.10, but often nonlinearly related to follower size. We examine the factors that predict variation in these FEI curves and show how firms can use these results to better determine influencer partnerships.

Keywords: influencer marketing, causal inference, deep learning, representation learning, heterogeneous treatment effects, video data

[*]Zijun Tian is an Assistant Professor of Marketing at Olin Business School, Washington University in St. Louis. Ryan Dew is an Assistant Professor of Marketing at the Wharton School of the University of Pennsylvania. Raghuram Iyengar is the Miers-Busch, W'1885 Professor, Professor of Marketing at the Wharton School of the University of Pennsylvania. Zijun Tian is the corresponding author: `ztian@wustl.edu`

[†]Equal contribution.

1. Introduction

Influencer marketing has become an integral part of advertising strategy as companies recognize the importance of collaborating with social media personalities (e.g., Hughes, Swaminathan, and Brooks 2019; Rajaram and Manchanda 2023; Wies, Bleier, and Edeling 2023). In its most basic form, companies pay influencers to share content on social platforms with the twofold goal of reaching a target audience and keeping them engaged (e.g., Leung, Gu, and Palmatier 2022). Although there are many outlets for such sponsorships, TikTok has emerged as the most popular.[1] It is a video-focused social networking service and is largely attributed with popularizing the short-form user video. With typical lengths ranging from 15 to 180 seconds, these videos are ideal for holding the attention of an audience with a shrinking attention span. Recent statistics indicate that TikTok has more than 150 million users in the United States and nearly a third of this installed base uses the app everyday for an average of just under 100 minutes per day.[2]

Given the large audience on social media, companies are exploring different ways to engage with them. To this end, many firms have been promoting *challenges*. For example, a typical challenge on TikTok contains a name in the form of a hashtag (e.g., #MakeMomSmile from Colgate), a sentence or two encouraging users to create content that matches a theme (e.g., "Make Mom Smile this Mother's Day by doing something special for mothers"), and a few user-generated videos, sometimes sponsored by the company, that can jumpstart the campaign.[3] Figure 1 shows examples of sponsored posts from two challenges on TikTok. The left panel shows a fashion influencer who created a dancing video for Dettol's #HandWashChallenge that aimed to promote its hand-washing products. The right panel shows a family influencer who participated in Walmart's #UnwrapTheDeals campaign, which featured a special effect animation which TikTok users could add to their videos to advertise Walmart's Black Friday deals. Other channels such as Instagram and Snap also encourage companies to offer challenges with the use of pictures and filters.

As companies continue to invest in this new form of marketing, there is a growing need to assess the value of sponsoring different types of influencers to promote their campaigns. An ongoing debate is the importance of influencers with many followers (e.g., so-called "mega influencers," with millions of followers) as best suited for promoting a campaign, as compared to those with smaller followings (e.g., "micro influencers," with thousands of followers). Although the former

[1] https://www.emarketer.com/content/tiktok-influencer-marketing

[2] https://wallaroomedia.com/blog/social-media/tiktok-statistics/

[3] https://www.tiktok.com/tag/MakeMomSmile

have a larger potential audience, they are typically more expensive than the latter.[4] There is also anecdotal evidence suggesting that mega influencers are good for awareness but not so much at later stages of the purchase funnel.[5] We contribute to the discussion around the effectiveness of influencer marketing by offering a framework for advertisers to assess how the impressions of an influencer's post would change with respect to their follower size.

Estimating the relationship between the impressions of an influencer's post and their follower size poses a few challenges. First, the content of an influencer's videos affects both their number of followers and the popularity of their videos. Consider the two influencers in Figure 1: the influencer in the right panel of the figure tends to post content related to families, which may have broad appeal on TikTok, while the influencer in the left panel of the figure typically posts about fashion, which may have more niche appeal. Without accounting for the differences in content, one may erroneously infer that follower size, rather than content, is driving the popularity of the videos. However, statistically controlling for the content of a video is not straightforward: video data are unstructured, containing different modalities (images, audio, text) that work in synergy to convey the essence of the content. Second, the impact of follower size on video impressions may be nonlinear, and may differ by the type of content. For example, there may generally be decreasing returns to popularity, although a social challenge may benefit more from each additional follower than an instructional fashion video. Finally, other unobserved confounders may be present that drive both an influencer's popularity, and the impression counts of their videos. As an example, suppose that the fashion influencer posts their videos simultaneously on multiple social media channels, while the family influencer does not. This behavior, termed cross-posting, may generate both more followers and more views for the first influencer, yet may also be difficult to systematically measure. Thus, to draw valid conclusions about the relationship between number of followers and video impressions, we need to use a method that is robust to potential unobserved variables, or *confounds*, that drive both followers and impressions.

In this paper, we develop a modeling framework for determining the causal relationship between followers and impressions that addresses all of the above issues. We do so through two components. First, to summarize the content of a video, we employ a representation learning model based on the variational autoencoder (VAE). This part of our model takes as its inputs all the observed information about a post (e.g., video, hashtags) and outputs a vector representation

[4] https://www.businessofapps.com/marketplace/influencer-marketing/research/influencer-marketing-costs/

[5] https://shanebarker.com/blog/influencer-marketing-statistics/

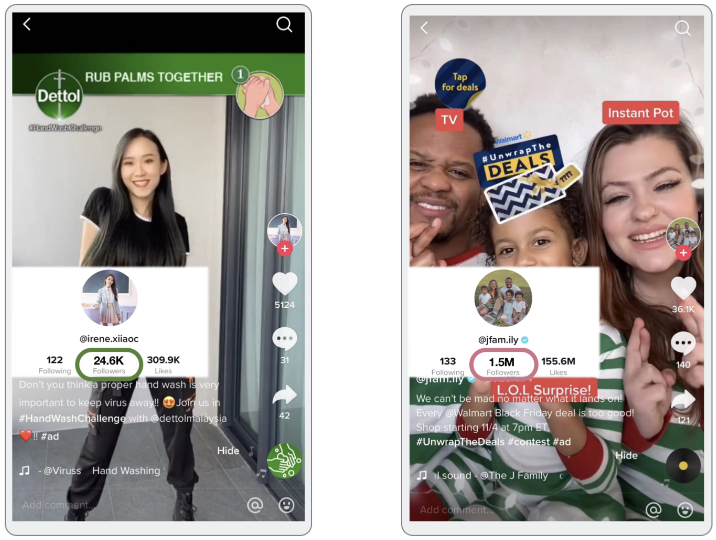

Figure 1: Examples of Sponsored Videos on TikTok
At left, a screenshot of a video posted by a mid-sized fashion influencer under Dettol's #HandWashChallenge. At right, a screenshot of a video posted by a mega influencer under Walmart's #UnwrapTheDeals campaign.

of the content. Second, conditional on the learned representation and other metadata about the creator and post, we leverage a machine learning-based causal inference framework called Deep Instrumental Variables (Hartford et al. 2017). A key benefit of using the Deep IV approach is that it allows for flexible estimation of a nonlinear relationship between the number of followers of an influencer and impressions of their posts. As implied by its name, Deep IV requires an instrument, or a variable that is correlated with a video's impressions only through its impact on an influencer's follower size, to infer the causal link between the two. Importantly, our representation learning framework allows us to reason about the similarity of influencer posts, which plays a central role in the construction of such an instrument.

We estimate our model using a dataset of TikTok videos that appeared on TikTok's Discover page over a period of six months. The Discover page on TikTok shows users a wide variety of videos grouped by categories that are currently trending in the TikTok community. The page is updated daily with 1-2 new hashtags, giving us a broad snapshot of popular content on TikTok. Our dataset consists of more than 200 hashtags, and just over half a million videos that are publicly accessible under those hashtags. In addition to the videos themselves, our data also include metadata about the videos and their creators, allowing us to record our two variables of interest: the follower count of an influencer at the time a video was created, and the impression count of a video two weeks after its posting. Finally, TikTok provides an ideal setting for our analysis because it has revealed in public posts aspects of how its recommendation algorithm works. As we explain later, this information helps us to create a valid instrument that we can use in conjunction with the Deep IV approach.

Using our modeling framework, we derive a key quantity of managerial interest: the follower elasticity of impressions (FEI). Our proposed metric is inspired by research on advertising elasticity (e.g., Danaher and Dagger 2013) and recent work on the impact of marketing spend on customer engagement in influencer campaigns (Leung et al. 2022). For a given video, the FEI captures the expected percentage change in that video's impressions corresponding to a unit percentage change in its creator's follower size. As influencers are typically paid based on the size of their follower base, and impressions are one of the key metrics of interest in digital campaigns, FEI can inform firms of the effectiveness of their influencer marketing.

Estimating FEI for posts in our TikTok data yields three key results. First, the FEI, averaged across all videos, is positive and nonlinearly related to influencer popularity. Specifically, the average FEI curve, which captures how FEI varies as a function of followers, has an inverted U-shape. This shape suggests that, on average, mid-tier influencers have the highest gain in impressions for each incremental follower. When translated into the expected number of impressions for a video, this FEI pattern provides evidence for an S-shaped response curve, similar to the S-shaped response curve often postulated (but not well documented) in the advertising literature (e.g., Johansson 1979; Jones 1995; Cannon, Leckenby, and Abernethy 2002). Interestingly, a descriptive analysis that does not control for any confounders points to the opposite pattern, suggesting that the biggest returns from additional followers come from either very small or very large influencers. Thus, firms may erroneously amplify the importance of mega influencers and make suboptimal decisions about sponsorships, when video content and other confounders are not controlled for. Second, FEI curves systematically vary with how the firm is trying to engage customers (e.g., by entertaining them or encouraging them to socialize with each other) and the topics covered in the video (e.g., food or gaming). Thus, a firm should collaborate with influencers with different levels of popularity when they have different goals, which is in line with industry practice (albeit done in an ad hoc manner based on intuition).[6] Third, counterfactual predictions for how video impressions will grow based on follower size for sponsored campaigns show a spectrum of different growth curves, ranging from S-shaped, to purely concave, to almost

[6]https://advertisingweek.com/how-to-pick-the-right-influencers-for-your-campaign-to-ensure-success-and-avoid-wasted-spend/

linear. The linear growth curve is noteworthy as it suggests that there is merit to the notion that mega-influencers can expand the number of impressions for some, but not all, campaigns.

The remainder of the paper is organized as follows. First, we review past work related to influencer marketing. Next, we describe our curated TikTok dataset. Then we describe our modeling framework, including representation learning and the causal inference specification. After that, we present our results and their implications for both theory and practice. Finally, we conclude with a discussion of the main contributions, limitations of our model, and ideas for future research.

2. Related Literature

Our paper is related to three streams of literature: social networks and influencer marketing; advertising; and causal inference with heterogeneous treatment effects and unstructured data. In this section, we discuss how our work complements previous research in each of these areas.

2.1. Social Networks and Influencer Marketing

The seminal work of Granovetter (1973) has served as an inspiration for researchers in fields such as marketing, economics, and sociology on how the structure of social networks can impact downstream micro- and market-level outcomes (e.g., Katona, Zubcsek, and Sarvary 2011; Susarla, Oh, and Tan 2012; Liu-Thompkins 2012; Lee, Hosanagar, and Nair 2018). Of direct relevance is work that has explored the relationship between the number of followers of an individual (their in-degree) and the diffusion of new products. For example, in the context of YouTube, Yoganarasimhan (2012) finds that the number of first- and second-degree connections of a content creator is positively related to the total viewership of their videos. A few recent studies have found a positive relationship between in-degree and engagement with content on social networks (e.g., Hughes, Swaminathan, and Brooks 2019; Valsesia, Proserpio, and Nunes 2020).

One aspect of the above studies worth noting is that their applications involve creators with relatively small numbers of followers (e.g., Yoganarasimhan (2012) has a maximum of $1,000$ followers). In the context of influencer marketing, there is large variation in follower size: some influencers have millions of followers (mega-influencers), while others have only a few thousand (micro-influencers). Such a broad range enables the test of more nuanced hypotheses that are relevant for marketers. In this vein, Wies, Bleier, and Edeling (2023) is a notable example of recent work that does so. The authors find that the in-degree of a creator has an inverted-U shape impact on engagement with their content. In particular, in their proposed conceptual model, it is assumed that the reach of an influencer is linearly related to followers, while engagement (conditional on reach) may decrease due to a weakening of average tie strength (e.g., Katona, Zubcsek, and Sarvary 2011). The notion that there are differences between micro- and mega-influencers on various outcomes resonates with findings in other contexts. In the medical literature, for example, the importance of local leaders as opposed to national opinion leaders is well documented (e.g., Kuo, Gifford, and Stein 1998; Doumit et al. 2007; Keating et al. 2007).

We contribute to this stream of research along a few dimensions. First, rather than assuming linearity, our model specification flexibly captures the relationship between the number of followers of a creator and the number of impressions garnered by their content, which past work and anecdotal evidence suggest is warranted. For example, followers might not be equally motivated to view and engage with an influencer's content (Farrell, Campbell, and Sands 2022), leading to differences in how often users seek out and share an influencer's content, depending on the influencer's popularity. Anecdotal evidence suggests that smaller influencers are viewed as more authentic and relatable,[7] which, when combined with what Wies, Bleier, and Edeling (2023) find in their paper, suggests that their followers may be more likely to interact with and share their content, leading to a higher marginal effect of followers for small versus large influencers. The flexibility of our modeling framework allows us to let the data speak and assess the nature of this relationship.

Second, we investigate whether the content of the post moderates the above relationship. Anecdotal evidence suggests that topics (e.g., food or gaming) have different popularities on social media. For example, on TikTok, food videos are less viewed than videos that highlight skills.[8] However, the role of influencer size in generating viewership for different topics is not well understood. Some types of content — for example, gaming — are inherently more social than others. In such cases, we might imagine that the popularity of the influencer matters more for generating engagement. It is important for companies to be aware of any synergy between the type of content and the popularity of the influencer for generating views, and identifying such synergies is a key benefit of our empirical framework.

[7] https://www.businessinsider.com/brands-turning-to-micro-influencers-instead-of-instagram-stars-2019-4

[8] https://www.statista.com/statistics/1130988/most-popular-categories-tiktok-worldwide-hashtag-views/

Third, we assess whether the manner in which firms engage with their customers (e.g., informational versus emotional) moderates the relationship between number of followers and impressions. A growing stream of research suggests that emotional appeals are more effective in engaging consumers compared to informational appeals (Berger and Milkman 2012; Akpinar and Berger 2017). With that said, recent surveys suggest that people increasingly prefer to search on social media for informational content rather than using traditional search engines (e.g., the how-to videos on TikTok).[9] Just like with content, our model allows us to understand how the nature of a firm's appeal may moderate the link between popularity and impressions. In particular, we classify appeals based on whether they are entertaining, socializing (two main forms of emotional appeal), or informative (Dolan et al. 2019), then study how an influencer's popularity differentially matters across these types, adding evidence to the importance of considering appeal type when deciding influencer sponsorships.

2.2. Advertising

Our work is fundamentally an effort to quantify the impact of popularity on the effectiveness of influencer marketing, which can be seen broadly as a form of social media advertising. In this sense, our study also connects to the long literature on quantifying the impact of advertising on sales (e.g., Dekimpe and Hanssens 2007; Shapiro, Hitsch, and Tuchman 2021). For instance, Lodish et al. (1995) analyze data from 389 split cable experiments to assess how different aspects of TV advertising (ad copy and media budget) impact sales. More recently, Shapiro, Hitsch, and Tuchman (2021) find substantially smaller effects of own-advertising compared to the results documented in the extant literature, as well as a sizable percentage of statistically insignificant or negative estimates. In a similar vein, Gordon et al. (2019) assess the causal impact of Facebook advertising in 15 experiments and find that it has a differential impact based on the purchase funnel. Hanssens (2018) notes that empirical generalizations are critical for elevating marketing from a cost center to an investment function. Another stream of research has focused on identifying the functional form of the advertising response curve (e.g., Rao and Miller 1975; Simon and Arndt 1980; Vakratsas et al. 2004). An ongoing (and not satisfactorily resolved) debate is whether the advertising response curve is S-shaped or concave (as discussed in Cannon, Leckenby, and Abernethy (2002)) and whether there is a threshold effect (Jones 1995). Some work finds that there is a threshold effect, but only for certain product categories (e.g., Vakratsas et al. 2004).

[9] https://blog.hootsuite.com/how-to-get-more-views-on-tiktok/

We contribute to both streams of research. We generate the follower elasticity of impressions (FEI) curve associated with a focal video and then assess if there are any generalizable patterns across videos. In doing so, we highlight the similarities as well as the differences between influencer marketing and traditional TV advertising. We find that the average FEI is 0.10, close to the short-term brand advertising elasticity of 0.12 noted in Sethuraman, Tellis, and Briesch (2011)'s meta-analysis, which focused on TV advertising. An important caveat with this comparison is that we have access only to the number of impressions of a campaign video and not to the sales associated with it. Our comparison is reasonable when the sales from an influencer campaign are proportional to the impressions. As a point of difference, we find that some influencer campaigns show a linear growth curve for the number of impressions, which is rare to find in the context of traditional advertising. Finally, just as ad elasticities give firms a tool for assessing the cost effectiveness of their traditional advertising (Danaher and Dagger 2013), our method for computing follower elasticity is general and provides a tool for firms to think about the cost effectiveness of their influencer campaigns.

2.3. Unstructured Data and Causal Analysis

The third stream of literature to which our work contributes is analysis of unstructured data, particularly from a causal perspective. Video analysis has emerged as a promising tool for understanding the link between the content of video marketing campaigns and firm-relevant outcomes (e.g., Li, Shi, and Wang 2019; Yang, Zhang, and Zhang 2021). Our approach considers how different modalities (audio, video) interact with each other to drive an outcome and is thus similar in spirit to other recent studies (e.g., Rajaram and Manchanda 2023). The way we address the problem, however, is different: we leverage a multimodal representation learning framework, which allows us to embed videos in a lower-dimensional vector space. Our learning framework is based on the variational autoencoder, first introduced by Kingma and Welling (2013), and built on recently in marketing by Dew, Ansari, and Toubia (2022) and Burnap, Hauser, and Timoshenko (2023). Our model is similar in structure to that proposed in Dew, Ansari, and Toubia (2022), but differs in the way multimodal features are processed. Our work uses modality-specific transfer learning to process raw video, audio, and text data, rather than purely hand-crafted features, and uses a neural network architecture to merge these domains together.

The second component of our modeling framework, causal machine learning, has gained significant attention in recent years. An example of such a method is the causal forest, which has been used to infer causality, and, in particular, heterogeneous treatment effects, in secondary data contexts (Guo, Sriram, and Manchanda 2021; Zhang and Luo 2023). A key underlying assumption while applying this framework is that of unconfoundedness; that is, that there are no unobserved variables that affect both the treatment and the outcome. This assumption is reasonable in contexts where there are many control variables or where institutional details indicate that unobserved confounders are unlikely. However, in our setting, unobserved confounders are of great concern. We therefore adopt an alternative causal inference framework based on machine learning, Deep IV (Hartford et al. 2017). The use of Deep IV also allows for the estimation of a flexible response curve for how a continuous variable (e.g., follower size) affects a continuous outcome (e.g., the number of impressions).

Our modeling framework uniquely combines the two components, namely, VAE for addressing the unstructured nature of videos and deep IV for inferring the casual impact of number of followers on video impressions. In this way, our model captures the unbiased, nonlinear, and context-dependent causal effect of followers on the number of TikTok video impressions. Our representation learning framework also facilitates the development of an instrumental variable to use with the Deep IV specification. While we focus on impressions as the outcome of interest, the proposed framework is general and can accommodate other outcome variables (e.g., likes).

3. Data

The goal of our research is to determine the causal relationship between the follower size of an influencer and their video impressions, with an eye to guiding firms in selecting an ideal level of popularity for their influencer partners. To do so, our context and data should satisfy several conditions: First and foremost, posts should be collected from a platform where we can gather data on the content of posts and where a valid instrumental variable (IV) can be leveraged to infer the focal relationship. Second, the posts should span a wide range of follower sizes and content types, from which we can identify any potential heterogeneity in the relationship between the popularity of the influencer and the popularity of a video. Third, the growth curve for impressions of posts should be tracked from their introduction (no left censoring) to maturity. Finally, the platform should allow brands to collaborate with influencers, and the data should contain an identifier for sponsored videos. TikTok provides an ideal empirical context for studying influencer marketing, as it not only meets all of these criteria but is also one of the most popular social media platforms on which firms sponsor influencers.

3.1. Data Collection

We curated a dataset with all videos, both sponsored and organic, that ever appeared on the public Discover page of the American version of TikTok over a period of six months (from Oct 2020 to April 2021). Our choice of TikTok as the social media platform and the Discover page, in particular, was motivated by how users are exposed to posts: compared to other browsing channels, the Discover page on TikTok contains videos grouped by categories, termed hashtags, that are not personalized to any single user's content preferences. The Discover page also updates daily with 1-2 new hashtags, leading to a variation in the observed content. Moreover, the Discover page hashtags reflect topics that are currently trending in the TikTok community, and include content creators whose follower size ranges from several hundreds to several millions. Our final dataset consists of 216 hashtags, 30 of them sponsored, and just over half a million videos that were publicly accessible under those hashtags.[10]

For each hashtag that appeared on the Discover page during the data period, we collected key characteristics for every video displayed under that hashtag, including its content, the caption, and metadata capturing information about the video's creator and the number of views of the video. It is the metadata about creators, specifically the average number of hearts (TikTok's equivalent of "likes" on Facebook) that their videos receive, that allows us to form our instrumental variable (described later). As some videos might contain multiple hashtags, we tracked each video's earliest appearance on the Discover page.[11] Importantly, we continued to track the number of views for each video on a daily basis for at least two weeks after its first appearance. This approach allows us to measure, first, how many followers the video's creator had on the day of the video's posting, and second, how many impressions the video attained subsequently.[12] Moreover, since the Discover page generally features *new* hashtags, our panel of observations consists of new videos, for which we can track the entire history. Our two-week period of observation accounts

[10]During our observation window, there were a total of 58 sponsored hashtags. However, a few sponsored campaigns on TikTok have shorter durations (e.g., several days), which made it impossible for us to apply our tracking procedure. The 30 hashtags that we include in our analysis are those that we were able to track over a period of at least two weeks, which captures more than 90% of a video's one-month views.

[11]Since we collect all hashtags and all videos associated with them in chronological order, the earliest appearance would capture a video's best possible position to get exposure on the Discover page (i.e., either under the most trending hashtag or with the highest ranking if the related hashtags are equally trending).

[12]We use the terms views and impressions interchangeably.

	Mean	Std. Dev.	Min	Median	Max
Followers	486,169	1,335,649	2	60,102	71,000,679
Impressions	1,038,728	2,468,832	109	4,444,500	165,200,000

Table 1: Summary Statistics of the Data
Summary statistics of our key independent variable, follower size ("Followers"), and key dependent variable, impressions ("Impressions")

for more than 90% of a video's one-month views.[13]

TikTok posts are multifaceted and differ from each other in many ways, including the imagery in the video, the background audio, edits and effects added to the video, and the caption under the video. We refer to this set of features, which all fall under the umbrella of unstructured data, as the content of a post. The other features that we collect about a post and its creator we call the post and creator statistics, which include measures like how many hashtags were used, whether those hashtags were trending or not, and how many videos a creator has posted in the past.[14]

3.2. Descriptive Statistics and Exploratory Analyses

In our analysis, the main outcome of interest is the total number of impressions for a video two weeks after its posting. The focal explanatory variable is the content creator's follower size when the video is posted. Table 1 presents the summary statistics for these variables, calculated across videos in our data. The follower size variable shows considerable variation, suggesting that our sample contains a wide range of influencers. The same is the case for the number of impressions, suggesting that our data contains a few very popular videos. Due to the right-skewness of both variables, we use their log transformations in our analysis. Figure 2 plots their marginal distributions.

Before describing our framework for inferring the causal relationship between followers and impressions, it is worthwhile to consider any correlational patterns that exist in the data. Figure 3 shows the relationship between (logged) follower size and (logged) video impressions, in terms of both the raw scatterplot, and a smoothed estimate of the relationship between the two. We see that on average there is a nonlinear relationship between follower size and number of impressions: influencers with either very small or very large number of followers seem to generate the most significant growth in impressions. A priori, this finding seems reasonable, as mega influencers

[13] See Web Appendix A.1 for more details on how soon videos gain most of their impressions on TikTok and why we choose a period of observation of two weeks.

[14] See Web Appendices A.2–A.3 for more details.

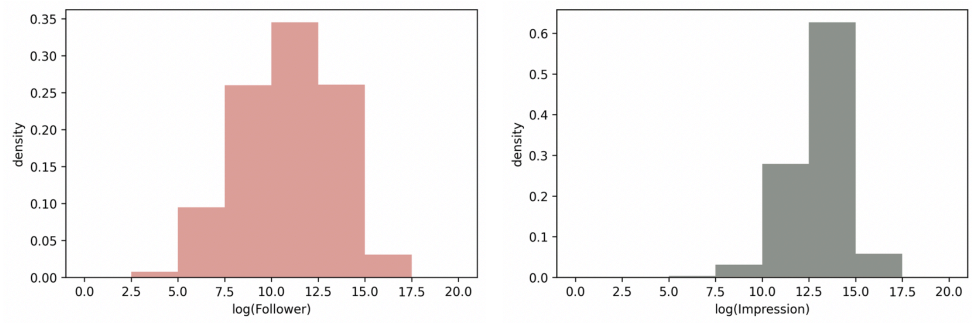

Figure 2: Marginal Distributions
Histograms of the same variables from Table 1, (logged) followers and impressions.

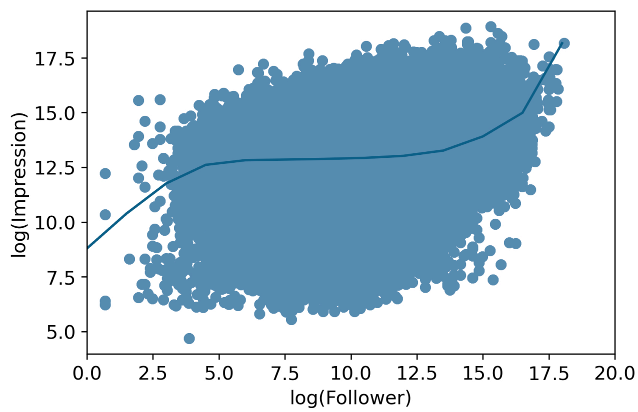

Figure 3: (Correlational) Relationship Between Followers and Impressions
The joint distribution of logged followers and logged video impressions in our data, along with a smoothed estimate of the relationship.

are extremely popular and often a target for firm partnerships. However, this result is misleading: as we show later, once we control for content and the possibility of unobserved confounders, the effect in Figure 3 changes, such that the highest marginal returns are for mid-tier influencers. This finding also suggests that, as we develop our framework to more accurately characterize this relationship, we should ensure that our model is sufficiently flexible to capture potentially nonlinear relationships.

4. Methodology

To begin, we give an overview of our empirical approach, which is rooted in causal inference. The question we seek to address is: how does an influencer's popularity affect the impressions of their videos? To better understand why this is a causal question, it is helpful to imagine the "experimental ideal" that our work seeks to emulate with observational data. The experimental ideal imagines that a company knows the exact type of influencer with which it wants to collaborate, and can decide to pay for an arbitrary number of followers for that influencer. Given the ability to pay for an arbitrary number of followers, the firm could then measure the effect a change in followers has on the impressions of their sponsored videos, to determine an optimal popularity. Furthermore, if this experiment were possible, the firm could run many such experiments to determine how the number of impressions varies as a function of the number of followers, conditional on different types of influencers, campaigns, and content.

Of course, such an experiment is *not* possible. In most cases, companies cannot arbitrarily endow influencers with a given number of followers and measure the resulting impressions.[15] However, in reality, the firm's decision is a close proxy of that imagined scenario: for a given campaign, there are typically many influencers with whom the company could collaborate, who all create similar content, and who all match the company's desired influencer profile, but who vary in their follower counts. Thus, under an assumption of exchangeability of the influencers, the counterfactual question of interest remains the same: if the firm paid for more followers, by selecting a more popular but otherwise similar influencer, how would that affect the performance of their campaign?

To accurately determine the effect of follower count on impressions, using purely observational data, our analysis must address four key issues:

1. **Nonlinear treatment effects:** Our discussion about prior findings, and our exploratory analysis in Figure 3 both suggest that the relationship between followers and impressions may be nonlinear. Thus, we must accordingly model that relationship in a flexible way.

2. **The key confound of content:** The content of an influencer's videos may affect both how popular that influencer is and how many views their posts receive. In the language of causal inference, a variable that affects both the treatment (i.e., followers) and the outcome (i.e., impressions) is termed a confound. As an illustrative, simple example, consider cat videos: because of widespread interest in cats, cat videos are likely to receive many views on TikTok, and creators who create those cat videos may end up with more followers. Since content can determine the popularity of both the influencer and the video, our causal analysis must carefully control for it.

[15] The only possible exception where endowing an influencer with followers is with the use of fake followers, as in Toubia and Stephen (2013), or in artificial social networks, as in Centola (2010).

3. **Heterogeneity in treatment effects:** The causal effect of followers on impressions may differ depending on the type of content. Imagine, for instance, a campaign that is informational, like a beauty product demonstration, versus a campaign that is inherently social, like TikTok's "duet" style videos, that require one user to build on another's video. We might expect the latter to depend much more on having a very popular influencer post the video than the former, given its inherently social nature. Therefore, to decide which creators to partner with, firms must assess the relationship between followers and impressions *for the specific type of content they want to produce*.

4. **Unobserved confounders:** Other potential confounders beyond video content may arise in estimating the relationship between followers and impressions. A subtle but important example is the practice of influencers cross-posting content: many influencers have accounts across multiple social networks, and it is common practice to use one social network to promote their content on another network.[16] In these scenarios, posting on the outside platform may lead to both an increase in followers for the cross-posting creator and an increase in views for the particular video being advertised. Importantly, such cross-posting behavior is difficult to observe, especially at scale: it is not uncommon for creators to use slightly different usernames across platforms, which, when combined with various platform restrictions on data scraping, makes it very difficult to systematically gather and match content across platforms. In general, unobserved confounders like cross-posting can lead to biased estimates of how valuable each follower is, as the estimate of the effect of popularity may be conflated with these unobserved factors.

We address all four issues with a modeling framework based on two components: first, to summarize the content of the video, we propose a representation learning model for TikTok posts. This model takes TikTok posts as its inputs, and outputs a dense vector representation of the post's content, which, for the remainder of this section, we will denote $\boldsymbol{r}$. Second, conditional on $\boldsymbol{r}$ and other observables about the creator and post, we leverage a machine learning-based causal inference framework to model the relationship between followers and impressions (Hartford et al. 2017). This framework uses instrumental variables to address the potential issue of unobserved confounders.

[16] https://blog.hubspot.com/marketing/cross-posting

4.1. Learning Representations of TikTok Content

We need to distill the essence of a TikTok post into a vector of features that captures both the content and quality of the post, and that can be integrated with our causal model. We learn such a feature vector by combining two technologies: transfer learning and multi-modal representation learning. With transfer learning, we extract a set of features that describe all aspects of a TikTok post, including its video and accompanying caption. The extraction of features using pre-trained models simplifies the subsequent representation learning process, and allows us to explicitly incorporate features relevant for TikTok posts (e.g., how the posts are edited). Then, we develop a multi-modal representation learning framework to condense this set of features, and their interlinkages, into a single lower-dimensional vector, denoted $\boldsymbol{r}$, which we use in our causal model to control for content. We now briefly describe the two parts, and defer a detailed discussion to Web Appendix B.

Extracting Features with Transfer Learning From each post on TikTok, we extract features from four modalities: (1) textual features, from the video and its caption, which capture what is said in the video and how the video is described; (2) image features, from the video's image frames, which capture the objects and sentiment in the video, and other high-level visual features; (3) audio features from the video's soundtrack, which capture features of the video's raw acoustic properties; and (4) editing features, which capture how creators edit their videos on the platform, including things like the presence of overlaid images (i.e., "stickers") and image filters. While the first three are standard in video analysis, the fourth is a unique modality for TikTok. Together, these features comprehensively capture the content of TikTok posts. For a complete description of the features we extract, see Web Appendix B.1.

Representation Learning While many of the features we extract are relatively high-level, the full set of features is still high-dimensional, and each modality is treated separately in the feature extraction. To further reduce the dimensionality of these features, and synthesize them in a way that produces a high-level representation of the essence of each post as a whole, we apply a representation learning procedure based on the variational autoencoder.

Variational autoencoders (or VAEs) are based on traditional autoencoders, which are machine learning models with two components: an encoder that compresses the data to a dense vector representation, and a decoder that reconstructs the original data from that representation

(Kingma and Welling 2013; Rezende, Mohamed, and Wierstra 2014). The variational autoencoder is a probabilistic variant of this framework, where the generative process of the observed data, $\mathbf{x}_n$ for observation n, is modeled as a function of lower-dimensional latent vector representations, $\boldsymbol{r}_n$. The VAE has two parts: the encoder (or "inference network") specifies an approximate posterior distribution over $\boldsymbol{r}$, given data $\mathbf{x}$, and a standard Gaussian prior for $\boldsymbol{r}$, by learning a density $q_\phi(\boldsymbol{r}|\boldsymbol{x}) \approx p(\boldsymbol{r}|\boldsymbol{x})$. The decoder (or "generative model") models the data generating process as a function of $\boldsymbol{r}$, $\mathbf{x} \sim p_\theta(\boldsymbol{x}|\boldsymbol{r})$. From the decoder, we can see that $\boldsymbol{r}$ acts as a sufficient statistic for the data: given $\boldsymbol{r}$, we can reconstruct $\boldsymbol{x}$. In both components, the link between the data and the representation is parameterized by neural networks with parameters ϕ for the encoder and θ for the decoder. VAEs have been shown to outperform traditional (non-probabilistic) autoencoders in learning meaningful, low-dimensional representations of data, across a variety of domains (Hsu, Zhang, and Glass 2017; Yao et al. 2019).

The version of VAE we develop is a *multimodal* VAE. Its structure is similar to a standard VAE, with both encoders and decoders, but it incorporates special structure within each of these components that allows it to merge together information from each of the four modalities of the TikTok data (text, image, audio, editing). Specifically, in the encoder, we first use modality-specific networks that capture the specifics of a given modality. Next, we combine the modality-specific encoders using a neural network, which allows them to interact in a flexible (potentially nonlinear) manner. This joint encoder is then used to estimate the latent representation of a post. The decoder mirrors this structure. Given the highly structured nature of the modality-specific encoders and decoders, we refer to our proposed model as a structured multimodal VAE or SMVAE. Figure 4 shows the overall modeling framework. The primary output of this framework comes from the encoder: given a new TikTok post $\mathbf{x}_n$, the encoder gives us an estimate of $p(\boldsymbol{r}_n|\mathbf{x}_n)$, in terms of a posterior mean, $\boldsymbol{\mu}_n$, and variance, $\boldsymbol{\sigma}_n$. This posterior mean is the representation we use to control for the content of the post in our causal framework. Going forward, we slightly abuse our notation and use $\boldsymbol{r}_n$ to refer to the estimated value of $\boldsymbol{r}_n$, which we take to be the posterior mean ($\boldsymbol{\mu}_n$). We include a detailed description of each component of our VAE framework in Web Appendix B.2.

Validating the Representations There are two aspects of TikTok posts that we need our representations to capture: their content, meaning the actual visual, auditory, and textual features of the post, and their quality, meaning the more subjective sense that a video is "good." We validate that our representations capture the content and quality of TikTok posts in two ways: reconstruction

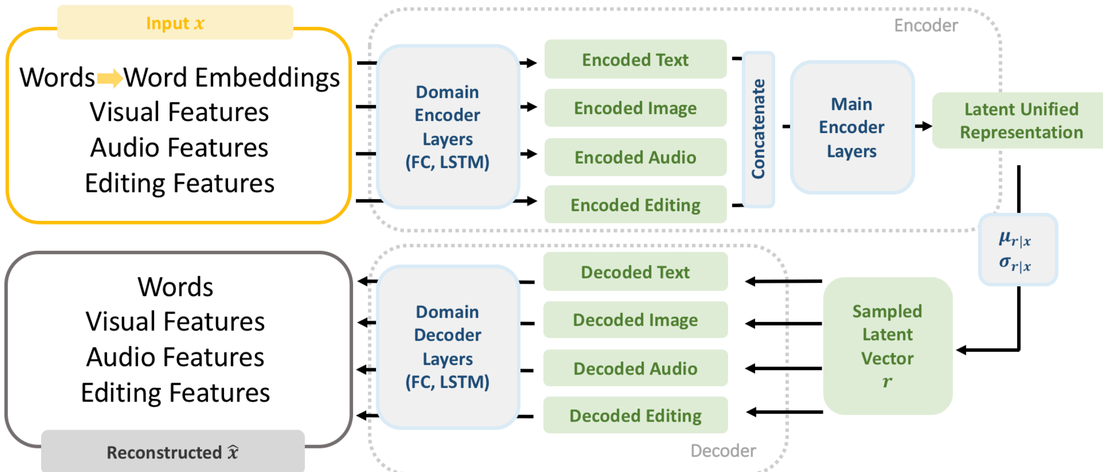

Figure 4: Network Architecture of the Proposed SMVAE
Illustration of the network architecture of our representation learning framework, the SMVAE. At top is the encoder, and at bottom, the decoder.

accuracy, and the ability of the representations to predict the success of a video.[17]

First, if the SMVAE is indeed capturing post content, then for a new post, we should be able to use the SMVAE to accurately reconstruct the post from its representation. We test this for a heldout set of videos, and report the results in Table 2, where we compare the reconstruction ability of our model to a no information rate (NIR). The NIR is a model that predicts the mean value for every feature. We include the NIR to provide a baseline for interpreting the magnitude of the reconstruction statistics.[18]

We see that our model's reconstruction ability greatly improves on the NIR, suggesting that our model is able to meaningfully reconstruct the data. That we can reconstruct the original post from its representation gives us confidence that the representation indeed acts as a "sufficient statistic" for the content of the post. In Web Appendix C, we also compare the reconstruction ability of the full model to submodels that use a subset of the four modalities. We find that the reconstruction accuracy for any given modality improves when other modalities are added to the model. For example, a VAE trained using only textual data achieves an accuracy of 0.14 in reconstructing which words are present in a video, while our full SMVAE achieves an accuracy of 0.38, suggesting that the SMVAE captures synergies between the different aspects of a TikTok

[17]We include a third "face validity" validation in Web Appendix C, illustrating that differences in representations map onto meaningful differences in videos.

[18]A useful analogy is thinking of an intercept-only model in regression analysis: comparing how much variability is explained by the full model, versus an intercept-only model tells us how much variability is explained by the model. The NIR is equivalent to the intercept-only model here.

	Text Accuracy	(Avg.) Image MSE	Audio MSE	Editing MSE
SMVAE	0.38	10.81	0.97	0.69
NIR	0.04	35.21	3.96	40.33

Table 2: Reconstruction Statistics under SMVAE and NIR
For text, we measure the percentage of the reconstructed words that correctly match the original words (a higher number indicates better fit). For the other three modalities, we calculate the mean squared error between the reconstructed and original features (higher numbers indicate worse fit). For image, we average the fit statistic over all image frames.

post.

The prior analysis illustrates that the representations can capture video content, in the sense of being able to reconstruct it, but it leaves open the question of quality. For example, there may be two videos with similar content — say, both about cats doing silly things — but one becomes very popular (i.e., a *good* cat video) and the other does not (i.e., a *bad* cat video). In this next analysis, we attempt to validate that our representations can also capture the quality of a post, by showcasing their usefulness in predicting which videos become surprisingly popular. By surprisingly, we mean that a video becomes more popular and grows in popularity faster than would be expected given its "initial conditions." By initial conditions, we mean covariates like which hashtag(s) it was first posted under or statistics about its creator. To model the popularity of posts over time, conditional on these initial conditions, we calibrate a hierarchical logarithmic growth model, with post-level parameters capturing the initial number of impressions for the post and its growth rate over time. We include the initial conditions of the video as predictors of these post-level parameters. We then test whether the representations $\boldsymbol{r}$ can incrementally explain the post's initial impressions and growth rate, beyond just the initial conditions. Our results, which we present in detail in Web Appendix C.1, show that our representations can predict with over 70% accuracy which videos over and underperform their initial conditions. We also show that the full SMVAE model is superior to the other simpler VAE models in this test. This result suggests that there is enough nuance captured in $\boldsymbol{r}$ to meaningfully predict which videos become surprisingly popular. That is, our representations are meaningful indicators of quality beyond basic content.

4.2. Deep IV and the Follower Elasticity of Impressions

While the representation learning procedure provides a way to control for video content, there are three remaining concerns for understanding the link between popularity of the influencer and a video: (1) the potentially nonlinear relationship between these variables; (2) potentially heterogeneous response functions across content types; and (3) the possibility of unobserved confounders that may simultaneously drive both popularity and impressions, and thus bias our understanding of that relationship. To account for these concerns, we turn to deep instrumental variables.

Instrumental variable methods have been widely used for causal effect estimation with observational data when there may be unobserved confounders. Most existing approaches (e.g., two-stage least squares) do so under two strong assumptions: linearity and homogeneity (e.g., Angrist, Imbens, and Rubin 1996). That is, they assume that the treatment affects all individuals in the same, constant way. In our case, as described above, both of these assumptions are unlikely to be true. To overcome these limitations, a recent literature has emerged that combines standard causal inference approaches with flexible machine learning models (e.g., Athey, Tibshirani, and Wager 2019; Farrell, Liang, and Misra 2021). In the context of instrumental variables, Hartford et al. (2017) have proposed Deep IV, a model that generalizes two-stage least squares by estimating both the first and second stages of the IV framework through deep neural networks, thereby allowing both heterogeneous and nonlinear estimation of causal effects.[19]

We adopt Deep IV to estimate the relationship between influencer popularity and video impressions, by assuming the following data generating model for video impressions:

$$Y_j = g(F_{i(j)}, X_j) + e_j. \tag{1}$$

Here, Y_j is the log number of impressions of video j, measured two weeks after its posting; $F_{i(j)}$ is the log follower size of creator i at the time of posting video j; $i(j)$ is the one-to-one mapping between video j and the influencer who posts video j; X_j contains all of our covariates, including both the metadata of the post and its creator (described in Web Appendices A.2 and A.3), and the content representation, $\boldsymbol{r}_j$; and finally, e_j is an error term. The function $g(\cdot)$ is assumed to be unknown and potentially nonlinear. We use a log transformation for impressions and followers to account for the skewness of their distributions (see Figure 2), which has the added benefit of simplifying some of the subsequent derivations. Importantly, we do not assume that the error term, e_j, is independent of $F_{i(j)}$; indeed, in the presence of unobserved confounders, they will be non-negligibly correlated.

In light of this potential correlation, to estimate counterfactual quantities like our focal question — how many impressions would this video have gotten if it had been posted by a similar but more popular influencer? — we need an instrument for the number of followers. We will describe

[19] We include a detailed overview of how Deep IV works in Web Appendix D.1.

such an instrument in detail in the next section. Assuming we have a valid instrument, Deep IV outputs an estimate of what Hartford et al. (2017) term the *counterfactual prediction function*. Denoted $h(F, X)$, the counterfactual prediction function captures the variation in impressions due only to exogenous variation in F, conditional on X, the covariates.[20] It is parameterized through a neural network, which allows us to flexibly capture both nonlinearities in the relationship between F and Y, and how features, X, moderate that relationship (i.e., heterogeneous effects).[21]

From this counterfactual prediction function, we derive our key metric of interest: the *follower elasticity of impressions*, or FEI. In simple terms, the FEI captures the marginal effect of changing the number of followers of an influencer on the impressions of their posts. To compute these marginal effects, we first take differences in h. More specifically, suppose that we are considering how the impression count of a video with features X would change if the log number of followers of its creator increased from F to $F + \epsilon$, $\epsilon > 0$. The expected change in the impressions of the video resulting from that ϵ increase can be determined by:

$$\mathbb{E}[Y \mid F + \epsilon, X] - \mathbb{E}[Y \mid F, X] = h(F + \epsilon, X) - h(F, X). \tag{2}$$

Dividing both sides by ϵ and taking $\epsilon \to 0$ yields:

$$\frac{\partial}{\partial F}\mathbb{E}[Y \mid F, X]. \tag{3}$$

Finally, since both F and Y are logged quantities, this partial derivative exactly corresponds to the percentage increase in impressions associated with a unit percentage increase in followers, which is our follower elasticity of impressions. Since FEI can vary with the number of followers, when examining a video's FEI, we must consider the entire FEI *curve*. Moreover, since the heterogeneous marginal effect depends on post features (X), the FEI curve may also differ between videos. As we show in the next section, these (heterogeneous) FEI curves are crucial for determining the benefit a firm may receive from collaborating with an influencer.

4.3. Instrument: Past Video Success

Our modeling framework requires a valid instrument for the number of followers that influencers have at the time they post their videos. Such an instrument should cause variation in influencers'

[20] Note that, in Web Appendix D, when deriving Deep IV, the estimated counterfactual prediction function is denoted h_{ζ^*}, rather than just h, to capture its dependence on a set of estimated parameters. We suppress that notation here for simplicity.

[21] We describe the details of our neural architecture in Web Appendix D.2.

follower sizes at posting time (the "relevance" condition), but not directly affect the impression count of the video being posted (the "exclusion" condition), and not be correlated with the error term in the outcome model, i.e., Equation 1 (the "exogeneity" condition). The instrument we use that achieves these three goals is *the number of likes received historically on posts by a similar influencer*, following Leung et al. (2022). For ease of exposition, we will refer to this quantity as "similar influencer likes." We first describe the intuition and context of this instrument, before carefully detailing how we construct this IV and how it meets the three conditions.

On TikTok, viewers can indicate that they like a video by clicking a small heart icon next to the video. These likes (or "hearts") are one way to measure the performance of a video, in terms of generating engagement. Our instrument is based on looking at *historical* likes: that is, likes on other videos that occurred prior to the time a focal video was posted. Intuitively, historical likes for the focal influencer should be predictive of the focal influencer's follower size. However, our identification strategy requires that the historical likes of a *similar* creator be correlated with the follower size of the focal influencer. We argue that if two creators are similar, they will produce similar types of videos. Thus, general interest in this type of influencer should manifest in both historical likes for the similar influencer, and in followers for the focal user, resulting in a correlation between the two. This correlation is the basis of our argument for the instrument's relevancy. Another important feature of this instrument is that TikTok has explicitly stated that the historical success of a creator's videos, in terms of number of likes, is not used in its recommendation system. This fact, combined with our controls for the content of videos, is central to our arguments for the instrument's exclusiveness and exogeneity, as we describe in more detail below.

Construction of the Instrument Our instrument is a similar user's *historical* likes: that is, likes of other videos that happened prior to the time a focal video was posted. It is constructed in a two-part process: first, for a given video j, posted under hashtag H (denoted $j \in H$), we compute the distance between video j's representation vector $\boldsymbol{r}_j$, and the representations of all other videos posted under the same hashtag, $\boldsymbol{r}_k$, $k \in H, k \neq j$. We select the creator of the second closest video to j as our similar influencer, excluding any videos made by the focal influencer. We index this second closest video by s. We choose the second-closest video's creator rather than the closest to preserve relevance, while mitigating concerns about influencer imitation and fan overlap, which we detail below.[22] From the metadata we collected about each video, we are able to see how

[22]This approach is inspired by Leung et al. (2022). Note that, while we choose the second-most similar influencer for reasons we describe subsequently, the results are largely similar if we use the most similar influencer. The results are

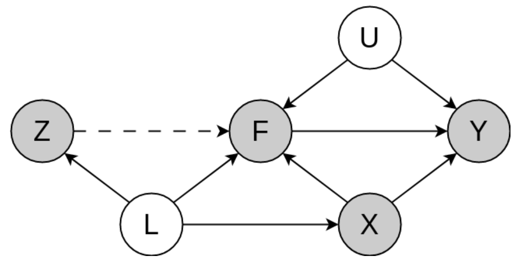

Figure 5: Causal DAG
Illustration of our assumed causal pathways: We assume the instrument Z is correlated with the key treatment variable, followers (F), as both are driven by a latent interest in a type of influencer (L). Controlling for content X is important both because it confounds the relationship between F and the outcome Y (impressions), and because it blocks a causal pathway from L. Finally, U captures any potential unobserved confounders, which motivate the use of the instrument in the first place.

many likes this similar influencer has gotten, in total, on their videos, including the number of likes on video s. Thus, our instrument, denoted z_j, is the number of historical likes of this similar influencer; that is, the total number of likes for the similar influencer less the number of likes on video s.[23]

Relevance We now describe in more detail how this IV meets the relevance, exclusiveness, and exogeneity criteria, starting with relevance. The essence of our claim is captured by the causal directed acyclic graphic (DAG) (Pearl 2009) shown in Figure 5. In this DAG, we illustrate the focal causal pathways between variables with arrows, and label each variable using the same notation as before: Y is our dependent variable, impressions; F is our treatment variable, followers; and Z is our instrument, which is meant to address the possibility of unobserved confounders, denoted U. The variable X captures our controls, which primarily relate to the content of the focal set of videos. Finally, L refers to latent interest in a type of influencer. Shaded nodes are observed, while white nodes are unobserved.

As shown in the DAG, the relevance of historical likes for a similar influencer (Z) to current followers of a focal user (F) comes from the assumption that both are generated by shared interests among TikTok users in a type of influencer (L). By type, we mean just that the influencer tends to post a particular type of content. We motivate this idea further with an example: Consider a hypo-

also robust to whether we measure similarity just based on the video representations or on the video representations plus our creator statistics.

[23]The construction of our instrument involves an assumption that the similarity of videos under the focal hashtag is representative of similarities between the influencers. We argue that this assumption is reasonable: if two influencers are of a similar type, they should generate similar content, which means that their videos posted under the same hashtag should be similar. Reversing this logic, seeing two very similar videos under the same hashtag should be predictive of similar influencer types. However, we cannot empirically support the claim that these influencers have historically posted similar types of video, since we only observe cross-sectional data at hashtag level.

thetical hashtag "#WhatIEat" where TikTok users post videos describing their meals. Suppose we see our focal influencer posting a video showing some avocado toast, and we find the second-most similar video, also about avocado toast, posted by another influencer. By our definition above, we will take the creator of that second-most similar video as our second-most similar influencer, and use their historical likes to construct our instrument. We argue that the similarity between these videos is a sign that these influencers are of a similar latent type L. In this example, both may be health influencers. The heart of our relevancy argument is that the level of interest among TikTok users in health topics causally determines both the success of the similar influencer's past videos (Z) and the number of followers of the focal influencer (F), leading to a correlation between these variables, even though we cannot observe L. This argument echoes the Hausman-type instruments often used in demand estimation: in that context, prices from other markets are used as instruments for prices in the focal market under the assumption that supply-related variables like material costs may jointly drive variation across markets (Hausman 1996). Here, the "other markets" are other influencers, and the supply-related factors are the latent influencer types, which drive their popularity on TikTok.

Empirically, we find evidence to support the relevancy of this instrument. Echoing a simple two-stage least squares analysis, we run a "first stage" regression of followers (F) on our instrument (Z) and content (X, which contains $\boldsymbol{r}_j$):

$$F_{i(j)} = \beta' X_j + \gamma' Z_j + \varepsilon_j. \tag{4}$$

We find that the partial F statistic for Z_j exceeds the typical threshold of 10 for determining relevancy, and that the partial R-square is reasonably high (for an instrument), at 0.129. In both cases, the word "partial" refers to the association between $F_{i(j)}$ and Z_j, net of X_j. Statistically, these results suggest that our instrument is relevant and not problematically weak.

Exclusiveness and Exogeneity Exclusiveness (or the exclusion restriction) means that the instrument drives the outcome only through the treatment, while exogeneity means that the instrument is uncorrelated with any potential unobserved confounders of the treatment and outcome. In our context, establishing exclusiveness and exogeneity boils down to arguing that the number of similar influencer (historical) likes does not affect the number of impressions a focal video receives, except through its correlation with the focal user's follower size, or in ways that we control for explicitly in our model. On TikTok, there are five possible channels through which users can discover content: (1) TikTok's personalized "For You" feed; (2) the Discover page, (3) the Following feed, (4) searching on TikTok, and (5) sharing. Thus, establishing exclusiveness and exogeneity in general requires establishing exclusiveness and exogeneity in each of these five channels.

The first channel is the *For You Page*, or FYP, which is a customized list of videos recommended by TikTok based on a user's activities and interests on the platform. While recommendation systems are often black boxes, TikTok has posted public details about how its FYP recommendations work, which support our exclusion and exogeneity restrictions.[24] According to TikTok, neither a creator's follower size nor whether an account has had previous high-performing videos are directly used by the recommendation system. Thus, historical likes, in general, do not affect the likelihood of a user being exposed to a creator's content. In our case, this exclusion is especially plausible, as we are not using the focal creator's historical likes, but a *similar* creator's historical likes.

While the pure number of historical likes of a video is not a factor in the recommendation system, specific historic interactions with a creator's content *are*. One factor that is explicitly used in FYP recommendations is a specific user's previous interactions with a specific creator's content, including viewing and liking that creator's videos. This means, for instance, that if a user had previously liked a creator's video, the user may be exposed at a higher rate to that creator's content on the For You Page in the future, even if the user did not follow the creator. In our case, this direct interaction channel is still not a problem, as again, we are looking at a similar creator, not the focal creator. The only way that historical likes of a similar influencer might directly affect impressions of a focal video is through similarities in content, which again, TikTok acknowledges as a potential factor in recommendations. However, our analysis controls for content, which blocks this potential causal pathway, and preserves the exogeneity restriction.[25]

The second channel is TikTok's *Discover page*, which shows users a list of hashtags (i.e., video categories) that are currently trending. After clicking on a hashtag, users are taken to a ranked list of videos associated with that hashtag. Both the list of hashtags and the videos associated with each hashtag are the same for every user. Thus, the main threat to our instrument's validity is if somehow the historical likes of a similar influencer drive the prominence of the focal video on the Discover page. The only plausible mechanism for such an effect is through content: if

[24] https://newsroom.tiktok.com/en-us/how-tiktok-recommends-videos-for-you

[25] This argument depends on our content representations, r_n, being able to comprehensively capture the content of videos. While we have tried to empirically justify that in Web Appendix C, no representation is perfect. If there are any features of videos that are systematically missing from these representations, then that content will not be controlled for in our analysis.

the similar user's content has gotten many likes in the past, this information could be used by TikTok's algorithms to rank the focal user's (similar) video more highly. Even if this is the case, we explicitly control for rank on the Discover page as one of our covariates. We also control for the content of the video, which rules out any mechanisms driven by the thumbnail shown on the Discover page. Thus, there is no reason to suspect that exposure on the Discover page violates exclusiveness or exogeneity.

The third channel is the *following feed*, which is the most straightforward to address: when a user follows a creator, that creator's videos will appear on the user's feed. It is not possible to appear on the following feed without following; thus, exposure to videos here happens exclusively through the follower count.

The next channel is *search*: If a user wants to see a particular type of content, they can use TikTok's search bar to find it. Similarly to searching on other platforms, users can search for content on TikTok using content-related keywords or using the creator's username. Again, we argue that historical likes of a similar influencer will not affect a focal video's exposure through this channel, after controlling for content. The validity concern would be that historical liking of the similar influencer may cause more people to search for the focal influencer, and thus, lead to impressions on the focal video. One way that could happen is by searching for content. However, our model controls for content directly, so this causal path is blocked. Therefore, the remaining concern would be that the historical liking of a similar influencer causes more people to search for the focal influencer by name. We argue that such a direct link from a similar creator's historical likes to direct name searches for the focal influencer is implausible: the only way a user can search for an influencer by name is if the user already knows the influencer, which in turn happens when either the influencer is famous enough to have name recognition more broadly, or the influencer also posts on other platforms, driving cross-platform users to search for the influencer on TikTok. Neither of these mechanisms invalidates our instrument: name recognition is unlikely to be driven by a *similar* user's historical likes, and, while cross-posting is common among influencers, for it to be a problem here would imply that cross-posting of the similar influencer is causing cross-posting of the focal influencer (or vice versa). Although such imitation is unlikely, our selection of the *second* closest influencer is intended to mitigate this concern.

The last channel is *sharing*: In addition to the four direct channels discussed previously, a video might also get impressions by being shared. The validity concern here would be that historical likes for a similar influencer may drive impressions for the focal video via sharing. Echoing previous arguments, one mechanism by which this could happen is through content, since influencers may create similar content over time. In this case, historical liking of a similar influencer's videos may be driven by a propensity to share that type of content, and may thus be non-trivially correlated with impressions of the focal video. In our case, this is not a concern, as we control for content. Alternatively, regardless of content, people may have a higher propensity to share because of the influencers themselves. For example, if the influencer is a celebrity or has a more loyal fan base, then users may be more likely to share that influencer's content. Such effects may be problematic if celebrity status or fan base loyalty are correlated across influencers. To control for celebrity status, we include whether an influencer is verified or not as part of our post statistics (i.e., part of X).[26]. Regarding fan loyalty, the argument would be that the propensity to share is correlated with L, the latent interest in influencer types, such that fans of similar influencers are themselves similar in terms of their propensity to share content, driving a correlation between Z and F. We partly address this concern, again, by controlling for content. We also attempt to mitigate this concern by our choice of the second most similar influencer: by adding distance between the focal influencer and the similar influencer, we reduce the potential for fan overlap.[27]

In sum, we have argued that our instrument, similar influencer historical likes, is not a driver of impressions of the focal video, except through its correlation with the number of followers, after controlling for key observables like content.

4.4. Benefits of Our Framework

Before describing our results, we first briefly highlight the key benefits of our framework. We have proposed a causal framework that generates an important metric to consider for assessing the effectiveness of influencer advertising: the follower elasticity of impressions (FEI). An important aspect of our modeling framework is that we account for potential unobserved confounders with an instrumental variable. If left unaccounted for, unobserved confounders like cross-posting may pose a problem: firms may see a positive link between followers and impressions, and thus believe they should pay more for more popular influencers, when in reality the positive link comes from cross-posting behavior. If the firm then pays for a more popular influencer, but who is not cross-posting, the benefits of that incremental popularity may be small. By using Deep IV, we ensure

[26]TikTok has a full list of criteria to determine whether an influencer could be rewarded as verified which contain a large variety of aspects that might make an influencer special, including being a celebrity (https://newsroom.tiktok.com/en-us/how-to-tell-if-an-account-is-verified-on-tiktok)

[27]The identities of followers are not observed in our data, so we cannot test this empirically.

that the link we capture is truly the causal link between followers and impressions, allowing for possible unobserved confounders.

Novel to our framework is the fusion of a representation learning model with Deep IV, and the simultaneous use of that representation learning model to construct a valid instrument. While each of these elements has been used in other papers and contexts, their merger here represents a step forward in doing causal analysis with unstructured data. The pre-processing of the unstructured data through representation learning methods helps us understand our data and construct our instrument. It also helps with the tractability of the framework, similar to how unsupervised pre-training helps in more classic deep learning contexts (Erhan et al. 2010). Compared to classic two-stage least squares and other nonparametric IV methods such as sieve or kernel-based methods (e.g., Newey and Powell 2003), deep IV does not require a strong prior understanding of the data generating process and is computationally tractable, especially when there are many inputs and a large number of training samples. More recent work has further established the convergence properties of deep learning-based frameworks under various network architectures and a general class of nonparametric regression-type loss functions (Farrell, Liang, and Misra 2021), giving still further theoretical justification to our architecture.

Finally, while our model is intended as a causal analysis, rather than a purely predictive analysis, it is worth noting that our approach also performs quite well at prediction: in Web Appendix E, we compare the predictive performance of our Deep IV model to benchmarks including a neural network (i.e., our model but without the IV), a random forest (again, without the IV), and a standard linear 2SLS specification (with the same IV). While Deep IV predictably does worse than the no-IV neural network model, as it leverages only limited variation in the number of followers (i.e., only the plausibly exogenous variation, via the IV), it still performs competitively, achieving equivalent predictive performance to the purely predictive random forest. The key benefit of approaching this problem from a causal, rather than predictive perspective, is that it provides better guidance on firms optimizing their ROI: it allows us to reason about *counterfactual* impressions garnered by a given level of popularity, accounting for the possibility of unobserved confounders, which is crucial for firms making decisions regarding popularity.

5. Results

We begin by describing the patterns that emerge in our estimated follower elasticity curves, and linking those patterns to prior literature. We then describe the implications of these findings in terms of designing effective influencer marketing campaigns.

5.1. Average Follower Elasticity

To begin, in the left panel of Figure 6, we plot the average FEI, calculated through our framework, averaged over all the videos in our dataset. In the right panel, we compute the counterfactual predicted average (log) impressions for each (log) follower count, again averaged over videos.[28] Recall that prior literature suggests that the relationship between popularity and different forms of engagement is often nonlinear, and that our simple correlational analysis in Figure 3 suggested a potentially nonlinear link between followers and impressions. In Figure 6, we see that the estimated relationship is indeed nonlinear, but of a very different shape than Figure 3: while FEI is always positive, suggesting that increasing follower counts are always associated with more impressions, the curve exhibits an inverted-U shape. This shape suggests that a percentage gain in followers yields, on average, the largest percentage gain in impressions for relatively small influencers, but not the small*est* influencers. In fact, the maximum average FEI is attained with only 6,500 followers. When translated into counterfactual predicted impressions in the right panel, we see that the relationship between followers and impressions is S-shaped, which is very different from the pattern suggested in the purely correlational analysis from Figure 3.

The nonlinear patterns we find are consistent with the idea that more popular influencers may suffer from weaker connections with their followers. Recall that one of the channels by which posts gain impressions on TikTok is by being shared. Prior literature suggests that there is an inverted-U shape in terms of engagement with an influencer's content, where the rate of behaviors like liking and commenting on a post first increases as an influencer becomes more popular and then decreases (Wies, Bleier, and Edeling 2023). Higher engagement may also increase the propensity of followers to share content, leading to the inverted-U shape we observe between FEI and popularity. Indeed, we find empirical support for this: in Figure 7, we show that the posts in

[28] As described in the previous section, the FEI is essentially the derivative of the counterfactual prediction curve. Hence, to convert an FEI curve to a counterfactual prediction curve requires a few assumptions about the overall level of impressions of the video (i.e., the constant from integration). For average curves, such as those shown in Figures 6, we assume that the constant is just the average number of impressions of all the videos in the data. For curves that we show later, we use the actual level of impressions observed for that video.

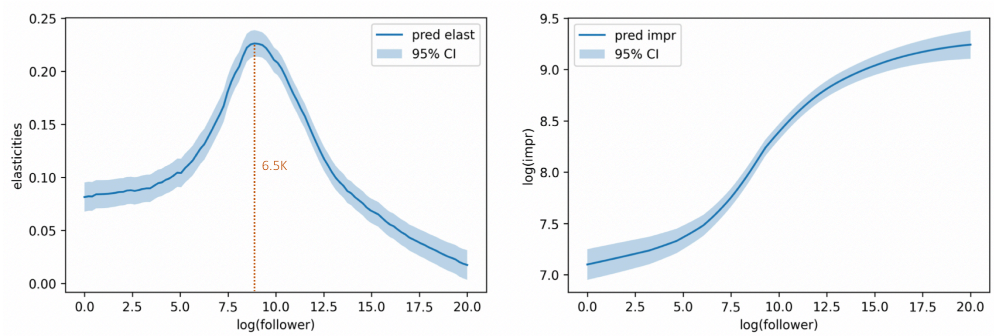

Figure 6: Average FEI and Counterfactual Prediction
At left, we plot the average follower elasticity of impressions, averaged over x, computed at different levels of log followers. At right, we plot the counterfactual prediction curve, which is the average predicted impressions of a video, given the number of followers of its creator. Uncertainty bands are 95% confidence intervals for the mean, computed by bootstrapping.

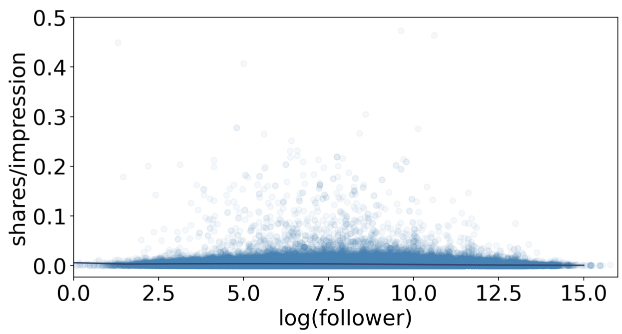

Figure 7: Sharing as a Function of Followers
Each point represents a single post. On the y-axis, we show the number of shares that post had per impression. On the x-axis is the log number of followers of the influencer who created the post.

our data that were shared the most frequently per impression are those that fall in the mid-range of followers, suggesting that the incremental boost in impressions that comes from sharing will be highest for influencers in the mid-tier of popularity. Beyond sharing, lower engagement for small and large influencers may also have an indirect effect on impressions, by means of TikTok's recommendation algorithm. Engagement, in terms of likes and comments, is explicitly used as a factor in TikTok's For You Page algorithm. Thus, the more frequently a user likes or comments on an influencer's posts, the more frequently that influencer's content will appear on the user's FYP.

These average FEI results also connect to the literature on advertising. Influencer marketing is, essentially, social media advertising done by third parties (i.e., influencers). Thus, it is interesting to consider how its effectiveness compares with traditional forms of advertising. From the left panel of Figure 6, we calculate the average FEI across videos by calculating the area under the FEI curve and dividing by the range of log followers. We find that the average FEI is 0.10, which is very similar to the overall short-term brand advertising elasticity of 0.12 reported in a meta-analysis by Sethuraman, Tellis, and Briesch (2011). The S-shaped response function in the right panel also connects to the advertising literature: while an S-shaped response curve has been postulated in advertising, it has been hard to detect in empirical studies (Simon and Arndt 1980). The argument for an S-shaped response in advertising is that ads must achieve a certain threshold of reach before they become effective. The lack of empirical evidence for such a curve is likely in part due to the fact that most ad campaigns are relatively large scale, and thus, that threshold is always exceeded. Given our rich data, which contain even very small influencers, we are able to find such an S-shaped pattern. That the threshold where impression growth accelerates is quite low lends support to the fact that, while threshold effects may exist in advertising-like contexts such as influencer marketing, they may be irrelevant for even modestly sized campaigns.

5.2. Follower Elasticities by Content

While the average FEI curve offers some insights for managers, the real power of our framework comes from its ability to compute the FEI curve for any video. Thus, in this section, we investigate how FEI varies across hashtags depending on content, and what that variation means for both firms and influencers.

Engagement Tactics The first dimension of content we consider is how users are encouraged to engage with a hashtag. Firms use influencer campaigns in various ways: in some cases, the campaign is designed to educate consumers on how best to use the firm's products. In other cases, firms may encourage consumers to participate in social activities that integrate their brands, or simply seek to entertain them. These different *engagement tactics* mirror a classification scheme previously proposed by Dolan et al. (2019) and summarized in Table 3 for how firms can engage with customers on social media by informing them, entertaining them, or socializing with them. Connecting these challenge types to FEI, we expect that socializing campaigns may rely more on large numbers of people becoming aware of them, and thus may be more effectively promoted by more popular influencers. In contrast, informational campaigns rely on signals of trust and authenticity, suggesting that using smaller influencers to promote them may be more effective, given prior research that suggests that influencers with a smaller number of followers have stronger ties with their followers (e.g., Katona, Zubcsek, and Sarvary 2011) and are viewed as more authentic (Park et al. 2021).

Engagement Tactic	Definition	Examples
Informational ($N = 81$)	Content that provides users with resources and helpful information (e.g., content about events, places, opportunities, people or celebrities).	#WomenInStem, #HomeOffice, #FallDIY
Entertaining ($N = 104$)	Content that is meant to be fun and entertaining, without explicit informational value	#GreenScreenScan, #FallGuysMoments, #OhNo
Socializing ($N = 87$)	Content that encourages users to interact with one another, and stimulates their desire for social integration and social benefits	#LaughingDuet, #PerfectMatch, #GroupChat

Table 3: Engagement Tactics
Definitions and examples of the three engagement tactics we use to classify hashtags. N refers to the number of hashtags exhibiting that engagement tactic.

To explore these ideas, we recruited independent raters to assign informational, entertaining, and socializing labels to each of our hashtags.[29]. Then, we computed an average FEI curve for all videos in each hashtag category. Figure 8 shows the results. We find significant differences in the FEI curve based on the engagement tactic. Consistent with our expectations, influencers with less than 10,000 followers achieve faster impression growth for informational hashtags, while influencers with a larger number of followers achieve faster growth for socializing hashtags. While we had no hypothesis regarding entertaining hashtags, we find that they exhibit FEI patterns much closer to the informational hashtags. In addition to these differences in shape, we also find differences in magnitude: socializing hashtags are overall a less effective tactic in generating impressions compared to entertaining and informational, given the much smaller average FEI for socializing (0.08) than for the other two (around 0.15). That being said, a benefit of our Deep IV framework is that we can see how FEI changes over the range of followers, in a nonlinear way. Based on the estimated FEI curves, while the average elasticities of entertaining and informational may be higher, for certain ranges of followers, socializing has a higher FEI.

Topic Beyond engagement tactics, hashtags also vary in their topic — that is, what the videos in that hashtag are actually about — which in turn, may moderate FEI. To explore the role of topic, we summarize common content topics across hashtags by clustering the word embeddings of words in the hashtags.[30] We find five broad content topics: life (e.g., #WeekendVibes, #GoodMorning),

[29]We do not require the categories of engagement tactics to be exclusive. A video can utilize one or multiple tactics.
[30]For more details, see Web Appendix F.

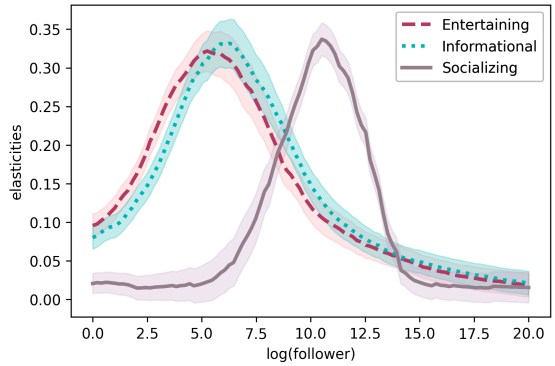

Figure 8: Average FEI by Engagement Tactics
The average FEI for videos that were labelled as coming from entertaining, informational, or socializing hashtags, computed at different values of log followers (averaged over x). Uncertainty bands are 95% confidence intervals for the mean, computed by bootstrapping.

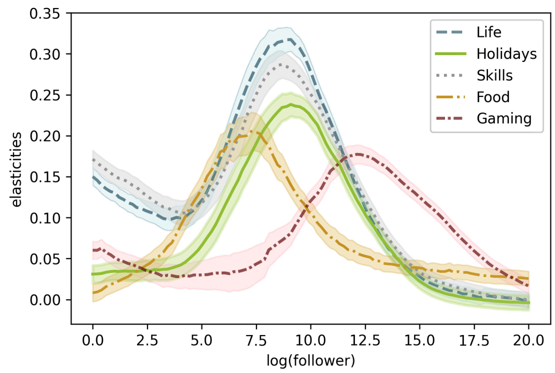

Figure 9: Average FEI by Content Topics
The average FEI for videos that were categorized by their hashtag as belonging to one of five content topics, reflecting what the videos are about at a high level. The FEI is computed at different values of log followers, averaged over x within the topic. Uncertainty bands are 95% confidence intervals for the mean, computed by bootstrapping.

holidays (e.g., #Christmas2020, #HappyHolidays), skills (e.g., #Yoga101, #InkDrawing), food (e.g., #HomeCooked, #HealthyCooking), and gaming (e.g., #GamerGoals, #GamingTikTok). Figure 9 shows the average FEI curve for each of these topics.

Across content types, while FEI always peaks in the midrange of followers, the timing of this peak, and the overall magnitude of the elasticity, varies substantially. Holiday posts, for instance, have a lower average elasticity than other video types. We hypothesize that this is due to the short lifespan of holiday posts. We operationalize lifespan as the time until a video reaches 90% of its total two-week impressions. In our data, the average lifespan of videos in holiday-related hashtags is 2.57 days, significantly shorter than all other types of videos, which have an average lifespan of 6.29 days ($t = -14.53$, $p < 1 \times 10^{-10}$). Intuitively, holiday content is most relevant during the holiday itself, a fairly short time, and any given video competes with all other holiday content during that narrow window. These two forces drive down the returns to popularity, on average, for holiday content.

Another interesting topic is gaming, which has a much later peak in FEI. This substantially later peak mirrors what we saw for socializing hashtags and suggests that more popular influencers may more effectively generate impressions for gaming-related content. Intuitively, gaming is a highly social context, in which network effects are very prominent (Tudón 2022). These prominent network effects suggest that gamers look to other gamers for signals of what games to play, a sentiment which seems to carry over to content on social media, in the form of strong social effects. Indeed, we find empirical evidence supporting that: in our data, we find the average number of comments per follower for gaming-related hashtags is 0.34, significantly higher than the 0.22 comments per follower achieved by non-gaming creators ($t = 2.82$, $p = 0.006$). Likewise, the number of shares per follower is higher in gaming: 0.45 versus 0.35, again a significant difference ($t = 2.20$, $p = 0.02$). These findings, together with the similarity between the FEI curves for gaming and socializing, provide convergent evidence that, in highly social contexts, FEI tends to peak for more popular influencers.

Finally, for two topics, life and skills, we see an apparent departure from the inverted-U shape, wherein FEI first decreases and then increases to its maximum. This "inverted-N" shape is driven by a substantial number of videos in these topics exhibiting purely decreasing FEI curves, which, when averaged with the more common inverted-U shape, yields an average that appears to have that initial peak. We hypothesize that the increased prevalence of purely decreasing curves for these two topics specifically is driven by the unique role of small-scale influencers in these areas. Life and skills videos are areas where authenticity is highly prized, and small scale influencers tend to be perceived as more authentic (Park et al. 2021). Thus, smaller influencers are marginally more effective at gaining attention in these contexts.

Implications for Firms and Influencers The variation in FEI curves observed across different types of hashtags has implications for both firms and influencers. Firms often design their influencer campaigns around a specific product, brand, or promotion, which means that each campaign may feature a different topic. In addition, when designing these campaigns, firms can choose how they engage with customers. Our results suggest that both topic and engagement tactic dictate the extent to which popularity matters in influencer selection. For social content, specifically socializing engagement tactics and the highly social context of gaming, FEI is higher for more popular influencers, suggesting that the firm will benefit more from sponsoring a more popular influencer. In other contexts, especially when authenticity and trust are concerns, FEI tends to peak for lower follower counts, suggesting that partnering with smaller influencers will yield higher returns. These findings also have implications for influencers, insofar as they suggest situations where an influencer may benefit more from seeking out followers. While many influencers view follower count as a statistic to be optimized, our findings show that, on the margin, more popular influencers may be less effective at promoting informative content, in areas like life and skills. On the other hand, for influencers who post about gaming, or more generally promote social content, additional followers matter much more.

5.3. Patterns Across Sponsored Campaigns

Until now, we have focused on establishing general FEI patterns across video types, regardless of whether these videos were part of a sponsored challenge or not.[31] In this section, we narrow our focus to just sponsored challenges, to understand if firm decisions in the structure of their influencer marketing campaigns drive different patterns of FEI. Recall that our data contain 30 sponsored challenges. As the content and structure of each of these campaigns are different, we expect to see different FEI curves. Therefore, for each campaign, we estimate an average FEI curve, based on the representations ($\boldsymbol{r}$) learned from the official videos posted under that challenge.

We observe three recurring patterns across campaigns: the FEI curves are either (a) inverted-U shaped, as before, implying an S-shaped counterfactual prediction curve; (b) relatively constant, corresponding to a linear counterfactual prediction curve; or (c) monotonically decreasing, corresponding to a concave counterfactual prediction curve. For example, Figure 10 shows the FEI and counterfactual prediction curves for Walmart's #UnwrapTheDeals campaign, which we discussed previously in the Introduction (Figure 1), and whose curve is similar to the average shape shown in Figure 6. On the plot, for reference, we also label regions of the x-axis corresponding to the six commonly defined tiers of influencers. For comparison, in Figures 11 and 12, we plot two other examples, from Apple and Dettol (the latter of which was the other example from Figure 1). Compared to the Walmart video, the Apple video, which was a workout tutorial with the hashtag #CloseYourRings, has a concave counterfactual prediction curve, suggesting that the highest

[31] As we illustrate in Web Appendix G, sponsored and organic hashtags are similar in terms of all observables, which is why we combined the two when thinking about broad topics and engagement tactics.

marginal effectiveness comes from smaller influencers. In contrast, the Dettol video's counterfactual prediction curve is nearly linear, suggesting that the marginal effectiveness is more consistent across influencer tiers. Even at high levels of popularity, this shape suggests that sizable impression gains are still possible with increasing follower sizes.

Although these are merely three case studies, we find that essentially every sponsored campaign follows one of these three patterns.[32] To better understand the relationship between the type of campaign and the three types of curves, we examined the content of the 30 campaigns, including their names, descriptions, and the types of videos posted under each. We give examples of campaigns for each curve type category in Table 4. We find that the S-shaped, linear, and concave growth shapes correspond to three salient campaign attributes: campaigns based on adding *special effects* to videos tend to exhibit S-shaped curves; challenges that encourage *self-expression* via either dancing or outfit change tend to exhibit linear curves; and *product demonstration* challenges tend to exhibit concave curves.

There are several explanations for these findings. The first relates to the audience size for particular types of campaigns: product demonstration videos, which exhibit a concave impression growth curve, can only attract an audience that is interested in their products, which is a smaller niche than, say, the community of TikTok users interested in dancing and outfit change videos. Consider, for example, the #MicellarRewind campaign from Garnier, which advertises a skin cleanser: regardless of whether this campaign is launched by a mega influencer or a more moderately popular one, there are only so many TikTok users interested in this product. Once that interested user base is saturated, there are limited gains from higher popularity, which leads to diminishing marginal returns. On the other hand, for special effects videos, the stronger S-shaped pattern suggests a more prominent threshold effect: a special effect will only catch on if a sufficiently popular influencer posts it. In this sense, the adoption of a special effect can be viewed as a classic adoption process, where the importance of hubs (i.e., highly connected individuals) has been well established (Goldenberg et al. 2009).

Another explanation for these diverse curve shapes is the quality of early followers. Previous work has established the immense heterogeneity in follower types (Farrell, Campbell, and Sands 2022). Although our model implicitly assumes that all followers are the same, an influencer's early followers may be more interested in their specific content than those who "hop on the bandwagon" once the influencer is already well established. This effect may be especially pronounced

[32] We show the complete set of results for all sponsored campaigns in Web Appendix H.

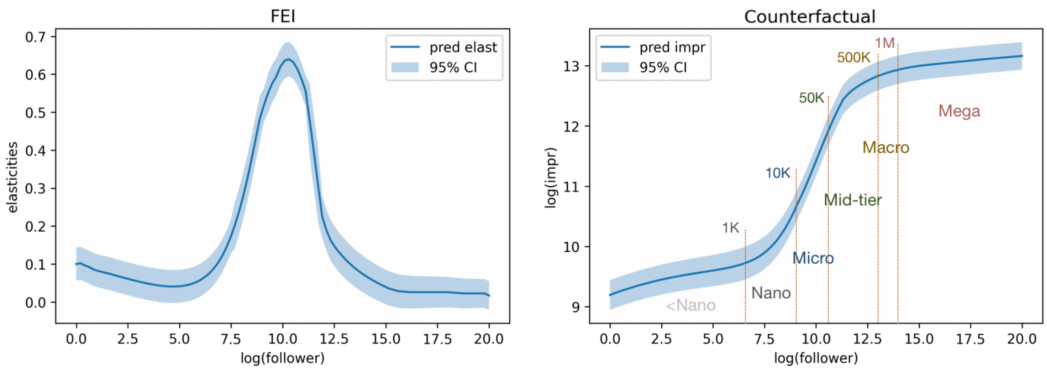

Figure 10: Walmart's #UnboxTheDeals FEI and Counterfactual Prediction Curves
On the left, the FEI curve, and on the right, the counterfactual prediction curve for impressions of Walmart's #UnwrapTheDeals video, given the (log) number of followers. The vertical lines in the right panel delineate the six influencer tiers on the x-axis.

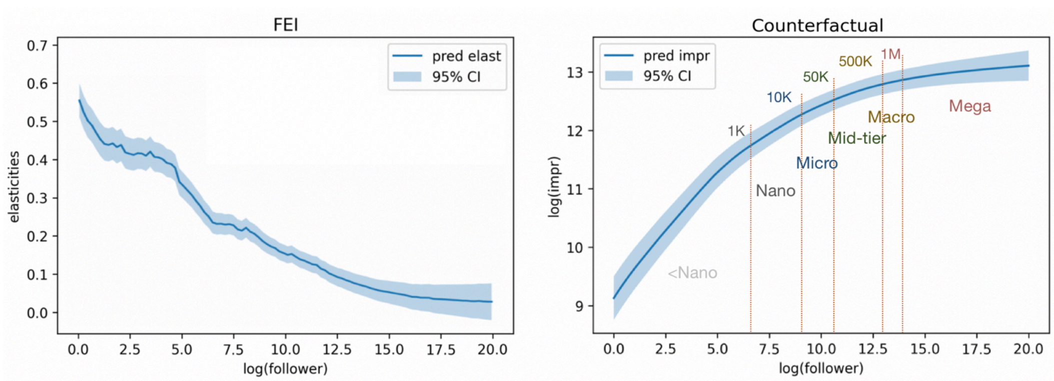

Figure 11: Apple's #CloseYourRings FEI and Counterfactual Prediction Curves
On the left, the FEI curve, and one the right, the counterfactual prediction curve for impressions of Apple's #CloseYourRings video, given the (log) number of followers. The vertical lines in the right panel delineate the six influencer tiers on the x-axis.

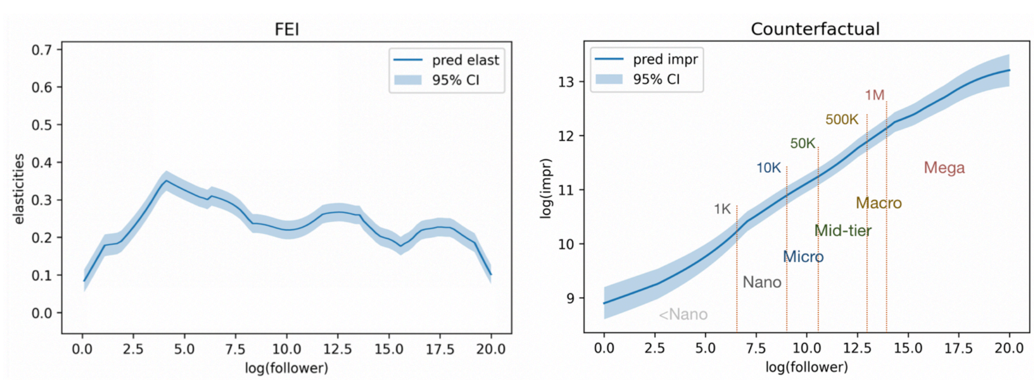

Figure 12: Dettol's #HandWashChallenge FEI and Counterfactual Prediction Curves
On the left, the FEI curve, and on the right, the counterfactual prediction curve for impressions of Dettol's #HandWashChallenge video, given the (log) number of followers. The vertical lines in the right panel delineate the six influencer tiers on the x-axis.

	Description	*Frames*
S-shaped		
#UnwrapTheDeals	Ready to #UnwrapTheDeals? Use our **#UnwrapTheDeals effect**, reveal your deal, then post with the hashtags #UnwrapTheDeals and #contest for a chance to win.	
#GetCrocd	It's time to #GetCrocd. Use our **Crocs effect** to try on three different pairs (with jibbitz!) and strut your stuff.	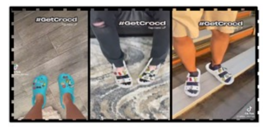
Linear		
#UptheBeat	Let's get down and #UpTheBeat. Challenge: throw on your favorite FILA look, learn this **dance** and put your own spin on it.	
#DoPacSun	You know how to dress, and you know how to **do transitions**. You did that. Now #DoPacSun. Show us your favorite PacSun looks, and let's see who owns this sound.	
Concave		
#MicellarRewind	**Rewind your routine** with Garnier Micellar Water! Show off your Micellar transformation in 3,2,1….	
#CancelTheNoise	This holiday season, it's time to tap into the joy of Bose QuietComfort Earbuds. Show us how you d**ouble tap to #CancelTheNoise** and feel it all.	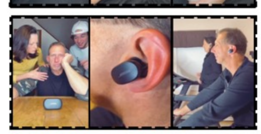

Table 4: Campaign Descriptions by Curve Type

For each of the curve types, we show two examples of campaigns with that curve type, including the hashtag name, a description of the campaign, and an example video. In total, 4 of the 30 sponsored hashtags are S-shaped, 16 are linear, and 10 are concave.

in contexts where the content is more specific, such as product demonstration videos. Thus, small-scale influencers in this area may have more engaged followers, who are more likely to share and engage with that content, leading to higher marginal returns for lower levels of popularity and concave impression growth curves. Indeed, we find this to be true: for product demonstration videos, the number of shares per follower is strictly decreasing in influencer popularity, ranging from 0.65 for sub-nano influencers (<1,000 followers), to 0.28 for mid-tier influencers (50,000-500,000 followers), to just 0.07 for mega influencers (>1 million followers). This pattern does not hold for the other campaign types. The prevalence of decreasing FEI curves in this context is also convergent with what we documented for life and skills videos, adding further evidence for the role of small influencers in contexts where connection and a sense of authenticity are especially important.

Until now, we have focused on S-shaped and concave counterfactual impression curves. These curve types are similar to those previously identified in the advertising literature. However, of our 30 sponsored campaigns, the most common curve type is linear: that is, where there seem to be almost no decreasing marginal returns to popularity. The fact that linear curves emerge for self-expression content is consistent with the idea that this type of content has broad, homogeneous appeal on TikTok, suggesting a minimal role for the effect of follower heterogeneity. Such broad appeal is supported in the data, where self-expression hashtags have the highest total number of impressions of any of the sponsored hashtags. More broadly, these linear curves indicate that for some content, a limit on the exposure potential for influencer advertising is yet to come. That we observe these response curves on TikTok, but not historically in advertising, may be attributable to the less mature content market on social media, which still experiences explosive growth in attention.

5.4. Optimal Influencer Selection

Our video-level FEI curves can guide firms in choosing an optimal popularity level for an influencer partner for a specific campaign. Doing so requires three inputs: first, the features of the desired video, x, from which we can compute the predicted FEI curve for that video; second, an assumption about how impressions translate into revenue; and third, the cost structure of the sponsorship. With these inputs, we can use the FEI to calculate an estimate of the marginal revenue associated with an additional follower. We can also estimate the marginal cost in terms of an additional follower. Setting the two equal yields the first-order conditions for optimal influencer selection.[33]

To illustrate, we focus on the case study of Walmart's #UnwrapTheDeals campaign, which was previously shown in the right panel of Figure 1. The focal question is, what level of popularity should Walmart target for an influencer to promote this campaign? To operationalize "this campaign," we use the features from the official sponsored videos under #UnwrapTheDeals, including the video depicted in Figure 1. We compute an average FEI curve for these official videos, and from that, a counterfactual prediction curve, which we previously displayed in Figure 10. To determine an optimal partnership based on this curve, we assume a dollar per impression rate of $0.02, which was set to match the average cost per view of a video ad on YouTube.[34] On the cost side, companies typically pay influencers based on their follower count.[35] Assuming that the payment is linear in the number of followers implies that the marginal cost of sponsoring an influencer for an ad post is constant. Following industry norms on TikTok, we assume a pay rate of $5 per 1,000 followers.[36] Based on these assumptions, we find that the optimal influencer for our Walmart example is a mid-tier influencer with around 275,000 followers.[37] Our analysis suggests that Walmart's maximum profit from paying the suggested influencer is 56% more than paying a micro influencer with 50,000 followers, and 300% more than paying a mega influencer with 1.5 million followers, which was Walmart's actual choice for the campaign.[38]

While the previous analysis was based on a single set of assumptions, the benefit of thinking about optimal influencer selection through FEI is that it can be generalized to any assumed cost and revenue structures. To illustrate, in Figure 13, rather than assuming a cost and revenue as we did before, we analyze under what conditions a given follower size would make sense, in terms of cost per thousand followers and revenue per impression. In each cell, we label the follower size (in thousands of followers) for the optimal influencer, given the cost on the x-axis and the revenue per impression on the y-axis. Note that in the context of the Walmart campaign, influencers with millions of followers (at the bottom-left corner of the heatmap) will be justified only if either their payment rate is low, or if the firm places a high value on TikTok impressions.

[33] This calculation makes an assumption equivalent to the "overlap" assumption in causal inference. Specifically, we assume that, for any content type, there are influencers across different popularities that can (or have) posted similar content. We do a robustness check, illustrating the plausibility of this assumption, in Web Appendix I.

[34] https://influencermarketinghub.com/how-much-do-YouTube-ads-cost/

[35] https://influencermarketinghub.com/influencer-rates/

[36] https://www.refersion.com/blog/influencer-marketing-cost/

[37] In reality, pay schedules may be nonlinear: it is possible, for instance, that marginal pay decreases as followers increase. If that is the case, the optimum would be greater than that implied by the linear cost structure assumed here.

[38] More details on profit maximization and the resulting profit for Walmart (as a function of followers) are in Web Appendix J.

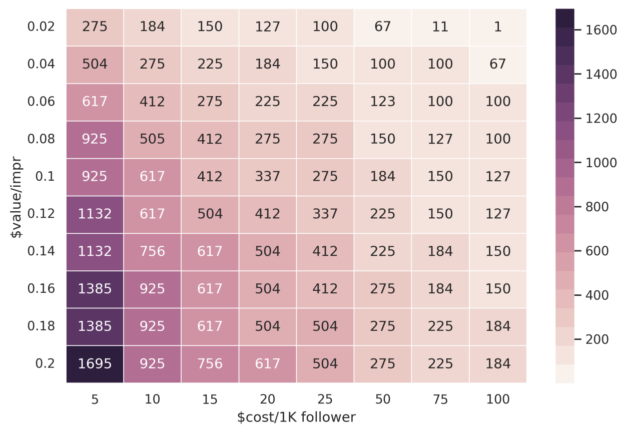

Figure 13: Walmart's Optimal Follower Size By Condition
The heatmap shows the optimal number of followers (in thousands) Walmart should target based on the cost per thousand followers on the x-axis, and revenue per impression on the y-axis.

Finally, we show how different growth curve shapes suggest different optimal influencer partners for different campaigns. Following the logic described previously, Walmart's counterfactual impression curve is S-shaped, suggesting that increasing followers in the mid-range yields the highest marginal gain in impressions. For concave curves, like Apple's #CloseYourRings campaign, increasing followers at the lower end of the distribution is more effective, suggesting that a collaboration with an influencer with a smaller number of followers may be optimal. Dettol's #HandWashChallenge, on the other hand, has an almost linear impression growth curve, suggesting that a collaboration with a mega influencer is optimal. With the assumptions regarding marginal revenue and marginal cost per follower as noted above, the optimal number of followers for Apple's campaign is around 100,000, substantially less than Walmart, while the optimal size for Dettol's campaign is around 505,000, substantially more than Walmart.

6. Conclusion

Influencer marketing is of central interest to modern marketers. In this work, we examine the relationship between an influencer's popularity and the resulting number of impressions of their posts on TikTok. Understanding this relationship is an essential input for firms considering investments in influencer marketing. Our work makes several contributions: first, we propose a modeling framework which combines representation learning with causal machine learning, to quantify the causal effect of creator popularity on the impressions for their content. From this framework, we derive the follower elasticity of impressions, or FEI, and showcase how it differs across many variables of interest to firms. We also illustrate how differences in FEI provide insights for the type of influencers with whom a firm should collaborate. Finally, our findings add to several ongoing debates about response to advertising and optimal advertising strategies in a new empirical context, social media advertising through influencers. Although our findings are based on TikTok data, our methodological framework can easily be generalized to other social media platforms.

Our results show that, absent causal analysis, firms may erroneously believe that very popular TikTok influencers are appropriate for all campaigns. In contrast, our causal analysis suggests that mid-tier influencers can be optimal to collaborate with for many campaigns. In addition, the content of a campaign plays an important role in how a focal video's impressions respond to its creator's popularity. We show that optimal influencer strategies vary in predictable ways, depending on how the firm is trying to engage with customers, what their campaigns are about, and how exactly their campaigns are structured. By characterizing these general FEI patterns and corresponding optimal partnerships along interpretable and actionable dimensions, we give managers a framework for influencer selection that goes beyond the simple case study analyses commonly pointed to in practice.[39]

While our research can help firms evaluate their influencer partnerships, it is not comprehensive. First, we only consider how influencers vary on the single dimension of popularity. Influencers can also vary in other ways, including how well the influencer matches a focal brand. Our content representations provide an avenue for firms to quantify the content that an influencer typically posts, but we do not provide any specific metric to assess how "good" of a match an influencer is for a company's brand image. Second, we analyze the case where a firm is collaborating with a single influencer. In contexts where multiple influencers are sponsored, there can be interactions among influencers that determine the eventual success of a campaign. However, the concept of FEI still provides a good starting point for the analysis. Third, while our data are a representative cross-section of trending topics on TikTok, we do not have complete panel data on influencers and their followers. Without such panel data, it is difficult to fully characterize what influencers typically post about and how such topics may evolve over time.

There are also several limitations to our framework. One limitation is that computing a precise FEI curve, and thus determining optimal influencer popularity, requires having access to content

[39]For example: https://www.tiktok.com/business/en-US/blog/branded-hashtag-challenge-harness-the-power-of-participation

(i.e., the influencer's post) in advance. While in some cases this is reasonable, like when firms elevate existing, organic content to official, sponsored status, in other cases, firms may not always have representative content before soliciting partnerships. In these cases, our general guidelines can still help firms think broadly about the importance of popularity for their intended campaign, even absent a specific predicted FEI. Finally, our framework is limited in its ability to reason about content: while our content representations control for content, they are not directly interpretable, and thus, our ability to directly optimize content to achieve a particular goal is limited. We hope that future work will build on the methods and metrics developed in our work to explore these and other related questions in more depth.

References

Akpinar, Ezgi and Jonah Berger (2017), "Valuable Virality," *Journal of Marketing Research*, 54 (2), 318–330.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin (1996), "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91 (434), 444–455.

Athey, Susan, Julie Tibshirani, and Stefan Wager (2019), "Generalized Random Forests," *The Annals of Statistics*, 47 (2), 1148–1178.

Berger, Jonah and Katherine L. Milkman (2012), "What Makes Online Content Viral?," *Journal of Marketing Research*, 49 (2), 192–205.

Burnap, Alex, John R. Hauser, and Artem Timoshenko (2023), "Product Aesthetic Design: A Machine Learning Augmentation," *Marketing Science*, Articles in Advance.

Cannon, Hugh M., John D. Leckenby, and Avery Abernethy (2002), "Beyond Effective Frequency: Evaluating Media Schedules Using Frequency Value Planning," *Journal of Advertising Research*, 42 (6), 33–46.

Centola, Damon (2010), "The Spread of Behavior in An Online Social Network Experiment," *Science*, 329 (5996), 1194–1197.

Danaher, Peter J. and Tracey S. Dagger (2013), "Comparing the Relative Effectiveness of Advertising Channels: A Case Study of a Multimedia Blitz Campaign," *Journal of Marketing Research*, 50 (4), 517–534.

Dekimpe, Marnik G. and Dominique M. Hanssens (2007), "Advertising Response Models," *Handbook of Advertising*, pages 247–263.

Dew, Ryan, Asim Ansari, and Olivier Toubia (2022), "Letting Logos Speak: Leveraging Multiview Representation Learning for Data-driven Branding and Logo Design," *Marketing Science*, 41 (2), 401–425.

Dolan, Rebecca, Jodie Conduit, Catherine Frethey-Bentham, John Fahy, and Steve Goodman (2019), "Social Media Engagement Behavior: A Framework for Engaging Customers through Social Media Content," *European Journal of Marketing*, 53 (10), 2213–2243.

Doumit, Gaby, Melina Gattellari, Jeremy Grimshaw, Mary Ann O'Brien et al. (2007), "Local Opinion Leaders: Effects on Professional Practice and Health Care Outcomes," *Cochrane Database of Systematic Reviews*, 1 (1).

Erhan, Dumitru, Aaron Courville, Yoshua Bengio, and Pascal Vincent (2010), "Why Does Unsupervised Pre-training Help Deep Learning?," *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 201–208.

Farrell, Justine Rapp, Colin Campbell, and Sean Sands (2022), "What Drives Consumers to Engage with Influencers?: Segmenting Consumer Response to Influencers: Insights for Managing Social-media Relationships," *Journal of Advertising Research*, 62 (1), 35–48.

Farrell, Max H., Tengyuan Liang, and Sanjog Misra (2021), "Deep Neural Networks for Estimation and Inference," *Econometrica*, 89 (1), 181–213.

Goldenberg, Jacob, Sangman Han, Donald R Lehmann, and Jae Weon Hong (2009), "The Role of Hubs in the Adoption Process," *Journal of Marketing*, 73 (2), 1–13.

Gordon, Brett R., Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky (2019), "A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook," *Marketing Science*, 38 (2), 193–225.

Granovetter, Mark S. (1973), "The Strength of Weak Ties," *American Journal of Sociology*, 78 (6), 1360–1380.

Guo, Tong, S. Sriram, and Puneet Manchanda (2021), "The Effect of Information Disclosure on Industry Payments to Physicians," *Journal of Marketing Research*, 58 (1), 115–140.

Hanssens, Dominique M. (2018), "The Value of Empirical Generalizations in Marketing," *Journal of the Academy of Marketing Science*, 46, 6–8.

Hartford, Jason, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy (2017), "Deep IV: A Flexible Approach for Counterfactual Prediction," *Proceedings of the International Conference on Machine Learning*, pages 1414–1423.

Hausman, Jerry A "Valuation of New Goods Under Perfect and Imperfect Competition," "The Economics of New Goods," pages 207–248, University of Chicago Press (1996).

Hsu, Wei-Ning, Yu Zhang, and James Glass "Learning Latent Representations for Speech Generation and Transformation," "Interspeech," pages 1273–1277 (2017).

Hughes, Christian, Vanitha Swaminathan, and Gillian Brooks (2019), "Driving Brand Engagement through Online Social Influencers: An Empirical Investigation of Sponsored Blogging Campaigns," *Journal of Marketing*, 83 (5), 78–96.

Johansson, Johny K. (1979), "Advertising and the S-curve: A New Approach," *Journal of Marketing Research*, 16 (3), 346–354.

Jones, John Philip (1995), "Single-source Research Begins to Fulfill Its Promise," *Journal of Advertising Research*, 35 (3), 9–17.

Katona, Zsolt, Peter Pal Zubcsek, and Miklos Sarvary (2011), "Network Effects and Personal Influences: The Diffusion of an Online Social Network," *Journal of Marketing Research*, 48 (3), 425–443.

Keating, Nancy L., John Z. Ayanian, Paul D. Cleary, and Peter V. Marsden (2007), "Factors Affecting Influential Discussions among Physicians: A Social Network Analysis of a Primary Care Practice," *Journal of General Internal Medicine*, 22, 794–798.

Kingma, Diederik P. and Max Welling (2013), "Auto-encoding Variational Bayes," *arXiv preprint arXiv:1312.6114*.

Kuo, David, David R. Gifford, and Michael D. Stein (1998), "Curbside Consultation Practices and Attitudes among Primary Care Physicians and Medical Subspecialists," *Journal of the American Medical Association*, 280 (10), 905–909.

Lee, Dokyun, Kartik Hosanagar, and Harikesh S. Nair (2018), "Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook," *Management Science*, 64 (11), 5105–5131.

Leung, Fine F., Flora F. Gu, Yiwei Li, Jonathan Z. Zhang, and Robert W Palmatier (2022), "Influencer Marketing Effectiveness," *Journal of Marketing*, 86 (6), 93–115.

Leung, Fine F., Flora F. Gu, and Robert W. Palmatier (2022), "Online Influencer Marketing," *Journal of the Academy of Marketing Science*, pages 1–26.

Li, Xi, Mengze Shi, and Xin Shane Wang (2019), "Video Mining: Measuring Visual Information Using Automatic Methods," *International Journal of Research in Marketing*, 36 (2), 216–231.

Liu-Thompkins, Yuping (2012), "Seeding Viral Content: The Role of Message and Network Factors," *Journal of Advertising Research*, 52 (4), 465–478.

Lodish, Leonard M, Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, and Mary Ellen Stevens (1995), "How TV Advertising Works: A Meta-analysis of 389 Real World Split Cable TV Advertising Experiments," *Journal of Marketing Research*, 32 (2), 125–139.

Newey, Whitney K. and James L. Powell (2003), "Instrumental Variable Estimation of Nonparametric Models," *Econometrica*, 71 (5), 1565–1578.

Park, Jiwoon, Ji Min Lee, Vikki Yiqi Xiong, Felix Septianto, and Yuri Seo (2021), "David and Goliath: When and Why Micro-Influencers are More Persuasive than Mega-Influencers," *Journal of Advertising*, 50 (5), 584–602.

Pearl, Judea (2009), *Causality* Cambridge University Press.

Rajaram, Prashant and Puneet Manchanda (2023), "Video Influencers: Unboxing the Mystique," *arXiv preprint arXiv:2012.12311*.

Rao, Ambar G. and Peter B. Miller (1975), "Advertising-sales Response Functions," *Journal of Advertising Research*, 15 (2), 7–15.

Rezende, Danilo Jimenez, Shakir Mohamed, and Daan Wierstra "Stochastic Backpropagation and Approximate Inference in Deep Generative Models," "International Conference on machine learning," pages 1278–1286, PMLR (2014).

Sethuraman, Raj, Gerard J. Tellis, and Richard A. Briesch (2011), "How Well Does Advertising Work? Generalizations from Meta-analysis of Brand Advertising Elasticities," *Journal of Marketing Research*, 48 (3), 457–471.

Shapiro, Bradley T., Günter J. Hitsch, and Anna E. Tuchman (2021), "TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands," *Econometrica*, 89 (4), 1855–1879.

Simon, Julian L. and Johan Arndt (1980), "The Shape of the Advertising Response Function.," *Journal of Advertising Research*.

Susarla, Anjana, Jeong-Ha Oh, and Yong Tan (2012), "Social Networks and the Diffusion of User-generated Content: Evidence from YouTube," *Information Systems Research*, 23 (1), 23–41.

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

Tudón, José (2022), "Distilling Network Effects from Steam," *Quantitative Marketing and Economics*, 20 (3), 293–312.

Vakratsas, Demetrios, Fred M. Feinberg, Frank M. Bass, and Gurumurthy Kalyanaram (2004), "The Shape of Advertising Response Functions Revisited: A Model of Dynamic Probabilistic Thresholds," *Marketing Science*, 23 (1), 109–119.

Valsesia, Francesca, Davide Proserpio, and Joseph C Nunes (2020), "The Positive Effect of Not Following Others on Social Media," *Journal of Marketing Research*, 57 (6), 1152–1168.

Wies, Simone, Alexander Bleier, and Alexander Edeling (2023), "Finding Goldilocks Influencers: How Follower Count Drives Social Media Engagement," *Journal of Marketing*, 87 (3), 383–405.

Yang, Jeremy, Juanjuan Zhang, and Yuhan Zhang (2021), "First Law of Motion: Influencer Video Advertising on TikTok," *Available at SSRN 3815124*.

Yao, Rong, Chongdang Liu, Linxuan Zhang, and Peng Peng (2019), "Unsupervised Anomaly Detection Using Variational Auto-encoder Based Feature Extraction," pages 1–7.

Yoganarasimhan, Hema (2012), "Impact of Social Network Structure on Content Propagation: A Study Using YouTube Data," *Quantitative Marketing and Economics*, 10, 111–150.

Zhang, Mengxia and Lan Luo (2023), "Can Consumer-posted Photos Serve as a Leading Indicator of Restaurant Survival? Evidence from Yelp," *Management Science*, 69 (1), 25–50.