Are Exposure Bucks Worth It?
Evidence from Instagram

Andrew T. Ching

Johns Hopkins University

Carey Business School

Email: andrew.ching@jhu.edu

Daniel Goetz

University of Toronto

Rotman School of Management

Email: daniel.goetz@rotman.utoronto.ca

March 20, 2025

Abstract

Brands often use the promise of exposure to their large audiences to incentivize content contributions from artists and creators. This study quantifies the value of exposure for independent creators that supply content to large brands' pages on social media. We leverage Instagram's recently introduced co-authorship feature, which lets brands credit creators in shared posts, and gather data on co-authored posts between 260 large beauty/clothing brands and small creators. We estimate changes in creators' follower counts after the co-authorship treatment in a difference-in-differences analysis, and address post content endogeneity by matching treatment and control posts on post-level visual content embeddings recovered from a Large Language Model. We find that exposure to a brand's audience increases a creator's cumulative followers by only 0.47%, or roughly one week's worth of growth. Although larger brand audiences lead to more growth all else equal, larger brands tend to have less engaged followers and lower quality co-authored posts, and we find constant treatment effects from co-authoring with brands with between 0.7m and 7m followers. The top 10% of collaborations average only 2.1% follower growth for the creator. Our findings have implications for understanding creators' dynamic behavior, incentivizing the creation of user-generated content, and informing creator-brand contract negotiations.

Keywords: Collaboration, Exposure Externality, Content Creation, Creator Economy, Influencer Marketing, Owned Social Media, Large Language Models, Image Processing

1 Introduction

> In exchange for your artwork, we will publish it on the birthday of the player assigned to you, with a mention on Instagram (8,800,000 followers) [. . .] We will tag you in the photo itself and in the caption.
>
> *– Solicitation of artwork by AFC Ajax, the Netherlands' largest football club, Sept. 2024*

> Exposure is something you die from, not something you get paid with.
>
> *– User ZombieButch, post on reddit.com/r/ArtistLounge, Sept. 2023*

Consider a brand that wishes to source artwork for an ad. One common approach is to solicit contributions by offering to credit the chosen contribution in a visible way. In the leading quote, a football club with a large social media following asks for club-related art, and in exchange, the artist will have their art shown and visibly credited to the millions of followers of the football club on social media. For brands that have small budgets, offering the creator exposure to the brand's audience as part of their compensation is seen as a very cost-effective motivator.[1] However, even for large brands such as Apple and Nike, offering visibility to collaborating artists is a common practice.[2] On the "largest stage in the world," the Superbowl Halftime show, performers are famously not paid for their performances, but the exposure to a large, engaged audience is claimed to provide them with enormous value.[3]

Despite brands' reliance on exposure to solicit creative contributions, independent creators themselves have a much more mixed view than brands on the value of exposure (Somerfield, 2023).[4] Exposure can be conceptualized as a type of extrinsic motivation, whereby a creator may expand their own audience through their increased visibility to a brand's audience (Roberts et al., 2014; Somerfield, 2023). There has been substantial work on other extrinsic motivators

[1] See "What are some creative ways to incentivize and reward user-generated content creators?", *LinkedIn*, 2024. Amongst the expert contributors to the article, most mention exposure as a way to reward creators, especially if the brand is on a budget.

[2] See "Connect Deeper, Sell Smarter: 7 UGC Campaigns That Transformed Brand-Customer Relationships", *PepperContent*, January 1, 2024. In print media, see *The New York Times*, "The Paradox of Art as Work", May 2014, "Magazines that once paid by the word and employed healthy rosters of staff writers are now recruiting readers to contribute articles, offering to compensate them with exposure and 'prestige."

[3] See "No, Performers Don't Get Paid for the Super Bowl Halftime Show — Here's Why," *People Magazine*, April 2024.

[4] For positive impressions of being credited by firms, see The State of User-Generated Video, 2024-05, " For exposure and awareness, it's good to get your name around and get your brand across." For negative impressions, see https://www.helloraderco.com/sacrificing-time-money-exposure-bucks/ and https://amandajspedding.com/2019/05/27/exposing-exposure-bucks-hint-its-not-a-real-thing/

to supply creative content like prestige, awards, and direct payments (Luca, 2015; Gallus, 2017; Burtch et al., 2022; Khernamnuai et al., 2018; Aridor et al., 2024). However, to our knowledge, there is no existing research which tries to quantify the value of exposure to creators due to at least three obstacles. First, one needs to observe a large, representative sample of brand-creator collaborations with creators being visibly credited. Second, one must be able to measure the benefits of exposure to the creator. Third, one needs to resolve the selection problem, whereby creators' contributed content may be endogeneous.

The goal of our study is to provide causal evidence about the value of exposure to creators (i.e., so-called "exposure bucks"). To address the first challenge, we turn to Instagram. More specifically, we take advantage of a new feature introduced by Instagram in late 2021 which allows a post to be explicitly recognized as co-authored by two accounts. This feature allows brands to visibly give credit to the co-authoring creator for the content they provide, and requires creators to agree to be credited, ensuring that the post represents a true two-sided collaboration. We collect a large, unique dataset consisting of the entire post history of the top 260 fashion and beauty brands on Instagram, from which we are able to identify almost 1,000 posts which are co-authored with independent creators.

Regarding the second challenge, we use the fact that the co-authored posts are simultaneously shared to the followers of both accounts. When a creator has significantly fewer followers than the brand they work with, such a collaboration could amount to a substantial increase in audience for the creator. For example, in Figure 1, the creator wrath_____ (27.6 thousand followers) co-authors with the fashion brand Converse (10.6 million followers), leading her work to appear in the Instagram feed of millions of potential new followers. As a result, her own follower count may expand, giving the creator a larger audience for their own content going forward. Building their follower base is the primary objective for influencers in social media (Abidin, 2016; Smith and Fischer, 2021), as it means more attention which may lead to more opportunities for monetization (Carpentier, 2023). Therefore, to address the second challenge of measuring exposure benefits, we collect the daily number of followers and complete post history for each creator who has collaborated with a brand —this allows us to quantify to what extent creators gain followers after the co-authored post.

The last empirical challenge which we address is the selection problem, where the content of co-authored posts is endogenous. To address the issue of non-random content, we propose a novel matching approach by using the powerful computer vision capabilities of the ChatGPT 4o

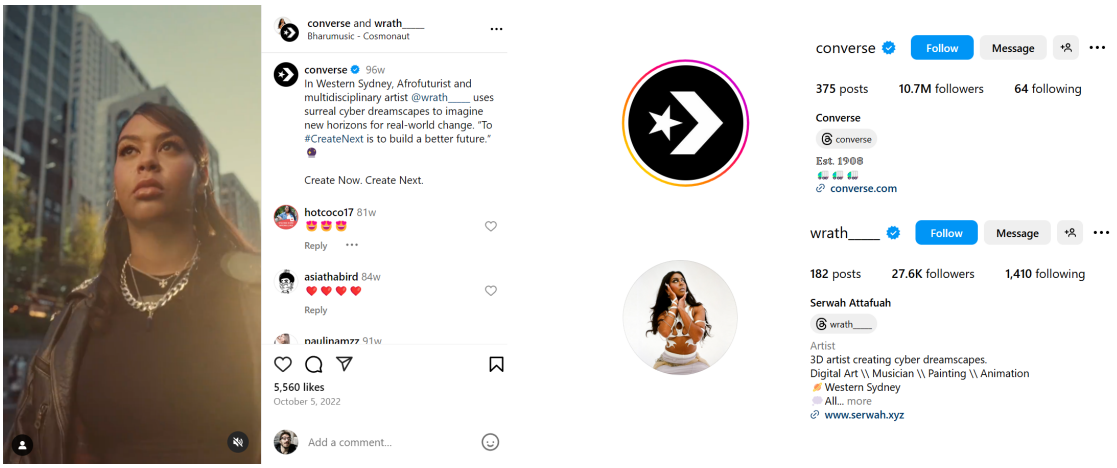


Figure 1: Example of a collaborative post

Note: The left image is the collaborative post, which is credited to both the brand (converse) and the creator (wrath_____). At time of posting, it is simultaneously displayed to the followers of both the brand and the creator. From the profile information, this implies an audience of over 10 million converse followers for the creator's post.

Large Language Model (LLM) to extract textual descriptions of images and video from creators' posts. We convert these descriptions to numerical embeddings, and match treated creator posts to posts by control creators with similar audiences and embeddings, posted at similar times of the day, week, and year, using a gradient-boosted propensity score matching function. We then conduct a difference-in-differences analysis to estimate the average treatment effect on the treated, leveraging our daily follower growth data and quasi-random post timing to rule out pre-trends and contemporaneous shocks.

Note that a priori, it is difficult for a researcher to anticipate the magnitude of such exposure benefits in this context. Although being exposed to a large audience should matter, the duration of the exposure can also play a crucial role. In social media like Instagram, a post typically has a very short "life-cycle" because new posts are constantly created to compete for users' attention, and brands also post new items regularly, pushing down the position of a post collaborated with a creator. Besides, such a co-authored post maybe primarily designed to help the brand instead of promoting the creator. Hence, even though a brand may have a large followers size, a co-authored post could still lead to a small exposure benefit to the creator.

Our results show that there is some value to creators from exposure to a brand's audience on average, with a total increase in the creator's follower count of 0.47% after the co-authored post. The positive directional effect of co-authorship is in line with previous research on how news content picked up by an aggregator leads to an expansion in traffic to the original newspaper

(Calzada and Gil, 2020). Our estimated exposure benefit might appear to be small given the audience size: the average brand has 7.6 million followers—over 120 times more followers than the average co-authoring creator—but we estimate that the creator acquires only between 115 and 277 new audience members after the average co-authored post. However, if one conceptualizes the co-authored post as an ad for the creator and a new follow as a conversion, our estimated new follow rate is in line with recent estimates of intention-to-treat advertising effects (Gordon et al., 2019).

We then turn to study some substantive issues. First, we address a key point in the debate on exposure: Does working with brands with larger audiences lead to a stronger exposure benefit? We estimate heterogeneous treatment effects, allowing the effect of co-authorship on exposure to vary by brand, post, and creator attributes, and find that all else equal, co-authoring with brands with larger follower counts benefits the creator more. Moving from the mean brand audience to one standard deviation above the mean in log followers roughly doubles the baseline lift. However, larger brands with more followers also tend to have lower follower engagement, and their co-authored posts gather relatively few likes. When we hold fixed only creator attributes, allowing large brands' engagement and post attributes to vary as they do in the data, we find that working with brands from 710,000 followers (American Vintage) to 8.6 million followers (Guess Jeans) all lead to roughly similar lifts in followers for the creator. Thus, while brand's audience size is important for exposure in isolation, taking into account other brand features that correlate with size suggest that a wide range of brand audience sizes are equally beneficial for exposure.

Second, we turn to creator heterogeneity, and ask whether proxies for creator quality increase the value of exposure. We find that larger creators benefit proportionately less, because the absolute increase in followers is approximately constant across creators. Moreover, we find that creators with more engaged audiences do not experience significantly greater benefits, suggesting a limited role for creator quality signals.

Third, we show that the more different a creator's co-authored post is to their usual content, the smaller the lift in new followers, and with a large enough difference the effect can even turn slightly negative. This finding is driven by the creator's own followers leaving, which spotlights a hidden cost of exposure: supplying content to a brand may alienate a creator's current audience if they treat the co-authored post as a sponsored post, and interpret that the creator may have lost their authenticity.

Finally, we use our heterogeneity results to estimate the full distribution of treatment effects,

and find that the top 10% of collaborations lead to only a 2.1% lift in followers on average. Using a back-of-the-envelope assumption on how creators monetize their social media accounts, we estimate the full distribution of the dollar value from exposure for our treated creators. Exposure is most valuable for small creators making content for large brands with highly engaged audiences, and for creators who frequently engage in sponsored posting pre-treatment; this is because these creators can leverage their increased follower count into a higher price-per-post.

1.1 Literature review

This paper directly contributes to a growing literature on influencer incentives. Early theory papers sought to explain the behavior of influencers, and coalesced around the idea that influencers build an organic following by offering entertaining content or informative product reviews (Fainmesser and Galeotti, 2021; Mitchell, 2021; Pei and Mayzlin, 2022), and then monetize this following by creating sponsored content for their followers. Recent empirical papers have found support for the idea that creating sponsored content "harvests" the goodwill and credibility developed by influencers (Cheng and Zhang, 2024; Bairathi and Lambrecht, 2023). We add to this literature by empirically exploring the possibility that influencers can leverage brand audiences to build their own followings through collaborations.

Beyond influencer marketing, there is a long literature on creators making user-generated content (UGC). This literature is quite broad, as UGC spans content from online reviews to forum comments to Instagram videos. We directly connect to a series of qualitative papers that ask why creators make UGC directly for brands. Roberts et al. (2014) find that individuals who co-create with video game companies are driven by the prospect of opportunities or goals to do more work in the video game industry, Berthon et al. (2008) identify self-promotion as a goal of creating unpaid ads for companies, and Somerfield (2023) show that gaining exposure motivates UGC, especially from creators who also work inside the social media industry. However, these papers have not tried to quantify those benefits, which is the goal of our paper.

From the perspective of platforms, motivating UGC is important, and there has been substantial work on understanding incentives to create UGC. There are two types of incentives: extrinsic rewards and intrinsic rewards. When it comes to extrinsic rewards, visible status or hierarchy markers on an online platform have been shown to be highly effective at generating more UGC in contexts including the creation of news by citizen journalists (Sen et al., 2024), contributions on a knowledge sharing website (Goes et al., 2016), editing articles on Wikipedia (Gallus, 2017), and

writing comments on Reddit (Burtch et al., 2022). Meanwhile, monetary payments have a mixed record of success, as paying people to contribute reviews may decrease new reviews (Sun et al., 2017) or increase review quantity but reduce review informativeness (Khernamnuai et al., 2018). Being closely related to the motivation of our research, Ghosh and McAfee (2011) show that it is theoretically possible to incentivize high-quality UGC from users who are eager for attention, by eliciting visible measures of UGC quality from platform participants.

Finally, our paper connects to a small literature that asks whether the creators of content can benefit from that content being shown on a larger platform. In a natural experiment where Google News was shut down in Spain in 2014, both Calzada and Gil (2020) and Athey et al. (2021) show that news aggregators increase page visits to news outlets, with the former showing that the aggregators increase consumers' awareness of news outlets' contents.

The rest of the paper is organized as follows. Section 2 discusses the data and institutional background. Section 3 presents our econometric model, explains the identification issues, and discuss results. Section 4 study the heterogeneous treatment effects and shed light on the situations under which creators can benefit more from such collaboration. Section 5 is about the implications for practice. We conclude in section 6.

2 Data and Context

2.1 Instagram and co-authored posts

Instagram is a platform for photo and video sharing with over 2 billion active monthly users, allowing users to follow accounts whose content they enjoy consuming, and, in turn, be followed by other users.[5] Instagram serves as a significant hub for social media marketing, where brands share content to their followers to promote engagement with the brand, a channel referred to as "owned" social media. Thousands of brands are active on Instagram, with many brands pursuing a "social media first" marketing strategy and building a loyal following through content on their brand account that maintains customer awareness and engagement.[6]

Collaborative or co-authored posts were introduced by Instagram in late 2021, allowing two accounts to co-author a post and share it with both of their follower bases simultaneously. From the creator's perspective, co-authored posts effectively increase their visibility, as the creator's

[5] See "60+ Instagram Statistics Marketers Need to Know in 2024,", *SocialPilot*, September 10, 2024.
[6] See "Launching Consumer Brands in the Age of Social Media,", *Meltwater*, November 11, 2018.

account name clearly appears alongside the brand's name (see Figure 1). We find that most co-authored posts are one-off collaborations between a brand and a creator, which provides a clean quasi-experimental setup for understanding the effect of a boost in visibility for the creator. From the brand's perspective, the collaborative post saves them from creating their own content entirely, and may be perceived as more authentic or community-building.[7]

Apart from co-authorship, tagging can be used to credit creator contributions on Instagram, and partnership/branded posts can be used to attach a brand name to a creator's post. In Appendix A.1, we discuss why co-authored posts have several advantages for a study of the value of exposure compared to tagging or partnership posts.

2.2 Data collection

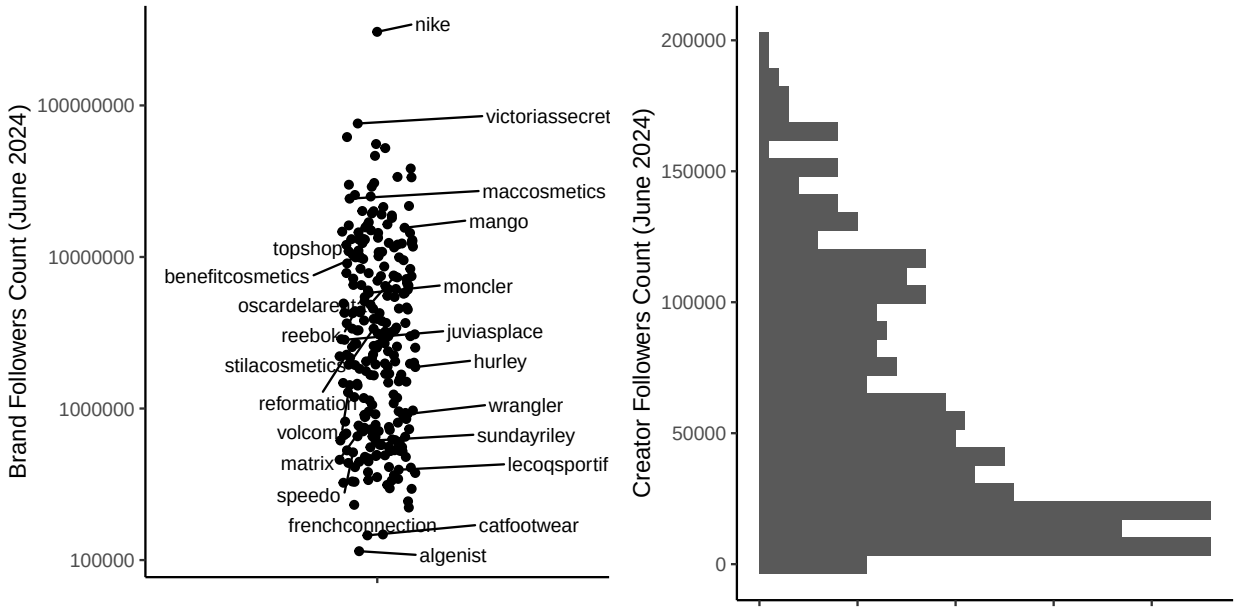
Brands. We start by identifying the top brands in beauty, clothing, shoes, and fashion that are active on Instagram. We use several websites, including FashionUnited.com and StarNGage.com, which list brands in decreasing order of number of followers on Instagram, and supplement this list with web searches.[8] We gather 260 brands in total, primarily based in the United States. The distribution of the number of followers across brands is shown in Figure 2a; brands range from as few as 114,000 followers for the beauty brand Algenist to over 300,000,000 followers for Nike.

Collaborative Posts and Co-authoring Creators. For each brand, we use a web-scraper to pull the entire post history of that brand on Instagram back to January 2021. We recover 15,340 co-authored posts and extract the complete set of 4,320 co-authoring accounts. We keep accounts that have 200,000 or fewer followers, that are categorized as a personal or creator account, and that are not co-owned by one of the 260 brands. For instance, @nikewomen or @nikebasketball would be dropped, as they are owned and operated by Nike. There are 953 Instagram accounts meeting our inclusion criteria who are treated with a co-authorship post, which we refer to as the co-authoring creators. Note that we exclude accounts with more than 200,000 followers because we expect that smaller creators will likely benefit more from exposure, making the effect more pronounced and measurable.

Also, we are interested in the effect of exposure from a co-authored post, and not the effect of

[7] See "How to Use Instagram Collab Posts to Get More Engagement," *Hootsuite*, August 14, 2023.

[8] See https://fashionunited.com/i/instagram-fashion-index. and https://starngage.com/plus/en/brand/ranking., accessed on 10/10/2024.



(a) Brand follower counts

(b) Creator follower counts

Figure 2: Distribution of follower counts on Instagram

Note: In the left figure, each point is a particular brand, and labels are the brand's Instagram account name. Points are jittered horizontally for readability. The interquartile range for the number of followers is [759,000, 8,211,000]. In the right figure, we plot the histogram of follower counts for creators. The interquartile range is [20,000, 96,000].

long-term partnerships between brands and creators. We assume that if a creator has more than one co-authored post with a brand, then their relationship is a long-term partnership, and we drop them from the data.[9] This leaves us with 692 treated creators who only have a single co-authored post with one of the 260 brands. We plot the distribution of creator follower numbers in Figure 2b.

Follower data. The key outcome of our study is how the number of followers for each creator changes after they have co-authored a post with a brand. We collect the complete followership history for each co-authoring creator account from SocialBlade. SocialBlade is a social media marketing company that helps brands identify influencers for partnerships. Importantly, for our purposes, they collect daily follower count information for tens of millions of accounts. Since the number of followers of each account tend to increase over time and is therefore non-stationary, we work primarily with the follower growth rate or the change in followers, to make the data amenable to fixed effects analysis.

We trim out any creators whose posts are preceded by "viral" follower growth, to remove col-

[9]These creators are also treated more than once, making it difficult to interpret the treatment effects.

laborations between brands and creators that may be sparked by a creator's sharply rising popularity. We do so by dropping any creator where, within the 30 days prior to their co-authored post, the creator experienced a jump in follower growth that was in the top 1% of daily follower growth rates for that creator. After this trimming step, we are left with 648 creators in the sample.

Posting data. We gather the complete posting history for all 648 treated creators from January 2021 through August 2024. In addition to post level metrics such as likes, comments, and video views, we extract the complete set of media elements including the caption, images, and videos associated with the post.

Pool of control creators and posts. Ideally, for each treated unit, we want to find a control unit which satisfies the parallel trend assumption and make a solo post that looks identical to the treated unit's co-authored post. To find such control creators and posts, we turn to two sources: (a) posts by treated creators pre-treatment, and (b) posts by never-treated creators (Callaway and Sant'Anna, 2021; Roth et al., 2023).

For the former, we restrict to non-collaborative posts created at least 90 days prior to their collaborative post. This restriction is made to avoid overlap in our observations of follower growth, as we will look at the evolution of follower growth in a window of 3 weeks before and 6 weeks after a post. This yields approximately 210,000 control posts by treated creators pre-treatment.

For the latter, our criteria for gathering these additional control posts are that (1) the posting account is sourced from the set of "related accounts" supplied by Instagram for at least one treated account; (2) the posting account is similar in category (e.g. artist, athlete) and follower count to a treated account; and (3) the posting account has follower data from SocialBlade. We are able to collect an additional 220,000 control posts by 625 never-treated creators.[10] Additional details are provided in Appendix A.2. The 430,000 total control posts comprise our complete control pool. In robustness checks in Section 3.4.2, we consider only pre-treatment posts by treated creators or only posts by never-treated creators as alternative control pools, and our main findings remain unchanged.

[10]During the entire sample period, the never-treated creators each have 352 posts on average, while the treated creators have approximately 250,000 posts total (including post-treatment), leading to 358 posts on average.

2.3 Controlling for post content

To estimate the effect of co-authorship on follower growth, one cannot simply match creators who co-author to similar creators who do not co-author. While this would create a set of control creators who resemble treated creators on average, the content and timing of the treated creators' co-authored posts could be very different than the average content for the controls. This would lead to a problem where comparing follower outcomes before and after the co-authored post with the average control post would conflate the effect of the content with the effect of exposure.

Instead, we pursue a strategy of matching treated creator-posts to control creator-posts, we create a matched sample of 648 solo posts by control creators that closely resemble the 648 co-authored posts by treated creators. To achieve this, we take the following three steps. First, we convert the rich, complex visual content of the images and videos into a numeric embedding space. Second, we reduce the dimension of these embeddings in a way that preserves the attributes of the content, but is also interpretable by humans. Third, we estimate a matching function that predicts treatment status based on post and creator features, and validate that the function yields control posts and creators that are observably similar to treated posts and creators. Combined with the assumption that the content features we extract are sufficient to solve the selection problem, we can then isolate the effect of co-authorship.

2.3.1 Embedding post content

The literature has taken two approaches to capture visual data in studies of social media content: (i) hand-selecting a set of features such as facial emotions and color hue and then using machine learning libraries to extract those features for each image (Zhang et al., 2022; Cheng and Zhang, 2024) , or (ii) using pre-trained image-recognition models such as VGG-19 (Simonyan, 2014) to embed visual content as numeric matrices before conducting dimension reduction (Tian et al., 2024; Yang et al., 2023; Rajaram and Manchanda, 2020).

Our proposed approach improves the latter method in two ways. We start by using the vision capabilities of ChatGPT 4o to recover textual descriptions of the visual content of a post. ChatGPT 4o is a large language model (LLM) with substantially improved image recognition capabilities relative to previous models. Compared to alternatives like CLIP or VGG-19 which prioritize object detection within an image, ChatGPT retains more information about activities (e.g., a person playing with a dog, versus simply a person and a dog) and context (e.g., exercising outdoors

versus in a gym). For each image in a post, we prompt ChatGPT 4o with "Describe this image," and recover the textual description. For each video, we take snapshots every 3 seconds for the first 15 seconds and perform the same prompt on each image. In aggregate, between images and video screenshots we have over 2 million images to be described. Before prompting ChatGPT, we downscale images to have a maximum size of 300 x 450 pixels.[11] In Appendix A.3, Table A.1 we provide examples of ChatGPT 4o supplied descriptions of images in our dataset.

We next convert each image description that we recovered to a dense vector embedding. We use *paraphrase-mpnet-base-v2* from the Sentence Transformers library, a pre-trained BERT-based algorithm with 109M parameters to contextualize information in an entire paragraph as opposed to single sentences at a time (Reimers, 2019). We take the pre-trained embeddings as a starting point, but add a fully connected layer to SBERT that predicts whether a post is a co-authored post or not. We unfreeze the top 4 layers of SBERT to allow the embedding model to re-weight parameters—a supervised learning process known as fine-tuning. This process will recover vectors that are most relevant to our goal of predicting treatment status. The fine-tuned SBERT model yields a 768-dimensional vector embedding that captures the semantic meaning of the description for each image.

2.3.2 Creating topic clusters from content embeddings

In theory one could match control posts to treatment posts using the 768-dimensional vector of SBERT embeddings (aggregated across images for posts with multiple images or video content). We pursue dimension reduction and topic clustering to improve robustness to overfitting on noise, increase interpretability of the matching variables, and conserve computational resources.

For each image, we reduce the SBERT embedding to 200 dimensions using Principal Components Analysis. We then use the Uniform Manifold Approximation and Projection (UMAP) algorithm to reduce each PCA vector to 5-dimensions, which preserves both local and global features of the embedded vectors. Finally, we use a simple k-means clustering algorithm to group vectors, and use the elbow method to select the optimal number of clusters, which is 50. Figure 3 details the complete dimension reduction and clustering pipeline for images in the dataset.

To understand the topic clusters and assess the performance of our model, we use Term Frequency-Inverse Document Frequency (TD-IDF) analysis to identify words that are especially

[11] We choose 300 x 450 pixels since there is a tradeoff between the cost of the API and the accuracy of ChatGPT 4o when faced with lower-resolution images; we find that ChatGPT 4o is able to describe 300 x 450 pixel images with similar levels of accuracy as images in their original resolutions.

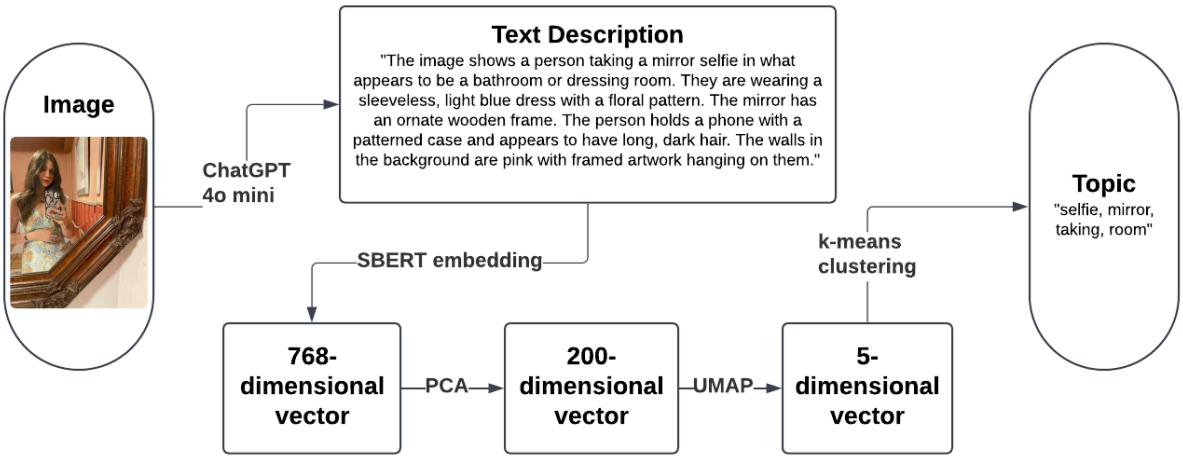


Figure 3: Topic creation flow

Note: ChatGPT 4o is accessed using the OpenAI API, using the prompt "Describe this image." SBERT is the sentence transformer model paraphrase-mpnet-base-v2, PCA is principal component analysis, UMAP is the uniform manifold approximation and projection, and k-means clustering is performed with $k = 50$. Topic text is included for descriptive purposes; the topic number itself (from 1-50) is included as an image's contribution to post-level content features.

prevalent within the descriptions in each topic cluster relative to the corpus of descriptions as a whole. Table 1 shows a sample of clusters, with the TF-IDF terms associated with each cluster, the number of images per cluster, and a few representative images from each cluster.

The comparison between the second and third rows showcases a benefit of our approach. While all six images feature individuals working out in similar attire, the LLM identified that images in the second row feature individuals focused on running on an outdoor track, while images in the third row feature individuals lifting weights or doing pliometric exercises in a gym. The SBERT embedding and dimension reduction preserves the content of the LLM's descriptive texts, and the TF-IDF analysis yields human-interpretable terms to understand how the topics are grouped.

2.3.3 Selecting matched control creators and posts

Our objective is to select a set of matched solo posts by control creators, that are visually similar to the treated posts by treated creators. We proceed by converting our image-level topics into post-level content attributes, adding several hand-selected features of posts and creators, and then detailing our matching function.

To convert our image-level topics to post-level attributes, we take the average of image topics

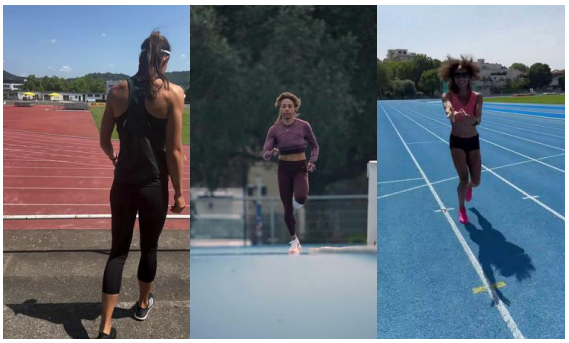
Topic	Topic Size	Image 1	Image 2	Image 3
salon, hair, stylist, cape	67154			
running, track, athletic, race	64153			
gym, exercise, performing, equipment	48624			
selfie, mirror, taking, room	23240			

Table 1: A sample of topics resulting from embedding and clustering

Note: Each row is a distinct cluster in the data. We present 4 clusters (out of 50 total) of varying sizes to provide face validity of the algorithm's ability to pick out meaningful topics from images. "Topic" includes the terms recovered using TF-IDF on the image descriptions, and can be thought of as providing a human-interpretable sense of what is captured by the cluster. "Topic Size" indicates how many images belong to that topic cluster. Image 1, 2, and 3 are example images drawn from that topic cluster, for each cluster.

over all images within a post to arrive at a post-level topic description. For example, for a post that includes both a video and an image, we sum the topic vectors for the image and each video snapshot, and then divide by the total number of images plus the number of snapshots. This 50-dimensional vector captures the topic weights within each post and sums to 1. As part of the attributes for matching treatment posts to control posts, we also include several post-level features that are likely relevant to a post's performance, including the precise timing of a post (year, month-of-year, day-of-week, and hour-of-day), and the type of post (Video, Image, or Sidecar, the latter of which denotes a post that contains multiple images, videos, or both). For each post, we also include the lagged value of the total number of creator's followers 7 days prior to the co-authored post.

As mentioned earlier, the pool of control posts includes the set of all non-collaborative posts created by the control creators. We also include as potential controls non-collaborative posts created by the treated creators, keeping only those posts that occur at least 90 days prior to their collaborative post. There are approximately 430,000 candidate control posts in total, for which we observe all the attributes detailed above.

To select our set of control posts from the pool of potential controls, we use Propensity Score Matching (PSM). PSM models the propensity to be treated as a function of observable characteristics, and for each treated post, selects a control post with the most similar propensity. We use the attributes described above and estimate the propensity score using a flexible gradient-boosted machine learning model. Note that PSM is essentially a type of dimension reduction, where a post's only relevant attribute becomes its propensity to be treated. We use greedy pairwise matching to select control posts sequentially. After each matched control post is selected, we remove from the pool all candidate control posts by the same creator 90 days before and 90 days after. In our empirical strategy we will track follower growth from 3 weeks (21 days) before each post to 6 weeks (42 days) after, and this trimming step ensures no overlap in observations. We should highlight that when a control post is selected, its creator is also selected as a control creator.

Appendix A.4 reports the standardized mean differences between covariates, and shows that the PSM model achieves a substantial balance improvement relative to the complete set of control posts. In Section 3.4.2, we consider alternative PSM-matched samples that (1) only include control posts from the creators who never co-author with brands and (2) only include control posts from the treated creators pre-treatment.

2.4 Controlling for cross-posting

An important confound for identifying the effect of co-authorship on Instagram is the tendency of brands to cross-post content to multiple social media platforms. It is possible that brands may simultaneously post about a creator on Instagram, Facebook, and Twitter/X. While other marketing papers that study the effect of creator content address the cross-posting problem using instrumental variables (Tian et al., 2024), we instead control directly for brand cross-posting after gathering cross-platform data.

For each creator j and for each brand in the set of treated creators, we find their account names on TikTok, Facebook, and Twitter/X. We start by collecting the brand names on each platform, to ensure a positive match for brands. Next, for each brand on each non-Instagram social media platform, we collect their entire posting history. We then filter to look at brand posts that occur within two weeks of the Instagram co-authored post on each non-Instagram platform. On TikTok and Facebook, we are able to gather all other accounts mentioned in each post using post metadata, while on Twitter/X we gather accounts by matching substrings within each Tweet that begin with "@". We perform fuzzy matching between these accounts and the co-authoring creator's name on Instagram and manually validate matches. With the account information in hand, we then code whether or not a brand post on Instagram mentions their co-authoring creator on another non-Instagram platform.

We find that the propensity to engage in cross-posting varies substantially across platforms, but in general is infrequent. We document 88 instances of a brand creating a post that mentions the co-authoring creator on TikTok, 44 on Facebook, and only 10 on Twitter/X. Of these, Facebook exhibits the highest proportion of same day posts—37 out of 44—which likely reflects the ease of cross-posting between two platforms owned by Meta. Seven of 10 Twitter/X posts are also simultaneous, while only 37 TikTok posts are simultaneous.

In Section 3.1 when we present the model, we discuss how we operationalize our controls for cross-posting by including dummy variables and lagged dummies for whether a brand posts about their co-authoring creator on another social media platform. We will estimate models both with and without these controls to assess how important cross-posting is for lifting creator outcomes on Instagram, and we will consider the sample of co-authored posts without any cross-posts in robustness checks.

2.5 Summary statistics and model-free evidence

Summary statistics. Our final sample consists of 648 co-authored posts, 648 control posts, and the follower growth rates for the creators making the co-authored posts and the creators making the control posts. In Table 2, we document the average rates of follower growth in our data around the time of the co-authored post.

Table 2: Outcome variable summary stats

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Follower Growth	49,142	0.0008	0.0034	−0.0545	−0.0001	0.0006	0.2270

Note: follower growth is daily growth rate. Sponsored post is the number of sponsored posts a creator posts on given day, excluding sponsored posts about the brand with which their co-author a post. These summary stats are computed using a window of 3 weeks prior to 6 weeks after the co-authored posts (or matched control posts).

Model-free evidence. To get a visual sense of the patterns of follower growth in the data, we provide raw plots of the average value by treatment and control before and after the co-authored posts. In Figure 4, we plot the normalized average follower growth rate for treated and control creators. We plot these variables for 3 weeks before and 6 weeks after the co-authored post (or the matched control post for the control creators). We choose these time windows conservatively based on the past empirical literature on influencer marketing: Cheng and Zhang (2024) find that follower counts change within a few days after a sponsored post on YouTube, while Huang and Morozov (2023) find that Twitch streams lead to changes in product engagement for streamed video games that decay within 24 hours.

The raw data suggests that prior to treatment, follower growth for treated and control creators behaves similarly, with a sharp divergence on the day of posting, peaking at a 5.5x higher growth rate for treated creators on the day after posting. This effect decays quickly, with treated creators largely returning to the same trend as control creators by 4 days after posting. Visually, for control creators there does seem to be a slight uptick in the daily follower growth rate at the time of the control post, which decays by 20 days after the post. Thus, both types of post seem to lead to growth, but with a large initial spike for co-authored posts before returning to the baseline growth trend. This trend is consistent with the fact that a typical post only stays on top of an active account for only a short time as new posts come out, and that Instagram keeps feeding new posts to users.

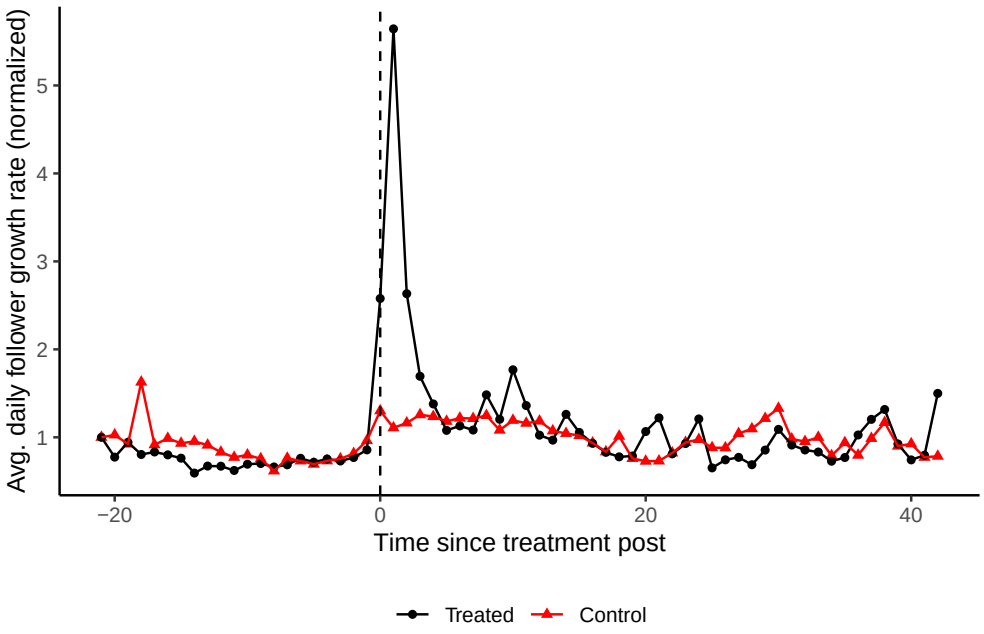


Figure 4: Normalized average follower growth rates around time of treatment

Note: The black circles denote the average outcome for treated creators, divided by its value 3 weeks (21 days) prior to the co-authored post. The red triangles denote the average outcome for control creators, divided by its value 3 weeks prior to the matched control post. For treated creators the treatment post is the co-authored post. For control creators, the treatment post is the control post selected through the propensity score matching model in Section 2.3.

In Section 3, we formalize this analysis in an econometric model, and provide evidence for our identification assumptions.

3 Model and Results

3.1 Model

Our analysis aims to identify the Average Treatment Effect on the Treated (ATT) for the subset of creators who are selected to co-author posts with brands. We estimate the effect of co-authorship on follower growth rate, using a modification of the standard two-way fixed effects (TWFE) model:

$$\text{FollowerGrowth}_{jt} = \sum_{s=0}^{T} \beta \cdot f(s,\theta) \cdot \text{Treat}_{j,t-s} + \boldsymbol{\gamma}'\mathbf{X}_{jt} + \alpha_j + \alpha_t + \alpha_{r(t)} + u_{jt} \tag{1}$$

where α_j is a creator-level fixed effect, α_t is a date level fixed effect, $\alpha_{r(t)}$ is a time-to-post level fixed effect, and $\mathbf{X}_{jt}$ is a vector of controls that we will discuss below.

$\text{Treat}_{j,t-s}$ equals 1 if j is a treated creator who co-authored a post with a brand on date $t-s$. s ranges from 0 to an exogenously set number T to include T lags of the treatment dummy. Including lags of $\text{Treat}_{j,t}$ allows the effect of treatment to persist for up to T days post-treatment.

The function, $f(s,\theta)$, which is parameterized by θ, captures the decay of the effect of co-authorship outcomes. As seen in the model-free evidence in Figure 4, the treatment effect does not peak on the day of treatment and then uniformly drop off, but rather peaks a day after the post before subsiding over the next few days. We therefore use a Poisson probability density function with $f(s,\theta) = \frac{\theta^s e^{-\theta}}{s!}$, which is parsimonious and still flexible enough to capture the non-monotonic decay effect.

The parameter β is our main focus, and it captures the effect of treatment on the outcome. Note that since $\sum_{s=0}^{T} f(s,\theta) \approx 1$ for large enough T, β is the *cumulative* effect of treatment, and $f(s,\theta)$ captures the share of the total treatment effect that manifests on day s post-treatment.

The vector of controls $\boldsymbol{X}_{jt}$ includes dummy variables that will control for cross-posting by brand j on other social media platforms, per the discussion in Section 2.4. Specifically, we construct a vector of three dummy variables, each of which is equal to one for creator j at time t if the brand they co-author with on Instagram posts about them on TikTok, Facebook, or Twitter/X respectively. We then include lags of these dummy variables up to seven days prior, which yields a 24-dimensional vector of control dummy variables for each creator that captures contemporaneous and lagged posts about the creator by co-authoring brands on other platforms.

In the next two subsections, we discuss justifications for our identification assumptions and the estimation procedure. Readers who are more interested in the results can skip to Section 3.4.

3.2 Identification

Equation (1) is meant to identify the causal effect of a co-authored post on a creator's follower growth. Our identification assumption is that conditional on the creator fixed effects, date fixed effects, time-to-post fixed effects, and controls, the treatment variable Treat_{jt} is orthogonal to the unobserved error u_{jt} in Equation (1).

Having controlled for post content through matching, the key remaining concern is that the timing of treatment may coincide with unobserved, creator-time-specific shocks that also drive

interest in the creator. The shocks must be creator-specific to cause an identification problem.[12] Although this problem is theoretically possible, we find it unlikely to happen in our institutional setting. The nature of co-authored posts is that although all included Instagram accounts appear as co-authors, one account must be designated as the post owner, which has control over post timing. In brand-creator collaborations, the brand is typically the post owner for reputation risk reasons, as only the owner has the ability to remove or modify the co-authored post.[13] Indeed, in our sample of co-authored posts, the post owner is the brand in all cases. Practitioners have found that brands have complex internal approval processes for posting content that frequently lead to delays.[14] Furthermore, content calendars for their social media accounts are scheduled weeks in advance.[15] Firms thus have little ability or incentive to quickly adjust content timing in order to take advantage of transient, creator-specific shocks. Post timing should therefore be orthogonal to such shocks.

To check our assumption that brands do not coordinate posts with external demand shocks for the creator, we gather a proxy variable that captures transient creator-specific demand shocks and estimate whether it is correlated with the timing of co-authored posts. Our proxy for demand shocks is weekly global search volumes scraped from Google trends for the creator's actual name. We successfully link 83% of treated creators to their actual name, and then collect global search volume data by week for each name. We regress Treat_{jt} on search volume with creator fixed effects to test whether treatment timing is related to changes in search volume, and find an insignificant partial correlation of 0.0001 ($p = 0.77$).[16] Moreover, we also perform a visual timing test: we first identify the week with peak search volume for a creator that occurs within six weeks of the co-authored post for that creator (Appendix B.1 provides visual examples of the size and timing of local peaks for several creators). We then plot the propensity of the brand to create the co-authored post from six weeks before to six weeks after that period of peak interest. If the

[12] Another endogeneity concern is that brands may tend to post co-authored content on dates that typically lead to larger exposure. However, our calendar-date level fixed effect α_t should address this problem by absorbing general unobserved time-specific heterogeneity.

[13] From https://www.facebook.com/help/instagram/291200585956732: "If the original author deletes the post, it will be removed from the collaborator's profile as well.", Accessed January 17, 2025.

[14] See "6 steps for an efficient marketing approval process", *Ziflow*, March 27, 2023, as well as "How to Speed Up the Content Approval Process in a Large Organization", *Kuno Creative*, February 3, 2017.

[15] See "Marketing Briefing: How brands and agencies are speeding up their pursuit of influencer trends", *Digiday*, December 12, 2023, as well as "How Far Ahead Should You Plan Your Marketing Strategy?", *WFSB Advertising*, August 11, 2021.

[16] Since treatment is observed at the daily level and search volume is at the weekly level, we impute the daily volume as the weekly volume in this regression.

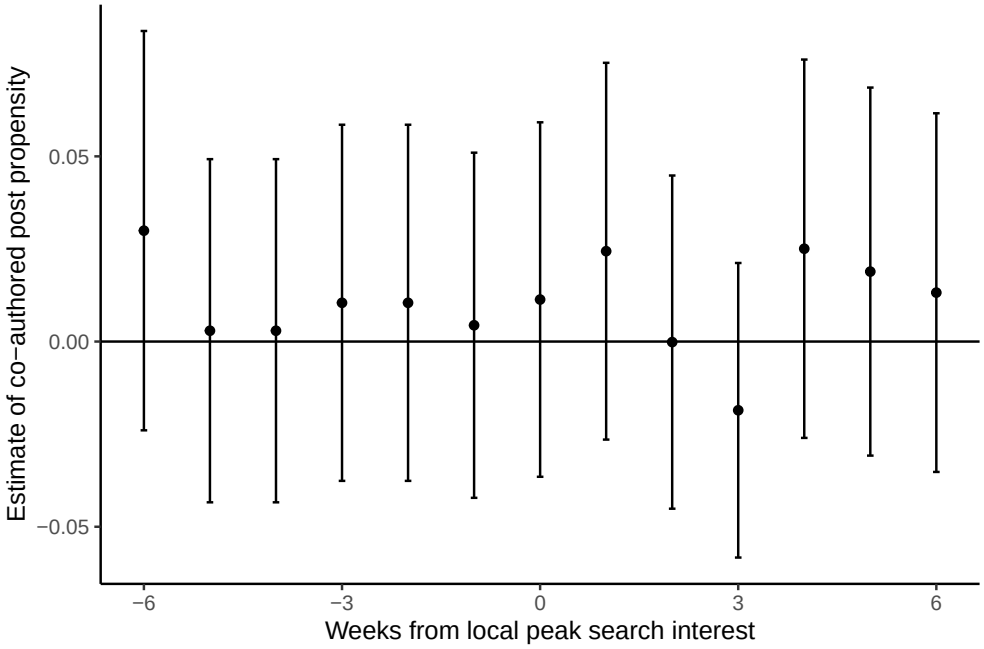


Figure 5: Co-authored posting around times of peak search interest in creator

Note: This figure presents coefficients from a regression at the creator-week level, where the dependent variable is whether the creator's co-authored post on Instagram was posted that week, and the independent variables are a set of leads and lags around the week of highest search volume for the creator's actual name. If brands time co-authored posts on Instagram with external demand shocks for the influencer, there should be a positive propensity to post in the week before, during, and after the external demand shock, which we do not find. Error bars are 95% confidence intervals clustered at the creator level. See Appendix B.1 for more details.

timing of co-authored posts is exogenous with respect to surges of search volume for the creator, then we should find that the propensity to post does not rise during the period of peak interest, or in any period around the period of peak interest. This analysis is presented in Figure 5, where we can see that there no greater propensity for the firm to post in the week of peak interest compared to any other week around the time of the co-authored post.[17] In Appendix B.1, we check the robustness of this exogeneity result to different timing windows.

One may be concerned that brands do not plan their content calendars far in advance, and instead select creators based on positive trends in followership growth. We conduct a formal event study analysis in Appendix B.2 to rule out that treated creators exhibit pre-trends prior to their co-authored posts, and also show that there is no differential linear pre-trend for treated creators in Appendix B.3.

[17]If the brand's co-authored post leads to Google search interest in the creator, then the timing of the post and Google search volumes may coincide, but not because the brand endogenously times co-authored posts to surges in interest for the creator. Since we find no relationship between post timing and search volume, both mechanisms are ruled out.

3.3 Estimation

Our assumption from Section 3.2 that treatment timing is exogenous conditional on our set of fixed effects and controls implies the following moment conditions:

$$\mathbb{E}[u_{jt}|\text{Treat}_{j,t-s}, \mathbf{X}_{jt}, \alpha_j, \alpha_t, \alpha_{r(t)}, \theta, \beta, \boldsymbol{\gamma}] = 0, \qquad s = 0, \ldots, T \tag{2}$$

We estimate Equation (1) using the generalized method of moments (GMM). The parameters to estimate are the non-linear Poisson parameter θ, the cumulative treatment effect β, the control parameters γ, and the fixed effects $\boldsymbol{\alpha}$. Since all parameters other than θ are linear, our estimation procedure first picks a value for θ, then estimates $(\beta, \boldsymbol{\gamma}, \boldsymbol{\alpha})$ using a fixed-effects regression. We recover $\hat{u}_{jt}(\theta, \beta, \boldsymbol{\gamma}, \alpha)$ and interact it with $\text{Treat}_{j,t-s}$ with $s = 0, \ldots, 21$, i.e. the treatment variable and twenty-one lags. The lags help match the dynamic shape of the response of the outcome variable to treatment over time. This yields a 22x1 empirical moment vector:

$$\hat{g}(\theta) = \frac{1}{N} \sum_j \frac{1}{T_j} \sum_t \hat{u}_{jt}(\theta, \beta(\theta), \boldsymbol{\gamma}(\theta), \boldsymbol{\alpha}(\theta)) \cdot [\text{Treat}_{jt}, \ldots, \text{Treat}_{j,t-21}]'$$

We use this vector to form the empirical GMM objective $\hat{g}(\theta)' \cdot \mathbf{W} \cdot \hat{g}(\theta)$, assuming the identity weighting matrix for simplicity.[18] We search for the optimal θ using golden section search with parabolic interpolation to find the minimum of the one-dimensional objective function.

Our estimation approach echoes prior work in the influencer marketing literature on estimating non-linear parameters in a panel data setting with fixed effects, including Li et al. (2023) and Huang and Morozov (2023) who estimate a geometric decay of influencer content on purchases. Consistency, unbiasedness, and asymptotic normality of the GMM estimator are guaranteed as long as Equation (2) holds, and the standard assumptions on full rank of the Jacobian of the expectation of the moment conditions are satisfied (Hayashi, 2011).

3.4 Results

Results from estimating Equation (1), the effect of co-authorship on the follower growth rate for creators, are presented in Column 1 of Table 3 without controls, and are presented in Column 2 controlling for brand cross-posting. The β parameter is positive and highly statistically sig-

[18] Using the identity weighting matrix will yields consistent estimates of θ, but may have larger standard errors than the asymptotically efficient continuously updated GMM.

nificant in both specifications, as expected given the patterns in the left panel of Figure 4. The estimate of β in Column 2 implies that the co-authored post leads to a *cumulative* 0.47% bump in the follower growth rate. The β coefficient in Column 2 shrinks slightly compared to Column 1 after including the controls. We find that this reduction only occurs after including the Facebook posting controls, which may be because cross-posting between Facebook and Instagram is relatively easy as both are owned by Meta.

For the Poisson distribution parameter θ, the estimate in Column 2 suggests that 26% of the cumulative follower growth is experienced on day 0, 62% by day 1, and 85% by day 2. This pattern accords closely with the visual evidence in Figure 4, where the majority of follower growth is earned the day after the post.

In Columns 3 and 4, we replicate our analysis and estimation using the change in followers as the dependent variable instead of the growth rate in followers.[19] After including controls, we estimate a cumulative increase in followers of 114.6 on average, with a Poisson distribution parameter whose magnitude implies most growth is earned on day 1.

The average effect of the collaboration is notable for just how small the magnitude is. The average creator in our data has 59,000 followers prior to collaborating, and the average brand has 7.6 million followers, an audience over 120 times larger than the creator's. From Column 2, creators earn a 0.47% cumulative bump implying only 277 new followers, while estimating the direct gain in Column 4 implies only 114.6 new followers.[20] Thus, for a creator to co-author with a big brand on social media and be exposed to a new audience of millions of potential followers does not lead to exceptional growth. Indeed, given that the average daily growth rate is 0.08%, the gain from exposure represents earning approximately an extra week's worth of growth with a single post. Brands' beliefs that there are large benefits to creators from exposure are thus not realized on average. This small gain could be because the duration of exposure is very short. However, there may be instances in which creators do enjoy substantial benefit, and in Section 4 we explore heterogeneity in the treatment effect.

3.4.1 Net effect versus pure exposure effect

One concern with our estimates may be that we are estimating a net effect, rather than a pure exposure effect. In particular, since we only observe each creator's total number of followers, the

[19] We replicate the event study plot using the growth in followers (not growth rate) in Appendix B.4.

[20] In general, $E[Y_t - Y_{t-1}|X_t] \neq E[(Y_t - Y_{t-1})/Y_{t-1}|X_t] \cdot Y_{t-1}$. The discrepancy in estimates is driven by the skewness of the follower growth variable.

Table 3: Average baseline effects of co-authored post

	Dependent variable:			
	Follower Growth Rate		Follower Growth	
	(1)	(2)	(3)	(4)
β	0.0053***	0.0047***	129.9***	114.6**
	(0.0013)	(0.0012)	(44.7)	(44.6)
θ	1.3562***	1.3291***	1.3060	1.2304
	(0.1224)	(0.1224)	(0.9021)	(0.9020)
IG Account FE	✓	✓	✓	✓
Date FE	✓	✓	✓	✓
Time-To-Post FE	✓	✓	✓	✓
Cross-Posting Controls		✓		✓
Avg. Value of Dep. Var.	0.0008	0.0008	42.2	42.2
Observations	47,524	47,524	47,521	47,521
R^2	0.1907	0.1932	0.2760	0.2770

Note: standard errors are clustered at the IG account level. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

estimated increase in new followers in Table 3 combines the increase in followers from the brand's audience, and a potential loss in followers from the creator's audience. If a creator's followers perceive the co-authored post as sponsored content, they may decide to disengage (Bairathi and Lambrecht, 2023), or even unfollow. While the relative sizes of the brand's and creator's audiences suggests that the exposure effect should dominate, as a robustness check we estimate the effect of a sponsored post on creator followers as in Cheng and Zhang (2024) to see whether a large drop in existing followers may be offsetting a large increase in new followers (see Appendix C.1 for details). Using almost 20,000 posts identified as sponsored in our data, we find a loss in followers of between 0.016% and 0.037% by three days after the post.[21] In this scenario, the growth in followers from pure exposure to the brand's audience is therefore slightly higher, between 0.49% and 0.51%.

[21] Cheng et al. (2023) find an average loss of 0.19% of followers by three days after a sponsored post, but estimate that smaller creators have a much smaller percentage loss, or even a percentage gain for small enough creators. Our average creator has 59,000 followers, while the average creator in Cheng and Zhang (2024) has over 1 million followers.

3.4.2 Robustness checks

Our approach to handling unobserved exposure generated by cross-posting to other platforms is to include a vector of cross-posting controls. As an alternative to controlling explicitly for cross-posting, we also try dropping all brand-post-creator treatment events where we detect cross-posting. Estimates in Appendix C.2 show that dropping these events has little effect on the size or significance of the treatment effect.

In Section 2.3, we draw on control posts from two sources: treated creators pre-treatment, and never-treated creators. We replicate Table 3 using two alternative sets of control posts. We try (1) dropping posts from treated creators pre-treatment from the control pool, and (2) dropping posts from never-treated creators from the control pool. In each case, we re-estimate the PSM function and reselect matching control posts from the more limited pools. Results reported in Appendix C.2 suggest that the estimated treated effect is unchanged.

Some creators are active across multiple social media platforms, and may attempt to publicize their co-authorship to followers on these other platforms, potentially leading to greater follower growth on Instagram. We estimate event studies for treatment effects after dropping creators who have large followings on other social media platforms in Appendix C.3, and find that omitting these creators has no impact on treatment timing or magnitude.

Finally, our sample includes creators who post immediately after the co-authored post. Since the largest follower growth does not manifest until a day after the co-authored post, one may worry that fast follow-on posts are driving the treatment effect. We therefore estimate Equation (1) using only those creators who do not create any new posts immediately after the co-authored post. We find no significant difference in the cumulative increase in the number of followers in Appendix C.2, and a slight increase in the Poisson distribution parameter compared to the baseline estimates.

4 When Do Creators Benefit More?

In this section, we seek to answer two questions. First, we wish to understand when creators benefit more from the exposure effect. The main dimension we analyze is whether it benefits creators more to work with larger brands. We are motivated by nuances in the debate between creators and brands, where creators tend to express beliefs that only working with large brands can be valuable (e.g., "anybody big enough to actually give you exposure is big enough to pay you

for your work"[22]), while social media managers suggest that even small brands can successfully incentivze creators through exposure (Somerfield, 2023).[23] Intuitively, supplying content to large brands should lead to more benefit for the creator as large brands have larger audiences, and therefore greater opportunities for potential followers to discover the creator. However, large brands also tend to have less engaged audiences, and this may constrain creator's prospect of gaining new followers. In addition to analyzing brand size, we also wish to understand how the exposure benefit varies with (1) post-level attributes, such as how different the co-authored content is from the creator's usual content, and (2) creator-level attributes like creator quality.

Second, we wish to estimate the complete empirical distribution of follower lift, to understand if and when it is possible for creators to experience higher growth in followers after co-authoring with a brand. To pursue both objectives, we will modify Equation (1) to allow for heterogeneity in the treatment effect with respect to brand size, as well as other attributes such as brand engagement.

4.1 Model specification and heterogeneity variables

Our approach for testing whether brand size moderates the treatment effect, and for recovering the complete empirical distribution of treatment effect sizes, will be to estimate the following model that allows for treatment effect heterogeneity through $\boldsymbol{Z}_j$:

$$\text{FollowerGrowth}_{jt} = \sum_{s=0}^{T} f(s,\theta) \cdot \boldsymbol{\beta}'\mathbf{Z_j} \cdot \text{Treat}_{j,t-s} + \boldsymbol{\gamma}'\mathbf{X_{jt}} + \alpha_j + \alpha_t + \alpha_{r(t)} + u_{jt} \tag{3}$$

where $\mathbf{Z}_j$ is a vector of variables that capture brand, message, and creator attributes associated with the co-authored post of creator j.

The key variable in $\mathbf{Z}_j$ will be a measure that captures brand size. We operationalize this variable by computing the log number of followers of the brand at time of data collection (August 2024), and normalizing it by subtracting the mean and dividing by the standard deviation. We expect that conditional on other brand, message, and creator level attributes, that larger brands should lead to more benefits for the creator. We discuss which additional attributes we control for below.

[22] See "Artwork for Exposure - when is it a good deal?", on https://www.reddit.com/r/ArtistLounge, April 8 2019.

[23] See "What are some creative ways to incentivize and reward user-generated content creators?", *LinkedIn*, 2024. Note that this belief may reflect that exposure from small brands can lead to meaningful exposure gains, or that small brands yield small exposure benefits but that even minimal exposure is still motivating to creators.

For additional brand-level attributes, we include in $\mathbf{Z}_j$ the engagement of the brand's audience. For almost all accounts on Instagram, engagement tends to fall as follower counts rise. Thus, the effect of larger brand size on creator's follower lift would be conflated with the effect of lower brand engagement unless it is held fixed. Engagement is likely to matter for follower lift as it captures audience attention; indeed, the Instagram algorithm prioritizes showing content from accounts that users consistently engage with.[24] We operationalize this variable as the average likes per post divided by brand followers in the 90 days prior to the co-authored post, and once again log and normalize this variable.

For message-level attributes, we include in $\mathbf{Z}_j$ three dimensions that may be important for lift, and that may vary across brands of different sizes: how surprising or novel the content of the co-authored post is relative to the brand's usual content, how closely the content of the co-authored post hews to the creator's typical work, and how engaging the content is. Novel or surprising content may grab audiences' attention (Kanuri et al., 2024) and be perceived as creative (Villanueva et al., 2008; Somerfield et al., 2018), or conversely may not be to the audience's tastes; it may also be more likely to originate at larger brands if they have a more sophisticated content team that works with the creator. For creators, a co-authored post that more closely resembles their typical content may be perceived as more authentic (Cheng and Zhang, 2024; Bairathi and Lambrecht, 2023), or may provide more accurate information about the creator, easing brand audience members' conversion journey in becoming new followers; small brands may also allow creators the freedom to craft the content in a way that resembles their usual output. Finally, engaging content is likely to lead to more follower growth, and acts as a proxy for the quality of the post.

To operationalize the message-level variables, we compute the cosine distance between the SBERT embeddings for the co-authored post and (1) embeddings for the brand's posts in the 90 days prior to the co-authored post, and (2) embeddings for the creator's posts in the 90 days prior to co-authored post. These two variables capture the distance between the co-authored post and the brand's usual content, and the distance between the co-authored post and the creator's usual content, and we log and normalize across all co-authored posts. For message-level engagement, we compute a measure of within-brand relative likes for the co-authored content, by computing the logged and normalized likes for the co-authored post relative to the logged normalized likes

[24] See https://socialbee.com/blog/instagram-algorithm/, accessed December 30, 2024.

Table 4: Correlation matrix: heterogeneity variables

	Brand Size	Brand Eng	Post BrandDiff	Post CreatorDiff	Post RelLikes	Creator Size	Creator Eng
BrandSize	1						
BrandEng	-0.357	1					
PostBrandDiff	0.028	0.087	1				
PostCreatorDiff	0.016	0.071	0.387	1			
PostRelLikes	-0.245	-0.151	0.005	-0.064	1		
CreatorSize	0.181	-0.022	0.052	-0.025	0.130	1	
CreatorEng	0.069	0.184	-0.024	-0.019	0.121	-0.117	1

Note: BrandSize is the log normalized number of brand followers, BrandEng is the log normalized ratio of likes to followers for brand posts in the 90 days prior to the co-authored post, PostBrandDiff is the normalized difference between the co-authored post content and brand post content in the 90 days prior to the co-authored post, PostCreatorDiff is the normalized difference between the co-authored post content and creator post content in the 90 days prior to the co-authored post, PostRelLikes is the engagement of the co-authored post normalized relative to the brand's non-co-authored posts in the 90 days prior to the co-authored post, Creator Size is the log normalized number of creator followers, and CreatorEng is the log normalized ratio of likes to followers for creator posts in the 90 days prior to the co-authored post. Normalization is Z-transformation across co-authored posts.

for the same brand in the previous 90 days.[25]

Finally, for creator-level attributes, we include the creator's size and engagement level. Creators with larger followers size and more engaged followers may be perceived as higher quality by the brand's audience, and may be more likely to be picked by larger brands. To capture creator size, we use the number of followers 7 days prior to the co-authored post. For creator engagement, we use the likes per post divided by number of followers, for non-co-authored posts in the 90 days prior to the co-authored post. Both variables are logged and normalized across co-authored posts as usual before inclusion in $\mathbf{Z}_j$.

In Table 4, we show the correlation matrix for our heterogeneity variables. Looking at the first column, we find that brands with larger follower counts tend to have substantially lower average engagement rates, with a correlation of -0.36. For message-level attributes, co-authored posts with larger brands tend to have a greater differences with the brand's usual content, and with the creator's usual content, although both correlations are small in magnitude. However, co-authored posts at large brands are substantially less liked than their typical posts, with a correlation of -0.24. Finally, we see evidence that larger brands tend to select higher "quality" creators, with positive correlations between brand size and creator size (0.181), and brand size and creator engagement (0.069).

[25]Note that if co-authored posts tend to receive very high engagement, then this variable will not be mean zero across brands, since the normalization happens within-brand, not across brands. In practice, we find that this variable is still centered approximately at zero.

Table 5: Treatment effect heterogeneity of the co-authored post

	Dependent variable:	
	Follower Growth Rate	Follower Growth
	(1)	(2)
Treat	0.0046***	99.4207*
	(0.0012)	(54.0879)
Treat×BrandSize	0.0051**	166.9360***
	(0.0020)	(56.1901)
Treat×BrandEng	0.0060***	164.2232***
	(0.0019)	(51.5226)
Treat×PostBrandDiff	0.0006	14.8038
	(0.0014)	(54.5093)
Treat×PostCreatorDiff	-0.0031*	-103.8308**
	(0.0016)	(48.1630)
Treat×PostRelLikes	0.0043***	153.8658***
	(0.0011)	(48.1018)
Treat×CreatorSize	-0.0070***	-64.6700*
	(0.0019)	(33.4160)
Treat×CreatorEng	-0.0019	-51.3994
	(0.0012)	(41.0093)
θ	1.4592***	1.6559***
	(0.3086)	(0.3070)
IG Account FE	✓	✓
Time-To-Post FE	✓	✓
Date FE	✓	✓
Controls	✓	✓
Avg. Value of Dep. Var.	0.0008	44.253
Observations	49,097	49,097
R^2	0.2060	0.2534

*Note: standard errors are clustered at the IG account level. *p<0.1; **p<0.05; ***p<0.01*

4.2 Heterogeneity results

In Table 5, we present the results from estimating Equation (3). The specification in Column 1 uses the growth rate in daily followers as the dependent variable. We also estimate Equation (3) using the change in daily followers as the dependent variable.

The average estimated baseline lift in row 1 of Column 1 is similar to that in row 1, Column 2 of Table 3. Since almost all interaction variables are mean zero by construction and we include controls in Table 5, this concordance is expected. The Poisson distribution parameter that governs decay in Column 1 is not statistically significantly different from its estimate in the baseline

regression.

We find that holding fixed all other attributes, a one standard deviation increase in the log number of brand followers leads to a 0.51% greater increase in cumulative follower growth from co-authorship. As expected, larger audiences translate into a larger benefit for the creator from working with the brand all else equal. In Figure 6, we plot the estimated treatment effect by brand size holding all other attributes fixed, and show that it is positive for the majority of brands. The estimated treatment effects for each brand size fall exactly on a straight line, since we only allow for the baseline treatment effect (intercept) and brand size coefficient (slope).

The estimated coefficients on other attributes largely accord with intuition. Similarly to brand audience size, brand audience engagement also positively moderates treatment effects, which reflects the benefits to the creator of an attentive audience. Novel content (PostBrandDiff) has no significant effect, possibly owing to the fact that although novel content may grab attention, it may not be to the audience's tastes. Meanwhile, a co-authored post that is very different from the creator's usual content (PostCreatorDiff) leads to lower lift, potentially since it is uninformative to potential followers and may even signal inauthenticity (Cheng and Zhang, 2024). The last post-level attribute, engagement, is associated with higher lift. Intuitively, high engagement captures residual attributes of the collaborative post that the audience values, such as quality or trendiness, which may reflect positively on the creator.

Turning to creator attributes, we find that creator size is strongly negative; however, this result is mechanical, as the absolute gain in followers is roughly independent of creator size, leading the growth rate to be smaller for creators that are ex ante larger. Interestingly, more engaging creators do not benefit more. Taking engagement as a proxy for creator quality, this result implies that higher quality creators do not enjoy a greater benefit from co-authorship. While this result may seem counterintuitive, it could be that the brand's choice to work with the creator is itself a very strong quality signal. In that case, there may be little room for brand followers' beliefs to be further shifted.

We next plot the estimated treatment effect by brand size, allowing other brand and post level attributes to vary as they do empirically in the data. This exercise captures what treatment effect a creator can actually expect from working with brands of different sizes. Note from Table 5 that brand audience engagement and the engagement with the co-authored post by the brand's audience both strongly positively moderate the effect of treatment. However, these variables are negatively related to brand size, and so one would expect the positive effect of brand size to also

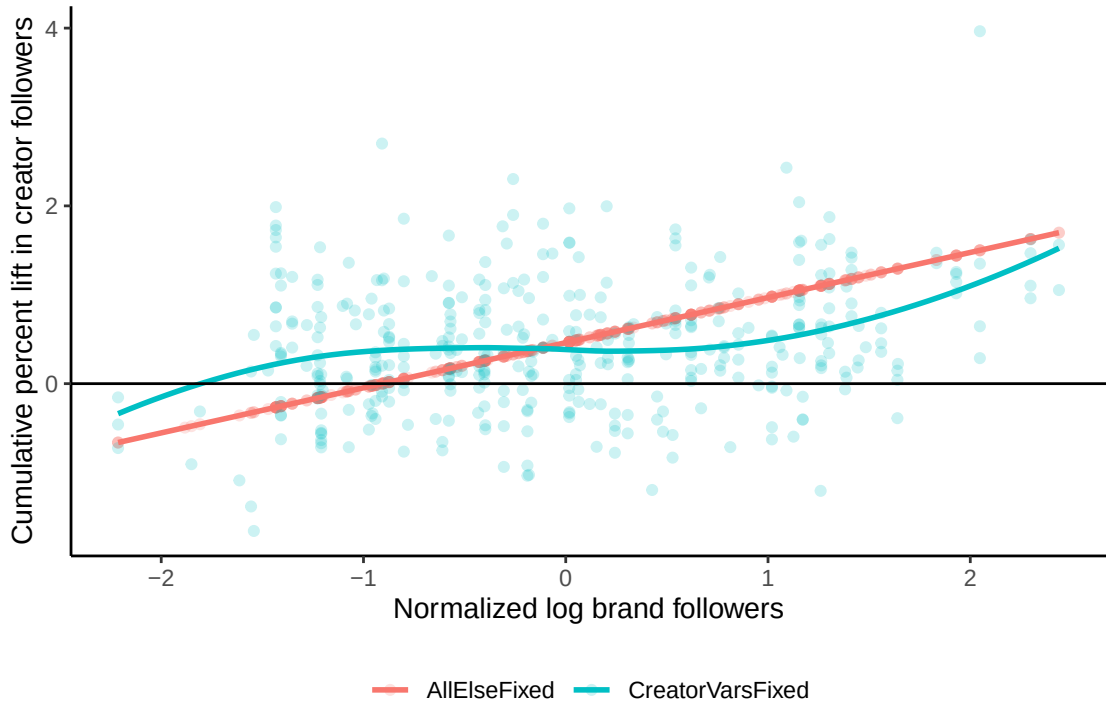


Figure 6: Treatment effects by brand size

Note: The straight line plots the marginal treatment effect across different brand sizes from Table 5, holding all other non-BrandSize attributes fixed. The points plot the treatment effect for each co-authorship allowing brand and post attributes to vary but holding only CreatorSize and CreatorEng fixed. The smooth curve is a local linear approximation to these estimated treatment effects.

moderate once these variables are allowed to co-move with size as they do in the data. In Figure 6 we plot the estimated treatment effect for each co-authored post in the data holding only creator attributes fixed, and fit a smoothed local linear curve to these points. We find that for a large part of the distribution of firm sizes (from roughly -1 standard deviation below to +1 standard deviation above the mean log followers) the average relationship between brand followers and the exposure benefit to the creator is essentially flat. Although the very largest (smallest) brands still drive higher (lower) follower growth, working with brands within the inter-quartile range of sizes, from 710,000 to 8.6 million followers, leads to approximately the same benefit from exposure on average.

4.3 Distribution of exposure bucks

In this section, we provide the full empirical distribution of treatment effects, and a back-of-the-envelope calculation to compute the distribution of dollar values to creators from exposure. This exercise is meant to provide information to creators, influencer agencies, and brands when

negotiating and evaluating compensation packages.

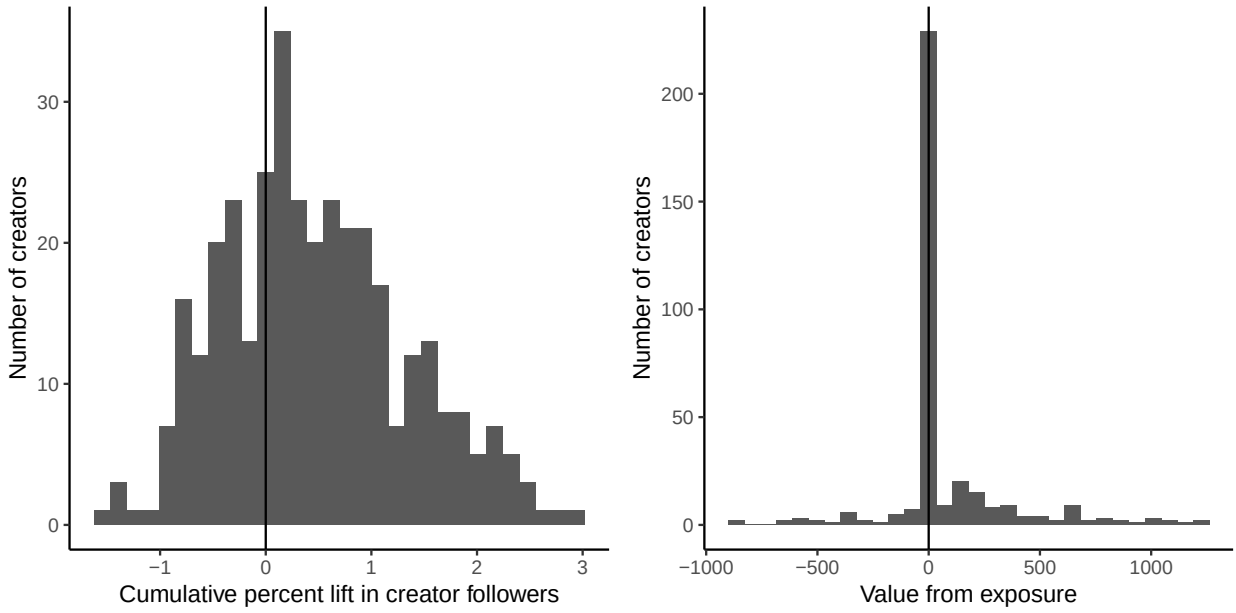
To convert follower lift into a dollar value, we assume that creators monetize their followers' attention by endorsing brands' products to their followers in exchange for payment. This endorsement usually takes the form of a "sponsored post", where the creator positively reviews, or is seen to use, a brand's products. This assumption on how creators monetize accords with theory models of influencer marketing (Fainmesser and Galeotti, 2021; Mitchell, 2021). We caveat however that this assumption may miss other ways that creators monetize their accounts, such as affiliate links, product co-creation, or using their account to advertise their paid services.

This assumption allows us to compute the value of a creator's account as the product of the price-per-sponsored-post multiplied by the frequency of sponsored posting, discounted over all future sponsored posts. The change in the value of a creator's account from a co-authored post depends on two things: the change in the price-per-sponsored post, which will increase as the creator's audience expands, and the increase in sponsored posting. We formalize this value in Appendix C.4.

For the increase in the price-per-sponsored-post, we estimate how sponsored post prices depend on follower counts using data from HypeAuditor, a leading influencer marketing agency (see Appendix C.5). For the increase in sponsored posting, we must first detect sponsored posts. We use the dictionary-based approach applied to the post captions that we detailed in Section 3.4.1 Using the detected sponsored posts, we estimate Equations (1) and (3) using sponsored posts as a dependent variable (Appendix C.6). We do not find any positive impact of co-authorship on sponsored posting within 6 weeks of the co-authored post. At least in the short run, the visibility from co-authoring with a brand is not correlated with new sponsorship business for the creator, and thus we can focus on the effect of increasing price through larger follower counts as the sole channel through which co-authorship affects the value of the creator's account.

Using our heterogeneity estimates in follower lift from Table 5, we first plot the distribution of the increase in the number of followers in Figure 7a. Combining this lift heterogeneity with the above assumptions on how followers are monetized, we quantify the dollar value of exposure across the distribution of brands, creators, and content observed in our data in Figure 7b.

The distribution of the values from co-authorship has a large spike at zero (58% of creators), and is otherwise predominantly above zero (70% of non-zero values are positive). A zero effect for creator value could derive from a null treatment effect on follower growth or a complete lack of sponsored posts on the creator's account, which are what allow the creator to monetarily profit



(a) Follower increases

(b) Dollar value of exposure

Figure 7: Distribution of treatment effects of co-authorship

Note: The data in the left figure comes from using Column 1 of Table 5 to predict the lift in followers for each treated creator in the dataset, and multiplying this number by 100. The data in the right figure comes from plugging a creator's lift in follower numbers into the relationship between sponsored post prices and followers (see Appendix C.5 and **??***), and multiplying that increase in price by creator-specific propensities to post sponsored content. We find that there is no marginal increase in sponsored posting from co-authorship (see Appendix C.6).*

from their followers. The distribution of follower growth rates is smooth around zero, indicating that the second mechanism explains the large number of creators who do not benefit monetarily from co-authorship.

Conditional on having a non-zero value from exposure, the median treated creator experiences an increase in the lifetime value of their account of only $149. However, at the 80th percentile the value of exposure is $526. The creators who are in the top of the value distribution tend to be smaller, work with brands that are larger and with better engagement rates, and have a better match between the content of the co-authored post and their own content. It is also the case that these creators tend to have a greater rate of sponsored posting: the top 20% of creators with non-zero values are almost twice as likely to create a sponsored post on a given day (0.09) compared to the bottom 80% (0.045).

5 Discussion

Our paper has several managerial implications. First, we attempt to resolve an ongoing debate about the value of exposure to creators, and whether working with smaller brands can still be valuable. Amongst creators, there are many who deride being compensated with exposure as "exposure bucks", and would prefer to be paid directly.[26] Meanwhile, social media marketers and social media managers at brands often think that the exposure brands can provide is very valuable (Somerfield, 2023).[27] Our paper can reconcile these two sets of beliefs as potentially arising from cherry-picking from different parts of the distribution of treatment effects. It is possible that, as with influencer marketing (Huang and Morozov, 2023), brands may only consider examples from the positive part of the distribution of treatment effects, leading them to have a false sense of how valuable exposure is. On average, for the set of creators and content we consider in this paper, exposure has a positive but small benefit for creators. We also provide nuance to the discussion on whether brand size matters, finding that larger audiences are better all else equal, but that larger brands systematically vary in other ways that actually lead them to provide a similar benefit from exposure for a broad range of sizes.

Second, we provide novel insights for influencers—or influencer agencies who negotiate con-

[26] See e.g. https://www.helloraderco.com/sacrificing-time-money-exposure-bucks/ and https://amandajspedding.com/2019/05/27/exposing-exposure-bucks-hint-its-not-a-real-thing/.

[27] See also "What are some creative ways to incentivize and reward user-generated content creators?", *LinkedIn*, 2024. Amongst the expert contributors to the article, most mention exposure as a way to reward creators, especially if the brand is on a budget.

tracts on behalf of influencers—on how they should value the offer of collaborations from brands. Influencers who hope to experience explosive growth in followers from co-authoring with a massive brand should moderate their expectations: even the top 10% of follower gains in our estimated distribution in Figure 7a only lead to a cumulative bump of 2.1% new follower growth. For the median co-authored post in our data, we calculate that exposure is not as valuable as influencers' expected sponsored payments. While brands may wish to offer exposure or credit in exchange for sponsored posts from influencers, only large brands with highly engaged customers tend to offer enough value for this to be a motivating proposition relative to direct monetary transfer.

Third, outside of influencer marketing, brands often solicit UGC to use in advertising or on their brand pages, and may wish to offer the promise of exposure as compensation. Campaigns like Apple's "Shot On iPhone" challenge are held up as success stories where authentic, cheap UGC was sourced partly through the promise of exposure. Based on our heterogeneity results, the offer of exposure is most motivating for small creators who wish to maintain consistency with their previous content. For brands who wish to ensure a steady supply of UGC for their social media accounts and advertising, they should avoid placing constraints that would force a creator to deviate from their usual content, since such deviations would mean far less value to the creator from the promised exposure. However, we note that for many small creators, prestige and recognition may also comprise part of the value of being credited, and more work is needed to understand the relative importance of these different extrinsic motivations.

Finally, we provide interesting new empirical patterns for the academic literature on influencer behaviour. Typically, influencers are conceptualized as creators on social media who build up a follower base through providing entertaining or informative content before then monetizing that base by engaging in sponsored posting (Fainmesser and Galeotti, 2021; Mitchell, 2021). Our analysis shows that it is possible to supply content to brands in a one-off collaboration and experience follower growth directly through exposure to the brand's audience. Whether growth from exposure

6 Conclusion

Soliciting content from creators in exchange for exposure is a common business tactic, but the value of exposure has been difficult to quantify. In this paper, we tackle this problem in the context of social media influencers on Instagram. We leverage a novel dataset that covers collaborations

between the largest fashion and beauty brands and relatively unknown influencers, where the influencers' work is shown to the brands' audiences in a one-off co-authored post. We find that as a result of the exposure from this co-authored post, influencers gain followers, but the average effect is quite small, does not vary much across most of the distribution of brand follower counts, and even for the most successful collaborations leads to only modest gains in followership.

Our research raises several questions for future work. While we have quantified the value of exposure on its own, it remains an open question as to how valuable exposure is relative to other elements of being recognized by brands like prestige or feelings of community. Understanding what is most motivating to creators can help brands craft their messages for soliciting creative content. More broadly, our paper contributes to an ongoing discussion about whether recognizing and attributing creative output—without actually paying for it—is enough to incentivize that output, when the original creator's work is re-broadcast through an aggregator, search engine summary, or LLM query response. The general equilibrium implications of whether such aggregators can sufficiently incentivize content with recognition and exposure, not payment, is a pressing question as individuals' attention online is increasingly captured by a few large players.

7 Funding and Disclosure

Daniel Goetz acknowledges funding from the TD Management and Data Analytics Lab. Both authors have no conflicts of interest to report.

References

Abidin, Crystal, "Visibility labour: Engaging with Influencers' fashion brands and# OOTD advertorial campaigns on Instagram," *Media International Australia*, 2016, *161* (1), 86–100.

Angrist, Joshua D and Alan B Krueger, "Empirical Strategies in Labor Economics. Handbook of Labor Economics, Volume 3A. o. Ashenfelter and D. Card, eds," 1999.

Aridor, Guy, Rafael Jiménez-Durán, Ro'ee Levy, and Lena Song, "The Economics of Social Media," *Journal of Economic Literature*, December 2024, *62* (4), 1422–74.

Athey, Susan, Markus Mobius, and Jeno Pal, "The impact of aggregators on internet news consumption," Technical Report, National Bureau of Economic Research 2021.

Bairathi, Mimansa and Anja Lambrecht, "Influencer Marketing: Sponsorship Disclosure and Authenticity," *Available at SSRN 4679776*, 2023.

Beichert, Maximilian, Andreas Bayerl, Jacob Goldenberg, and Andreas Lanz, "Revenue generation through influencer marketing," *Journal of Marketing*, 2023.

Berthon, Pierre, Leyland Pitt, and Colin Campbell, "Ad lib: When customers create the ad," *California management review*, 2008, *50* (4), 6–30.

Bronnenberg, Bart J, Jean-Pierre Dubé, and Robert E Sanders, "Consumer misinformation and the brand premium: A private label blind taste test," *Marketing Science*, 2020, *39* (2), 382–406.

Burtch, Gordon, Qinglai He, Yili Hong, and Dokyun Lee, "How do peer awards motivate creative content? Experimental evidence from Reddit," *Management Science*, 2022, *68* (5), 3488–3506.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Calzada, Joan and Ricard Gil, "What do news aggregators do? Evidence from Google News in Spain and Germany," *Marketing Science*, 2020, *39* (1), 134–167.

Carpentier, Chantal Line, "New Economics for Sustainable Development: Attention Economy," *United Nations Economist Network [pdf]. Dostupné také z: www. un. org/sites/un2. un. org/files/attention _economy_feb. pdf*, 2023.

Cheng, Mengjie and Shunyuan Zhang, "Reputation burning: Analyzing the impact of brand sponsorship on social influencers," *Management Science*, 2024.

__ , **Maximilian Beichert, Xitong Li, and Shunyuan Zhang**, "Influencer Selection Using Inverse Reinforcement Learning," *Working Paper*, 2023.

Ching, Andrew T and Daniel Goetz, "Consumption Responses to an Unpopular Policy: Evidence from a Short-Lived Soda Tax," *Marketing Science*, 2024, *43* (6), 1204–1225.

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

Gallus, Jana, "Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia," *Management Science*, 2017, *63* (12), 3999–4015.

Ghosh, Arpita and Preston McAfee, "Incentivizing high-quality user-generated content," in "Proceedings of the 20th international conference on World wide web" 2011, pp. 137–146.

Goes, Paulo B, Chenhui Guo, and Mingfeng Lin, "Do incentive hierarchies induce user effort? Evidence from an online knowledge exchange," *Information Systems Research*, 2016, *27* (3), 497–516.

Gordon, Brett R, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky, "A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook," *Marketing Science*, 2019, *38* (2), 193–225.

Hayashi, Fumio, *Econometrics*, Princeton University Press, 2011.

Huang, Yufeng and Ilya Morozov, "Video Advertising by Twitch Influencers," *Available at SSRN 4065064*, 2023.

Hwang, Serim, Xiao Liu, and Kannan Srinivasan, "Voice analytics of online influencers," *Available at SSRN 3773825*, 2021.

Kanuri, Vamsi K, Christian Hughes, and Brady T Hodges, "Standing out from the crowd: When and why color complexity in social media images increases user engagement," *International Journal of Research in Marketing*, 2024, *41* (2), 174–193.

Khernamnuai, Warut, Karthik Kannan, and Hossein Ghasemkhani, "Extrinsic versus intrinsic rewards for contributing reviews in an online platform," *Information Systems Research*, 2018, *29* (4), 871–892.

Li, Nan, Avery Haviv, and Mitchell J Lovett, "Let's Play Fair - Purchase and Usage Effects of Influencer Marketing on YouTube," *Available at SSRN 3884038*, 2023.

Luca, Michael, "Chapter 12 - User-Generated Content and Social Media," in Simon P. Anderson, Joel Waldfogel, and David Strömberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 563–592.

Mitchell, Matthew, "Free ad (vice): Internet influencers and disclosure regulation," *The RAND Journal of Economics*, 2021, *52* (1), 3–21.

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

Rajaram, Prashant and Puneet Manchanda, "Video influencers: Unboxing the mystique," *arXiv preprint arXiv:2012.12311*, 2020.

Reimers, N, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," *arXiv preprint arXiv:1908.10084*, 2019.

Roberts, Deborah, Mathew Hughes, and Kia Kertbo, "Exploring consumers' motivations to engage in innovation through co-creation activities," *European Journal of Marketing*, 2014, *48* (1/2), 147–169.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023.

Sen, Ananya, Tom Grad, Pedro Ferreira, and Jörg Claussen, "(How) does user-generated content impact content generated by professionals? Evidence from local news," *Management Science*, 2024, *70* (9), 6045–6068.

Simonyan, Karen, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

Smith, Andrew N and Eileen Fischer, "Pay attention, please! Person brand building in organized online attention economies," *Journal of the Academy of Marketing Science*, 2021, *49*, 258–279.

Somerfield, Kardi, "Sharing the love: Motivation and reward in brand-related user-generated content creation for social media," 2023.

_ , Kathleen Mortimer, and Geraint Evans, "The relevance of images in user-generated content: A mixed method study of when, and why, major brands retweet," *International Journal of Internet Marketing and Advertising*, 2018, *12* (4), 340–357.

Sun, Yacheng, Xiaojing Dong, and Shelby McIntyre, "Motivation of user-generated content: Social connectedness moderates the effects of monetary rewards," *Marketing Science*, 2017, *36* (3), 329–337.

Tian, Zijun, Ryan Dew, and Raghuram Iyengar, "Mega or micro? Influencer selection using follower elasticity," *Journal of Marketing Research*, 2024, *61* (3), 472–495.

Villanueva, Julian, Shijin Yoo, and Dominique M Hanssens, "The impact of marketing-induced versus word-of-mouth customer acquisition on customer equity growth," *Journal of marketing Research*, 2008, *45* (1), 48–59.

Yang, Jeremy, Juanjuan Zhang, and Yuhan Zhang, "Engagement that Sells: Influencer Video Advertising on TikTok," *Available at SSRN 3815124*, 2023.

Zhang, Shunyuan, Dokyun Lee, Param Vir Singh, and Kannan Srinivasan, "What makes a good image? Airbnb demand analytics leveraging interpretable image features," *Management Science*, 2022, *68* (8), 5644–5666.

Appendix

A Data construction

A.1 Tagging and Partnership/Branded Posts

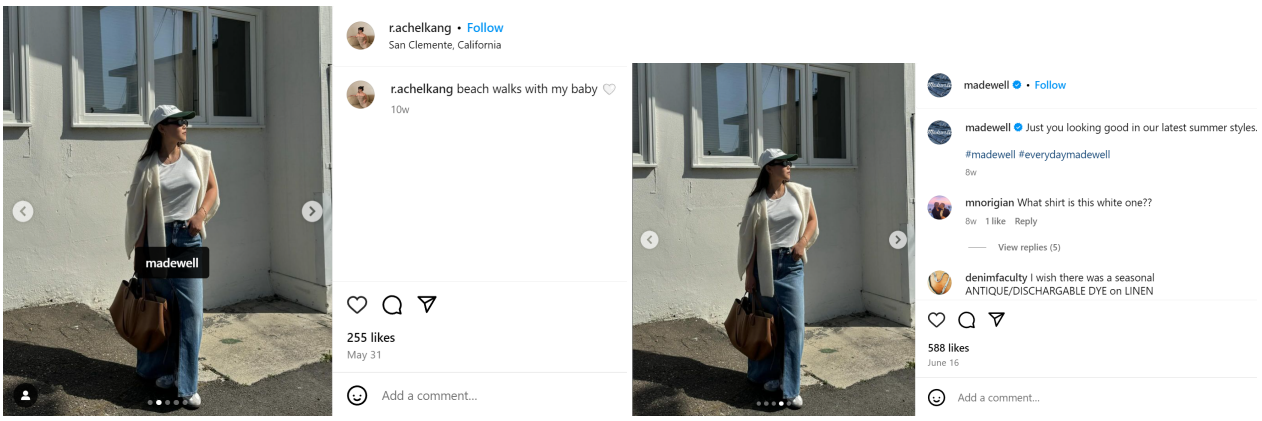
A.1.1 Tagging

Tagging is an alternative way that individuals on Instagram can credit other accounts in a post. As an example of tagging, in Figure A.1a, the content creator **rachelkang** tags the account **madewell**, which is a clothing company. The tagged company account name **madewell** is visible on this photo after a user taps it, and tapping the name will take the user directly to the brand account page.

Compared to co-authored posts, tagged posts have several disadvantages which make them less amenable to a study of the value of exposure. First, and as can be seen in Figure A.1a, tags are not always made in recognition of contributed content—**rachelkang** tagged **madewell** because she wanted to gain the brand's attention. Indeed, a tag can be made without the tagged account's consent. When examining brand posts, this property of tagging makes it very hard from a data perspective to parse which accounts that a brand tagged actually contributed content. Conversely, in co-authored posts, explicit consent of a co-authoring party is required before that party's account name appears on the post, so it is much more likely that the co-authored post reflects a contribution.

Second, if a brand tags a creator in recognition of contributed content, that tagged content on the brand account often duplicates an image or video the creator posted elsewhere. For instance, in Figure A.1a, the account owner of **rachelkang** posts an image of herself wearing **madewell** jeans and tags the brand account, and the brand then posts that content to their page in Figure A.1b, likely for free but with **rachelkang**'s consent. However, the Instagram algorithm tends not to promote duplicated content, which would lead to less exposure of that content to **madewell**'s audience. Compared with re-used tagged content, co-authored posts are not duplicates, but original content posted to both accounts simultaneously, which obviates this algorithmic bias.

Third, tagged content can often be UGC submitted with a hashtag or brand tag in response



(a) Original photo

(b) Photo with creator tag on brand account

Figure A.1: Original and tagged UGC on Instagram

Note: The left image is the original photo, posted by the creator rachelkang on her Instagram account and tagging madewell, the clothing brand. The tag brings the post to madewell's attention, and they repost the content on their brand page in the right image and tag rachelkang (not visible). Tagging is an alternative way of crediting content contributions compared to co-authorship.

to a contest. For example, in Apple's "Shot on iPhone" campaign in 2019, users were encouraged to use the hashtag #ShotOniPhone to bring photos to Apple's attention. Apple then selected photos with the hashtag to use as part of a marketing campaign (see iPhone challenge rules). If the contest hashtag is trending, then unobserved greater interest in this content will lead the algorithm to promote content connected to the hashtag—including the selected photo—which will contaminate measurement of the exposure effect. Conversely, co-authored posts tend to be one-off, and not part of trending hashtags, which provides a cleaner quasi-experiment to identify the effect of exposure.

A.1.2 Partnership/branded posts

Branded content ads (renamed partnership ads in May 2023) are posts made by the creator, but which tag or mention the brand. The creator can allow the brand to have access to the post for purposes of "boosting" the post. Boosting is a paid service that Instagram offers to business accounts whereby the post is promoted in the Instagram algorithm to gain additional views.

Partnership/branded posts are fundamentally different than collaborative posts. First, a partnership post will not be shown to the brand's followers. While a partnership post may be more widely viewed than a creator's typical post, the audience depends on the extent to which the post is boosted, which is unobservable. Second, a collaborative post is not necessarily an advertise-

ment, whereas branded posts represent an explicit contract between the brand and creator for the creator to promote the brand's products.

A.2 Control pool of posts

Sourcing additional control posts by never-treated creators proceeds in several steps. First, we generate a list of Instagram account names. Instead of randomly sampling from the universe of Instagram accounts, we use Instagram's recommendation algorithm to suggest related accounts for each treated creator. The algorithm lists up to 80 accounts per creator, and we gather 16,467 unique account names. We drop all accounts that are categorized as businesses, yielding 10,013 non-business accounts as candidate never-treated creators.

Gathering post histories and follower counts for such a large sample of accounts is prohibitive, and many of these accounts may not post content that is a good match for the treated content. We therefore select two unique related accounts per treated creator that have the same category (e.g. artist, athlete, creator, digital artist, etc.) and a similar number of followers at time of data collection. Specifically, we discretize the full distribution of follower counts into 20 intervals, and select accounts that are in the same interval as the original treated account to which they are related. If there are more than two related accounts with the same category and in the same follower interval as the original treated account, then we randomly select two. This procedure guarantees that the control posts we gather will be made by creators covering similar topics, and with a similar number of followers, as the co-authored posts by treated creators. Note that we are not explicitly matching control creators to treated creators, but rather selecting accounts that have the best chance of supplying control posts that strongly resemble the co-authored posts.

Of the related accounts, we are able to collect followership history on SocialBlade for 625 accounts. We then pull the complete post history for these 625 creators, which yields approximately 220,000 non-co-authored candidate control posts. Combined with the pre-treatment control posts from the co-authoring creators, we have approximately 430,000 non-collaborative posts in total from which to select matched control posts.

A.3 Post content

Before matching treatment posts to control posts, we must extract post content. We proceed by first using ChatGPT 4o's vision capabilities to describe images. We then embed these descrip-

tions, perform a dimension reduction of the embedding, and then perform clustering, selecting an optimal number of clusters using the elbow method. Finally, we use TF-IDF to identify human-interpretable topics.

A.3.1 Using ChatGPT 4o to describe images

Media on Instagram comes in two forms: images and videos. For videos, we take snapshots every 3 seconds for the first 15 seconds. For both images and video snapshots, we reduce the resolution to a 300 x 450 pixel resolution. We then prompt ChatGPT 4o to analyze each image and video snapshot, using the prompt "Describe this image" in all cases.

To showcase the outcome of this process, we select five example images from the set of treatment (co-authored) posts in Table A.1 along with the recovered image description from ChatGPT. ChatGPT 4o is able to recover not just the number of individuals and their emotional state, but also highly-idiosyncratic individual features like exact details of their clothes (e.g. "white checkered sleeveless vest", "heart-shaped belt buckle", "loose-fitting jeans"), their pose ("kneeling", "sitting"), detailed elements of the backdrop ("a snowy landscape with trees", "a cloudy sky with a warm, golden hue"), and text that appears on the image. The description of the images also often provides holistic, human-interpretable text that either directly captures both valence and arousal ("calm and inspiring", "smiling and ... having a good time") or provides context clues that allow these features to be understood.

The total corpus of creator media comprises over 2.24 million images in total. Using ChatGPT 4o, in November 2024 the price of generating all descriptions was $645 USD.

A.3.2 Embedding and clustering image descriptions

Our aim is to cluster the images into topics, and to use these topics as representations of the content of posts so that treated post topics can be matched to control post topics.

The first step is to transform the text descriptions into numeric embeddings. While in theory it is possible to skip the step of recovering the text description from ChatGPT and directly query ChatGPT to provide numeric embeddings of the content of images, at time of writing these visual embeddings were not available to be called from the API.

To recover embeddings of the text descriptions, we use the pre-trained *paraphrase-mpnet-base-v2* sentence transformer, a transformer model with 109 million parameters also referred to as SBERT. SBERT embeddings recover contextual information at the paragraph level, unlike base

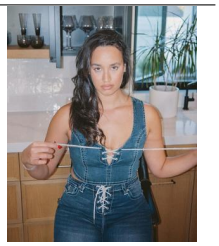
					
Post ID	CfKlh8SvyYq	Csb2EXCO62z	Ctjjv90gdyH	CuAGAqxLzO6	C0SGnBUPeeK
ChatGPT 4o description	The image features a tall stadium light pole against a backdrop of a cloudy sky with a warm, golden hue. There's a motivational quote overlaid on the image that reads, ""I've already won as long as I show up as myself."" The overall tone of the image is calm and inspiring.	The image shows two people outdoors enjoying the snow. They are smiling and appear to be having a good time, with their arms around each other. One person is wearing a hat and the other is in a T-shirt with colorful graphics. Both are wearing sunglasses, and the background features a snowy landscape with trees.	The image shows a person standing outdoors near a display of colorful flowers in pots. They are wearing sunglasses, a black and white checkered sleeveless vest, a white T-shirt, and black shorts with a heart-shaped belt buckle. The person has short, dark hair and visible tattoos on their arm and legs. The background includes various plants and a structure with a lattice roof.	The image is a black and white photograph of a person kneeling on the ground. They are wearing a cropped t-shirt and loose-fitting jeans. The individual's hands are placed behind their head. The setting appears to be minimal, with a simple, plain background.	The image shows a person with long dark hair sitting in a kitchen setting. They are wearing a denim top and jeans that both have lace-up designs. The person is holding a measuring tape. The background includes kitchen cabinets, countertops, and potted plants.

Table A.1: Example images with ChatGPT descriptions

Note: Post ID links to the post on Instagram, the link can be manually created using "https://www.instagram.com/p/" + post ID. Posts typically contain multiple images, the image shown here is a random selection from within the post.

BERT which analyzes text at the sentence level. This is important in our setting: for instance, in Table A.1, the final sentence in the second image starts with "Both are wearing sunglasses...", which references the "two people outdoors enjoying the snow". SBERT would correctly attribute "both" as referring to the two people outdoors since it uses the entire paragraph. For each image and each video snapshot, we recover the SBERT embedding, which is a 768-dimensional vector.

In theory we could cluster images into topics based on the SBERT embeddings. However, clustering requires holding all of the image embeddings in memory to form clusters. We have 2.24 million images, and the memory requirements are too large to perform the clustering using the 768-dimensional vector embedding. We therefore do dimension reduction. The Uniform Manifold Approximation and Projection (UMAP) is a popular technique for dimension reduction when the end goal is to cluster data, since it works to preserve local structure. UMAP is also memory intensive, since it constructs a complete graph of the data, where each point is a node and each vector encodes local similarity between that node and other nodes, and it is infeasible to perform using the complete data set with the SBERT embeddings. We therefore first use Principal

Components Analysis, a fast, linear method that identifies new axes in the data which maximize the remaining variance in the data, which helps preserve global structure by minimizing information loss. We use PCA to reduce the SBERT embeddings to 200 dimensions, then use UMAP on the 200-dimensional PCA vectors to reduce the embeddings to 5 dimensions.

A.4 Propensity Score Matching

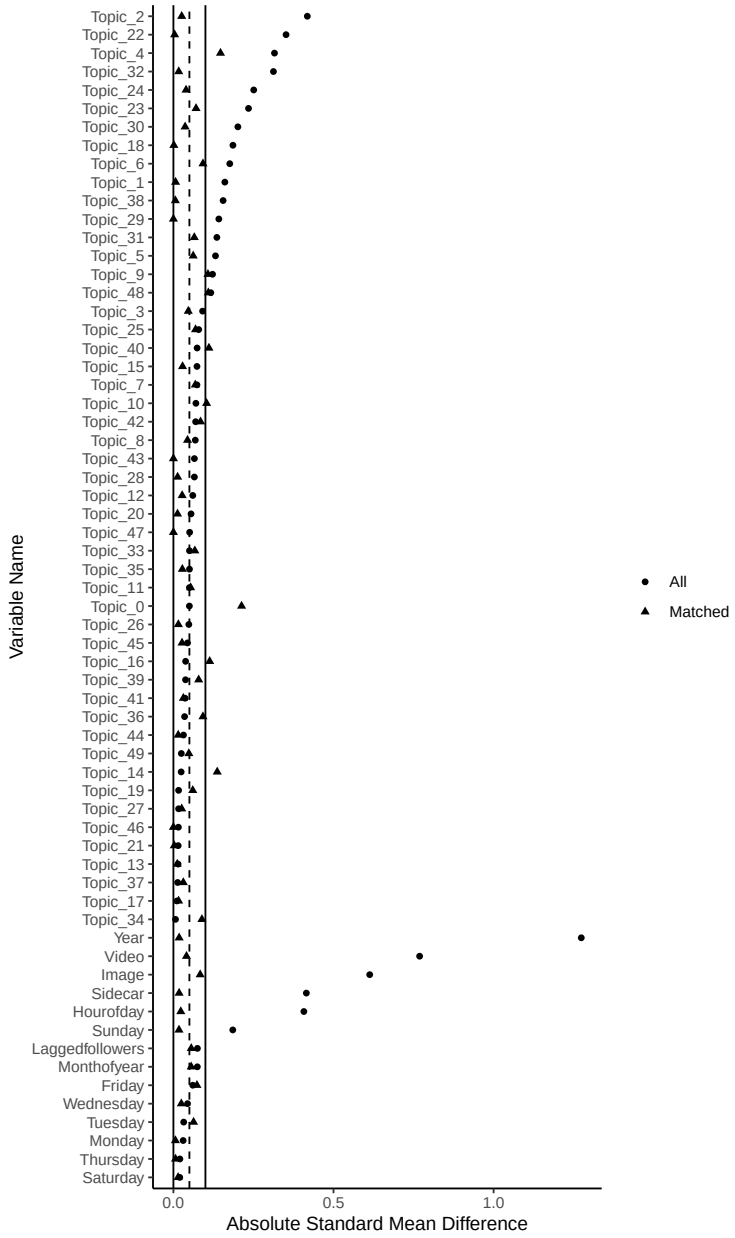


Figure A.2: Love Plot from propensity score matching

Note: Matching variables are grouped into the 50 topics extracted from the post media, and the hand-selected features below, and then sorted within each group in order of decreasing absolute standard mean difference in the unmatched (All) data. The absolute standardized mean difference is computed as $\left|(\bar{X}_{treat} - \bar{X}_{control})/\hat{\sigma}_{treat}\right|$.

Table A.2: Balance table from propensity score matching

	Unmatched controls			Matched controls		
	Means Treated	Means Control	Std. Mean Diff.	Means Treated	Means Control	Std. Mean Diff.
Distance	0.005	0.0011	0.7594	0.005	0.005	0
HourOfDay	15.4256	13.8221	0.4074	15.4256	15.3349	0.023
Friday	0.1744	0.1514	0.0606	0.1744	0.2023	-0.0735
Monday	0.1628	0.1516	0.0304	0.1628	0.1605	0.0063
Saturday	0.1116	0.1179	-0.0198	0.1116	0.107	0.0148
Sunday	0.0767	0.1261	-0.1855	0.0767	0.0721	0.0175
Thursday	0.1419	0.1489	-0.0202	0.1419	0.1395	0.0067
Tuesday	0.1628	0.1509	0.0321	0.1628	0.1395	0.063
Wednesday	0.1698	0.1532	0.0441	0.1698	0.1791	-0.0248
MonthOfYear	6.4163	6.1661	0.0747	6.4163	6.6023	-0.0556
Year	2023.2	2022.2	1.2738	2023.2	2023.2	0.018
Lagged Followers	59968	56393	0.075	59968	62620	-0.0558
Image	0.1047	0.2923	-0.6131	0.1047	0.0791	0.0836
Sidecar	0.1884	0.3507	-0.4152	0.1884	0.1953	-0.0178
Video	0.707	0.357	0.769	0.707	0.7256	-0.0409
Topic 0	0.0758	0.0662	0.0499	0.0758	0.1171	-0.2129
Topic 1	0.0813	0.051	0.1606	0.0813	0.0826	-0.0068
Topic 2	0.0148	0.0476	-0.4182	0.0148	0.0128	0.0259
Topic 3	0.0687	0.0525	0.0911	0.0687	0.0604	0.0465
Topic 4	0.0258	0.0591	-0.3157	0.0258	0.0413	-0.1471
Topic 5	0.0683	0.045	0.1315	0.0683	0.0573	0.0618
Topic 6	0.0143	0.0313	-0.176	0.0143	0.0233	-0.0924
Topic 7	0.0476	0.0352	0.0735	0.0476	0.0361	0.0684
Topic 8	0.0321	0.0227	0.0684	0.0321	0.0382	-0.0441
Topic 9	0.0212	0.0321	-0.1222	0.0212	0.0117	0.1079
Topic 10	0.0467	0.0352	0.0703	0.0467	0.0297	0.1036
Topic 11	0.0226	0.028	-0.05	0.0226	0.0169	0.0538
Topic 12	0.0205	0.0261	-0.0605	0.0205	0.0179	0.0273
Topic 13	0.0349	0.0325	0.0146	0.0349	0.0368	-0.012
Topic 14	0.0341	0.031	0.0242	0.0341	0.0514	-0.1368
Topic 15	0.0224	0.0296	-0.0736	0.0224	0.0252	-0.0286
Topic 16	0.0307	0.0265	0.0381	0.0307	0.0184	0.1133
Topic 17	0.026	0.0245	0.0112	0.026	0.0239	0.0162
Topic 18	0.0469	0.0159	0.1859	0.0469	0.0467	0.0014
Topic 19	0.0205	0.0223	-0.0162	0.0205	0.0142	0.0605
Topic 20	0.0254	0.0189	0.0551	0.0254	0.0239	0.0129
Topic 21	0.0205	0.0222	-0.015	0.0205	0.0208	-0.0028
Topic 22	0.0058	0.0204	-0.3518	0.0058	0.0057	0.0037
Topic 23	0.0069	0.0172	-0.2346	0.0069	0.0037	0.0708
Topic 24	0.0047	0.0183	-0.2509	0.0047	0.0025	0.0395
Topic 25	0.0092	0.0151	-0.079	0.0092	0.0144	-0.0691
Topic 26	0.0072	0.0103	-0.0482	0.0072	0.0062	0.0155
Topic 27	0.0153	0.0135	0.0162	0.0153	0.0124	0.0259
Topic 28	0.0197	0.0122	0.0655	0.0197	0.0182	0.0132
Topic 29	0	0.0142	-0.1419	0	0	0
Topic 30	0.0058	0.0182	-0.2014	0.0058	0.0081	-0.0365
Topic 31	0.0263	0.0111	0.1357	0.0263	0.019	0.0653
Topic 32	0.0026	0.0098	-0.3124	0.0026	0.003	-0.0167
Topic 33	0.0045	0.0071	-0.0502	0.0045	0.001	0.0669
Topic 34	0.0074	0.0069	0.0066	0.0074	0.0138	-0.089
Topic 35	0.0047	0.0074	-0.0502	0.0047	0.0031	0.0279
Topic 36	0.0091	0.0067	0.0349	0.0091	0.0027	0.0923
Topic 37	0.0099	0.0088	0.0133	0.0099	0.0073	0.0309
Topic 38	0.0031	0.0077	-0.1552	0.0031	0.0033	-0.006
Topic 39	0.0093	0.006	0.0378	0.0093	0.0023	0.0789
Topic 40	0.0131	0.0061	0.0741	0.0131	0.0236	-0.111
Topic 41	0.0085	0.0058	0.0364	0.0085	0.0062	0.031
Topic 42	0.0085	0.0033	0.0694	0.0085	0.0149	-0.0848
Topic 43	0	0.0039	-0.0655	0	0	0
Topic 44	0.0035	0.0051	-0.0315	0.0035	0.0043	-0.0151
Topic 45	0.0058	0.003	0.0442	0.0058	0.0076	-0.027
Topic 46	0.0023	0.0016	0.0156	0.0023	0.0023	0
Topic 47	0	0.0024	-0.0507	0	0	0
Topic 48	0.0047	0.0012	0.1168	0.0047	0.008	-0.1091
Topic 49	0.0008	0.0012	-0.0247	0.0008	0	0.0482

Note: Std. Mean Diff. is the standardized mean difference, $(\bar{X}_{treat} - \bar{X}_{control})/\hat{\sigma}_{treat}$

B Identification

B.1 Timing test - post coordination with external shocks

Figure B.1 provides examples of local maxima in search volumes near co-authored post timings. We aggregate the panel of co-authored posting timing to the creator×week level, and run the following regression:

$$\text{CoAuthorPost}_{jt} = \beta_0 + \sum_{s=-6}^{6} \beta_s \text{LocalMax}_{t-s} + \alpha_j + \epsilon_{jt}$$

where j indexes a creator, t indexes a week, CoAuthorPost_{jt} is an indicator for whether j has their co-authored post published in week t, and LocalMax_{t-s} is an indicator for whether the local maximum in search volume occurred at time $t-s$, for $s = -6, \ldots, 6$. To understand the regression, consider the coefficient β_{t-3}. If this coefficient were to be positive and significant, it would indicate that there is a greater propensity for co-authored posts 3 weeks after the local maximum in search volumes.

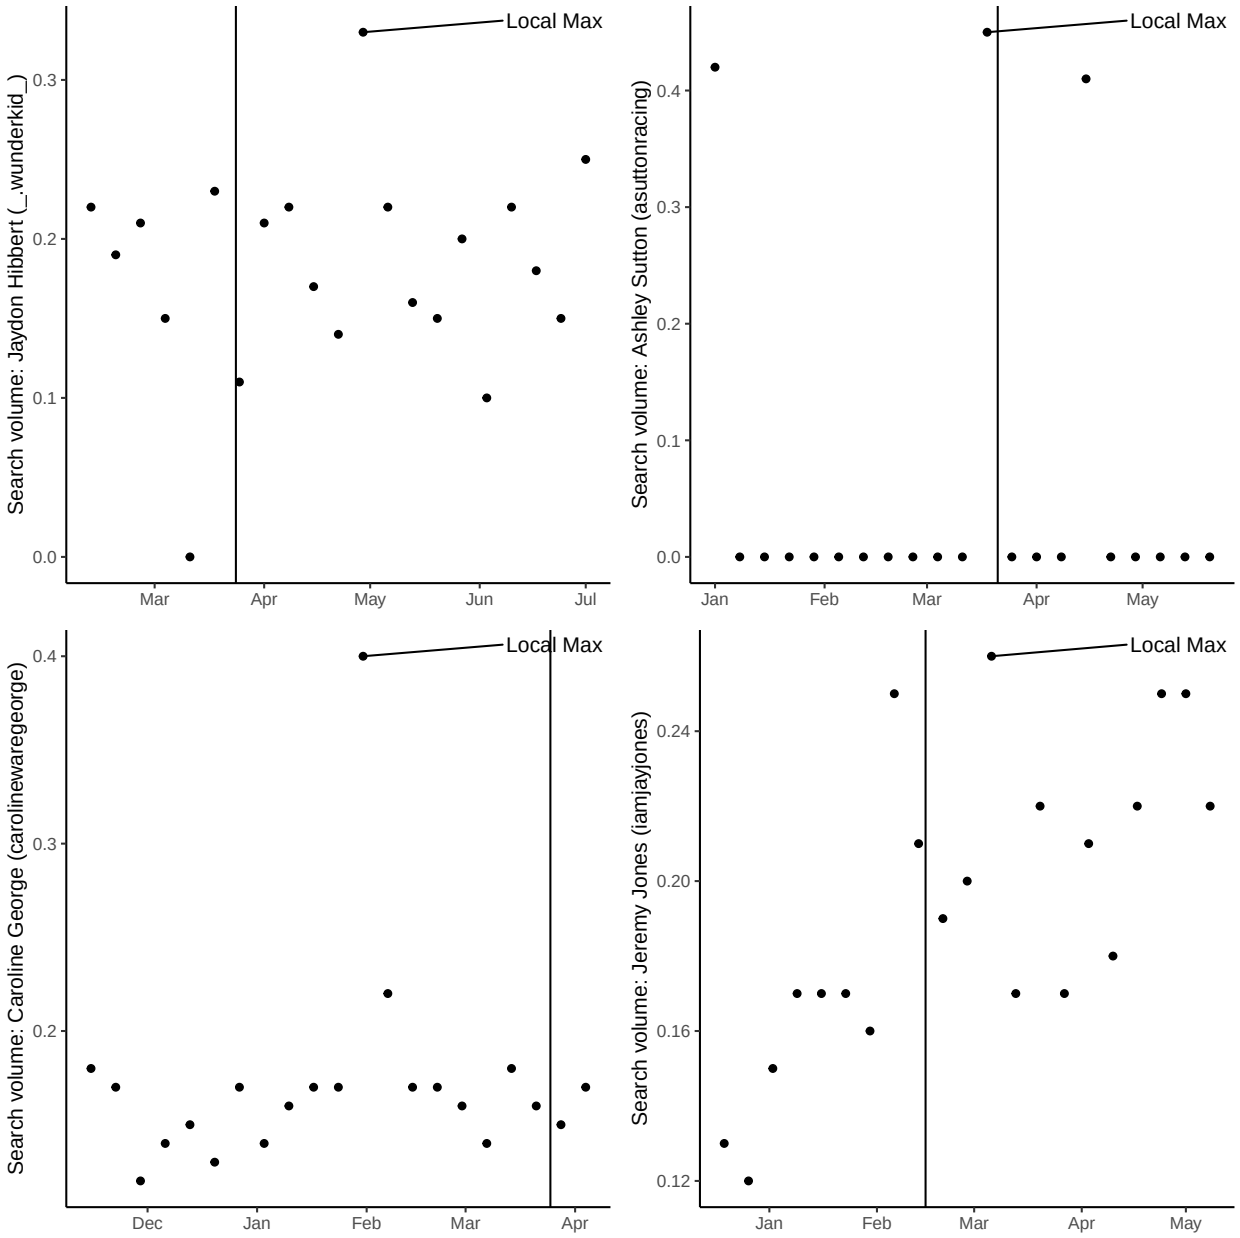


Figure B.1: Examples of local maxima in search volumes near co-authored posts

Note: Each vertical line is the date on which the co-authored post is published on Instagram for that creator. Search volumes are scraped at the weekly level from Google trends, and normalized by dividing by 100. Years are 2024, 2022, 2024, and 2023 moving from left to right starting at the upper left.

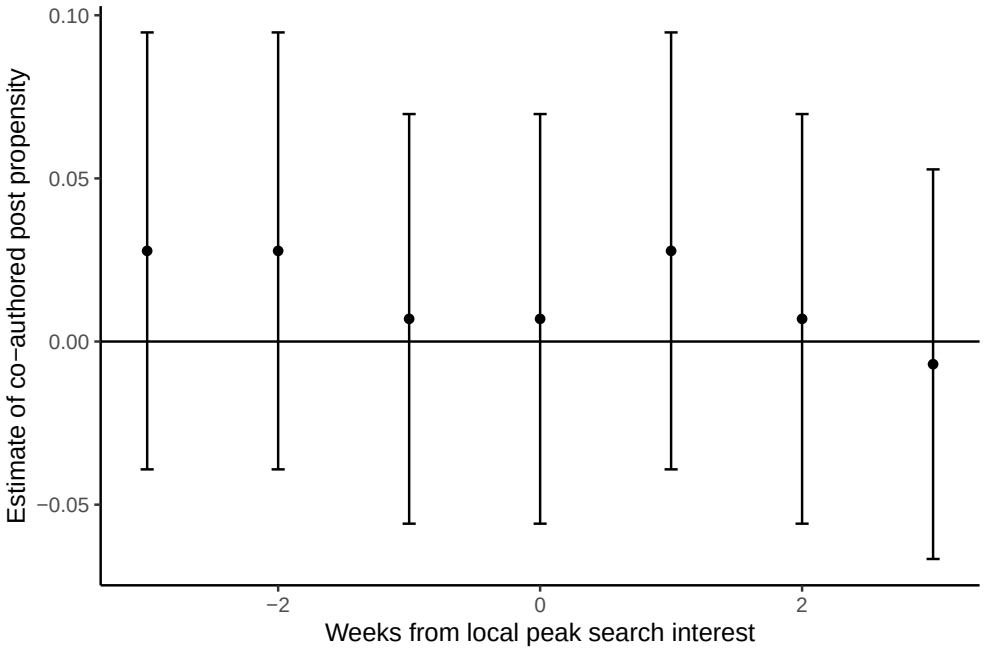


Figure B.2: Co-authored posting around times of peak search interest in creator, within a 3 week window

Note: This figure replicates the analysis in Figure 5, except defining the local search maximum within 3 weeks of the co-authored post. This figure is not equivalent to simply truncating Figure 5, since the timing of the local search maximum within 3 weeks of the co-authored post may be a different week compared to the timing of the local search maximum within 6 weeks of the co-authored post.

B.2 Event study analysis

To formally explore whether there are pre-trends in the data, we specify the following event study model:

$$\text{FollowerGrowth}_{jt} = \sum_{s=-21}^{42} \beta_s \cdot \text{Treat}_{j,t-s} + \alpha_j + \alpha_t + \alpha_{r(t)} + u_{jt} \tag{B.1}$$

Instead of estimating a single cumulative effect β with decay given by $f(s, \theta)$, we estimate a β_s coefficient for the relative effect of being in the treatment group up to 3 weeks prior to treatment, and 6 weeks post treatment. We expect that with no pre-trend, there should be no relative effect of being in the treatment group pre-treatment, and so β_s should not be statistically significantly different from zero for $s < 0$.

Coefficients are plotted in Figure B.3. The 95% confidence intervals for each β_s with $s < 0$ clearly cover 0, and thus we fail to reject the hypothesis of no pre-trends using this framework.

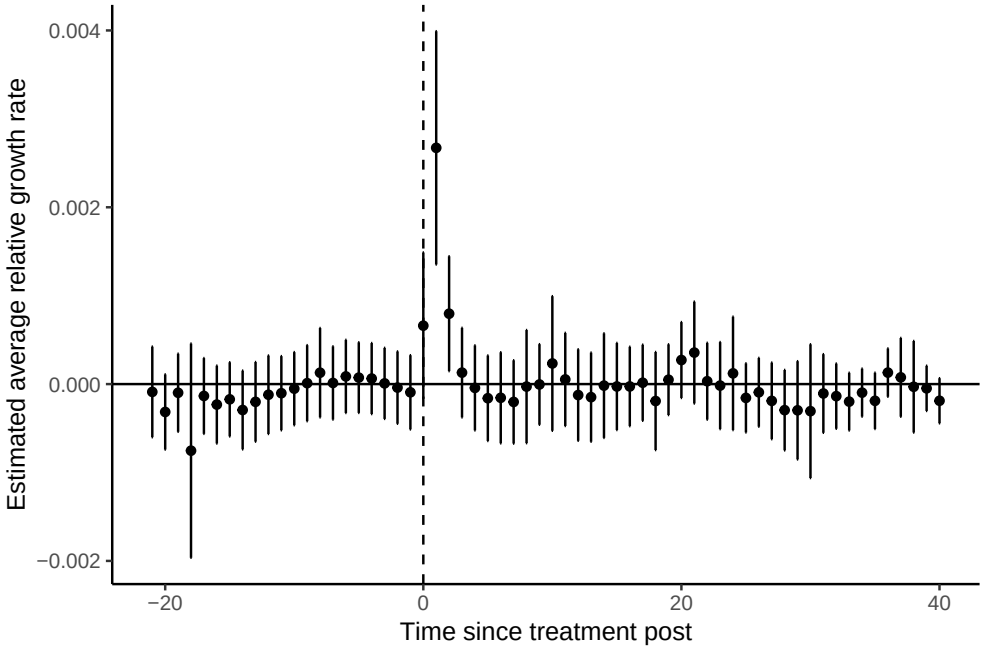


Figure B.3: Event study specification, follower growth rates

Note: This graph plots the estimated coefficients $\hat{\beta}_s$ from Equation (B.1) using follower growth as the dependent variable. Error bars are 95% confidence intervals, clustered at the level of the creator.

B.3 Linear pre-trend analysis

Angrist and Krueger (1999) propose a test for pre-trends, which involves estimating a differential linear trend on the pre-treatment data for treated and control units. In the marketing literature, Bronnenberg et al. (2020) and Ching and Goetz (2024) implement this test to help validate parallel trends assumptions. We estimate:

$$Y_{jt} = \delta_1 \cdot t + \delta_2 \cdot \text{Treat}_j \cdot t + \alpha_j + \alpha_t + u_{jt} \qquad \text{for } t \in [-28, -1] \tag{B.2}$$

for each of the outcome variables, t is the time-to-post, and Treat_j is an indicator variable that equals one if creator j is a treated unit. δ_1 will capture the underlying joint linear pre-trend for both treatment and control creators in the 30 days prior to posting, while δ_2 estimated significantly different from 0 would capture a differential pre-trend for treated creators.

Estimates from Equation (B.2) are presented in Table B.1. We find no differential linear trend for treated creators prior to treatment. In additional specifications we also try looking at shorter time windows before posting, of 10 days and 5 days, and continue to find no evidence of a differential pre-trend for treated creators pre-treatment.

Table B.1: Testing for linear pre-trends

	Dependent variable:		
	Daily Follower Growth		
	(1)	(2)	(3)
t	-0.00002***	0.00001	0.0001**
	(0.00001)	(0.00001)	(0.00002)
$t \times$ Treat	0.00001	-0.000005	-0.00005
	(0.00001)	(0.00001)	(0.00003)
No. Days Prior to Treatment.	30	10	5
IG Account FE	✓	✓	✓
Observations	24,688	8,393	4,232
R^2	0.3884	0.5838	0.7280

*Note: This table reports estimates from Equation (B.2). Column 1 includes 30 days for each treated and control creator prior to treatment, Column 2 includes 10 days, and Column 3 includes 5 days. Standard errors are clustered at the IG account level. *p<0.1; **p<0.05; ***p<0.01*

B.4 Event study with alternate dependent variable: follower growth

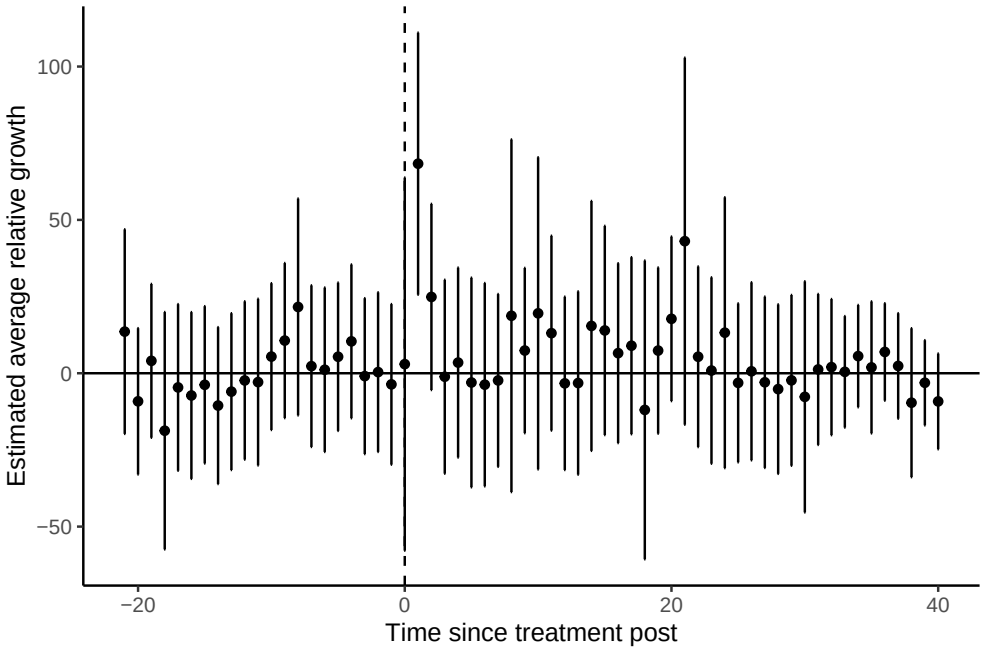


Figure B.4: Event study specification, follower growth

Note: This graph plots the estimated coefficients $\hat{\beta}_s$ from Equation (B.1), estimated using follower growth (instead of growth rate) as the dependent variable. Error bars are 95% confidence intervals, clustered at the level of the creator.

C Additional Analysis

C.1 Effect of sponsored posting on followers

In this section, we estimate the effect of a sponsored post on followers, as in Cheng and Zhang (2024). The purpose of this exercise is to determine whether creators experience a large drop in existing followers when they create a sponsored post. If they do, and if their follower base perceives co-authored posts as sponsored, then the exposure effect represents the sum of gaining new followers from the brand's audience, and losing existing followers from their current audience.

We first describe how we detect sponsored posts in the data in Section C.1.1, and then estimate the effect of these posts on follower growth in Section C.1.2.

C.1.1 Detecting sponsored posting

We first code a dummy variable that captures whether a post is sponsored or not. To do so, we follow Cheng and Zhang (2024) and Hwang et al. (2021) and use a dictionary-based approach based on disclosure in the post's caption. We lemmatize each caption, and then use search for matches using the stems of the following terms: "pay", "collaboration", "compensate", "#ad", "sponsored", "ambassador", "#partner", "gifted", "gifting". For any matches that are preceded by a negation ("not", "never", "no way", "disclaimer"), we code the post as not sponsored.

Since co-authored posts may be sponsored, there will be a mechanical positive relationship between co-authorship and sponsored posting if we do not remove sponsored posts with the co-authoring brand. We search for the brand's account name in the co-authoring influencer's caption, and recode the sponsored post dummy as zero if the name is detected. Thus, the sponsored post dummy captures relationships with brands other than the co-authoring brand. We find that on average, the propensity to create a sponsored post on any given day is 0.018. 58% of creators do not make any sponsored posts at all, and the prevalence of sponsored posts is approximately 6.1%.

C.1.2 Effect of sponsored posting on followers

To estimate the effect of sponsored posting on followers, we replicate the OLS analysis of Cheng and Zhang (2024) in our context. We also try using a modified of our TWFE specification.

For all creators in our dataset, we take the identified sponsored posts from Section C.1.1 as the treatment posts, and all other (non-coauthored) posts as potential controls. We match treatment posts to controls as in Section A.4, using propensity score matching on the same set of observable covariates as we use in the main analysis. We are left with a set of 22,172 treatment posts and 22,172 control posts.

Our OLS specification is:

$$\text{FollowerGrowth}_{jk} = \beta_0 + \beta_1 \text{SponsoredPost}_{jk} + \boldsymbol{\gamma}' \mathbf{W}_{jk} + u_{jk} \tag{C.1}$$

where j still indexes a creator, and k indexes a particular post by that creator. $\mathbf{W}_{jk}$ is the vector of matching variables, which we also include as controls. We observe follower growth on the day of the post, as well as for 3 days after as in Cheng and Zhang (2024). We will therefore multiply β_1 by 4 to get the total cumulative change in followers by 3 days after the post, and multiply by 100 to get the cumulative change in percentages.

Our TWFE specification is Equation (1) but dropping $\mathbf{X}_{jt}$—the controls for brands cross-posting to other social media platforms—and replacing Treat_{jt} with $\text{SponsoredPost}_{jt}$. The estimated β parameter from this specification will once again yield the total cumulative change in followers.

Results from both specifications are presented in Table C.1. In Column 1, a sponsored post is estimated to decrease cumulative followers by 0.013%, while controlling for matching variables in Column 2 leads to a reduction of 0.014%. In Column 3, the TWFE yields a reduction of 0.037%.

There are two key takeaways from this analysis. First, we replicate the direction of the estimates from Cheng and Zhang (2024), and find that sponsored posts do indeed have a reputation burning effect relative to similar-looking non-sponsored posts. Our estimates of the loss in followers are much smaller than theirs, as they estimate a loss of 0.19%. However, our creators are much smaller (the median YouTuber in their data as 1.18 million followers, while our median Instagram creator has fewer than 50,000 followers) and Cheng and Zhang (2024) find that small influencers suffer a smaller reputation burning effect. Second, this analysis sheds light on the relative magnitudes of the two different sources of changes in followers after a co-authored post: gaining followers from the brand's audience, and losing followers from one's own audience. In particular, the exposure effect we estimate of 0.47% is an order of magnitude larger than the reputation burning effect. Even if a creator's existing followers perceive a co-authored post as

Table C.1: Effect of a sponsored post on followers

	Dependent variable:		
	Follower Growth Rate		
	(1)	(2)	(3)
PostSponsored	-0.00003*	-0.00003*	-0.0004*
	(0.00002)	(0.00002)	(0.0002)
θ			6.5576*
			(3.0199)
Controls		✓	
IG Account FE			✓
Date FE			✓
Time-To-Post FE			✓
Avg. Value of Dep. Var.	0.0011	0.0011	0.0011
Cume effect	-0.013%	-0.014%	-0.037%
Observations	173,244	173,244	1,782,243
R^2	0.00002	0.0234	0.2807

Note: The first two columns shows results from estimating Equation (C.1), the last column shows results from estimating Equation (1) using sponsored posting as the treatment. For the first two columns, the cumulative effect is calculated by multiplying the PostSponsored coefficient by 4 (since it captures the 4 days post-sponsored post of growth) and then by 100 (to convert to percentages). For Column 3, we multiply the coefficient directly by 100, as it is the cumulative effect as in Equation (1). Standard errors are heteroskedasticity robust in Columns 1 and 2, and clustered at the IG account level in Column 3. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

sponsored and decide to unfollow, this loss of followers is quite small relative to the total exposure effect.

C.2 Other robustness checks

Table C.2: Robustness checks - Part 1

	Dependent variable:		
	Follower Growth Rate		
	(1)	(2)	(3)
β	0.0043***	0.0052***	0.0055***
	(0.0013)	(0.0013)	(0.0012)
θ	1.2878***	1.2316***	1.1724***
	(0.19)	(0.1454)	(0.1408)
IG Account FE	✓	✓	✓
Date FE	✓	✓	✓
Time-To-Post FE	✓	✓	✓
Cross-Posting Controls		✓	✓
Avg. Value of Dep. Var.	0.0008	0.0008	0.0008
Observations	41,176	35,503	38,638
R^2	0.1935	0.2283	0.2811

*Note: This table presents estimates of Equation (1) using different cuts of the data. Column 1 drops any creator-brand pairing where cross-posting by the brand to other social media platforms is detected. In Column 2, we only use posts by never-treated creators as controls, instead of never-treated and not-yet-treated as in the baseline estimates. In Column 3, we only use posts by not-yet-treated creators at least 90 days prior to treatment as controls. Standard errors are clustered at the IG account level. *p<0.1; **p<0.05; ***p<0.01*

Table C.3: Robustness checks - Part 2

	Dependent variable:	
	Follower Growth Rate	
	(3)	(4)
β	0.0039***	0.0037**
	(0.0015)	(0.0016)
θ	1.5973**	1.8030**
	(0.2676)	(0.7725)
IG Account FE	✓	✓
Date FE	✓	✓
Time-To-Post FE	✓	✓
Cross-Posting Controls	✓	✓
Avg. Value of Dep. Var.	0.0008	0.0008
Observations	29,336	20,724
R^2	0.1496	0.1423

*Note: This table presents estimates of Equation (1) using different cuts of the data. In Column 1 (Column 2) we drop any treated creator who creates a post within one day (two days) after the co-authored post. Standard errors are clustered at the IG account level. * p<0.1; ** p<0.05; *** p<0.01*

C.3 Robustness: omitting creators with large followings on other social media platforms

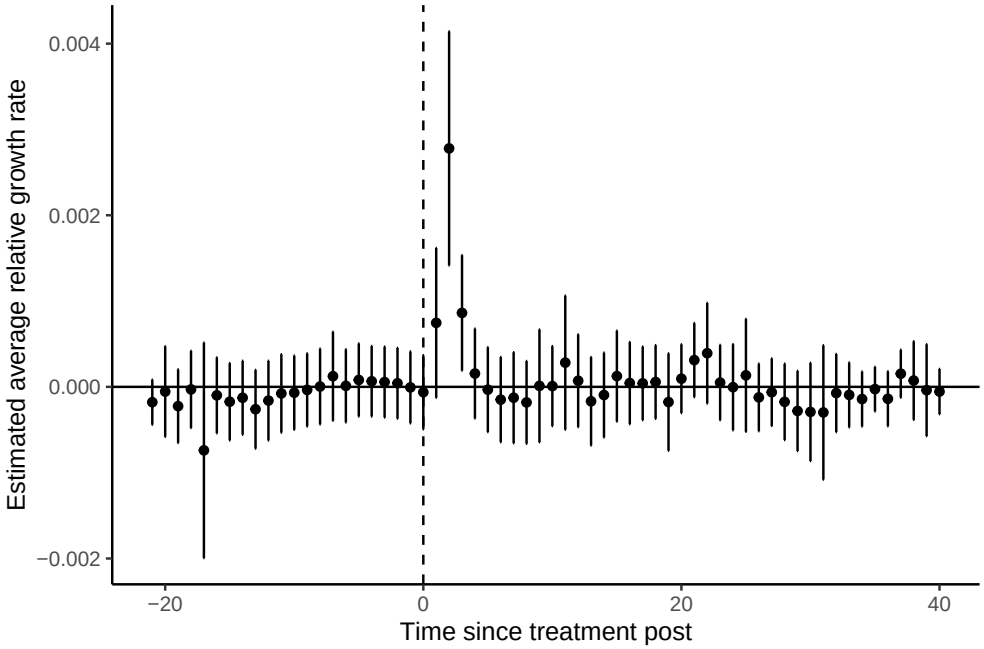


Figure C.1: Event study specification, follower growth rates, dropping creators active on other social media platforms

Note: This graph plots the estimated coefficients $\hat{\beta}_s$ from Equation (B.1). Error bars are 95% confidence intervals, clustered at the level of the creator. Creators who mention another social media profile in their Instagram biography are dropped from the sample.

C.4 Creator account valuation model

We assume that creators primarily monetize their Instagram accounts by endorsing products in sponsored posts in exchange for payments from sponsoring brands. The flow value at time t to a creator j of an account with Followers_{jt} number of followers at time t is:

$$\text{FlowValue}_{jt} = p(\text{Followers}_{jt}) \cdot \Pr(\text{SponsoredPost}_{jt}),$$

which is the price a creator would expect to receive from making a sponsored post, which depends only on the number of followers that creator j has at time t, multiplied by the probability of making a sponsored post, which may vary depending on the creator j and calendar date t. The total discounted value at time t is:

$$\text{TotalValue}_{jt} = \sum_{t' \geq t} \delta^{t'-t} \text{FlowValue}_{jt}$$

The exposure value for creator j at time t as the change in the total discounted value after making a co-authored post, which we calculate by taking the total difference with respect to co-authoring a post:

$$\begin{aligned}\Delta\text{TotalValue}_{jt} = \sum_{t' \geq t} \delta^{t'-t} [& \underbrace{\frac{\Delta p(\text{Followers}_{jt})}{\Delta \text{Followers}_{jt}} \Delta\text{Followers}_{jt}}_{(1)} \cdot \Pr(\text{SponsoredPost}_{jt}) + \\ & p(\text{Followers}_{jt}) \cdot \underbrace{\Delta Pr(\text{SponsoredPost}_{jt})}_{(2)}]\end{aligned} \tag{C.2}$$

Element (1) in Equation (C.2) is the change in the price received per sponsored post as a result of the gain in followers from exposure. It is the product of the relationship between price and followers, and the change in the number of followers. The change in the number of followers for each creator $\hat{\beta}(\mathbf{Z}_j)$ we recover from Table 5, while the change in the price-per-sponsored post as a function of followers we estimate below in Section C.5.

Element (2) in Equation (C.2) is the change in the probability of making a sponsored post, or a direct "business expansion effect' that manifests as the result of making a co-authored post. We discuss how we recover this quantity below in Section C.6.

This simple formula embeds a number of assumptions. First, we assume that the price a cre-

ator can charge for a sponsored post depends primarily on their follower count. Follower count is generally understood in academia as the most important input in influencer pricing (Tian et al., 2024; Beichert et al., 2023). Practitioners also place a key emphasis on follower count for pricing, see for example "Influencer Rates: How To Maximize Your Budget in 2024", *Hootsuite*, July 31, 2024. We validate this assumption in Section C.5 using data from HypeAuditor, a leading influencer marketing platform. Second, we assume an exogenous daily discount rate δ. With a yearly discount of 0.95, the daily discount rate is 0.9998. Third, we assume that sponsored posts are the only way that a creator can monetize their account. In reality, creators may engage in long-term "brand ambassador" arrangements, earn money through affiliate links,

C.5 Sponsored post price

As an input into Equation (C.2), we require an estimate of how sponsored post compensation depends on the number of followers. We use data from HypeAuditor, a leading influencer marketing platform, to recover information on predicted influencer compensation for sponsored posts. For a random sample of 50 treated and control creators in our dataset, we recover HypeAuditor's predicted per-post fee. We then run a simple regression of log per-post fee on log follower count, and find an elasticity of 0.64 with an R^2 value of 0.85. Results are reported in Table C.4.

Table C.4: Sponsored Post Pricing

	Dependent variable:
	log(Sponsored Post Price)
log(Followers)	0.6423***
	(0.0338)
Constant	−0.5917
	(0.3993)
Observations	51
R^2	0.8761
Adjusted R^2	0.8736
Note:	*p<0.1; **p<0.05; ***p<0.01

C.6 Relationship between co-authorship and sponsored posting

Do co-authored posts spark an increase in the number of sponsorships from other brands that creators receive? To answer this question, we estimate the correlation between co-authored posting and sponsored posting.

We first estimate an event-study specification in the style of Equation (B.1). Results in Figure C.2 suggest that on average, there is no visible relationship between sponsored posting and co-authored posting. There is thus no empirical guidance on how to model any decay in sponsored posting. There may still be a heterogeneous treatment effect however, and so we continue to model a Poisson-style decay, which can flexibly capture an increase in sponsored posting that may arise in heterogeneity analysis.

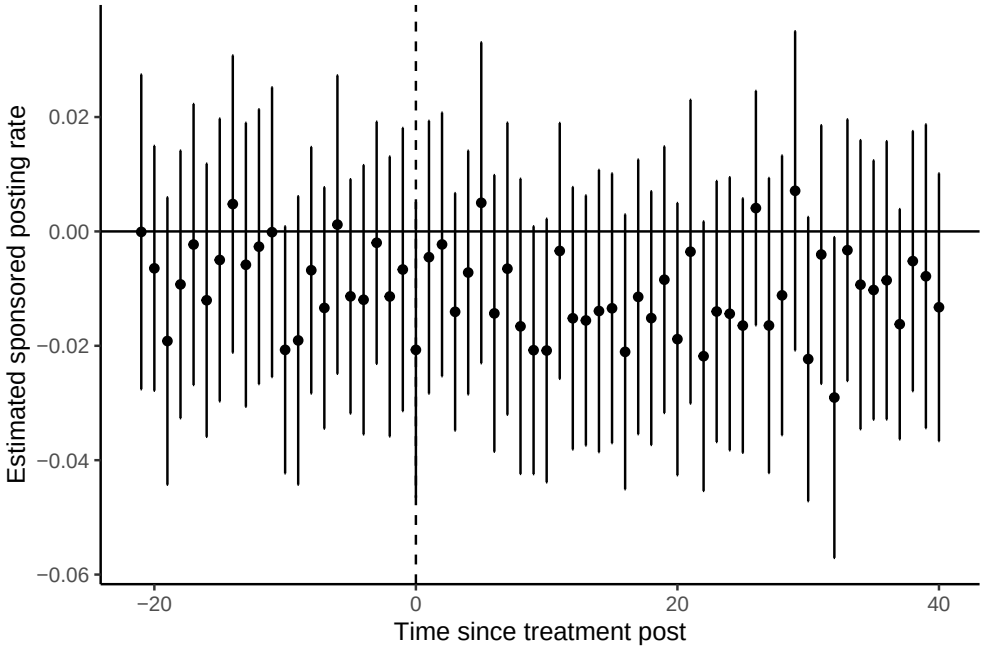


Figure C.2: Event study specification, sponsored posting rate

Note: This graph plots the estimated coefficients $\hat{\beta}_s$ from Equation (B.1) using sponsored posting rate as the dependent variable. Error bars are 95% confidence intervals, clustered at the level of the creator.

We take the sponsored posting dummy as the dependent variable and estimate Equation (3). Results are reported in Table C.5. Both the baseline and heterogeneity coefficients are uniformly insignificant. Altogether, the results suggest no relationship between co-authorship and the quantity of sponsored posts.

Given the findings in Table C.5, we set element (2) in Equation (C.2) to zero when computing the change in value to the creator from the co-authored post. The distribution of values thus

Table C.5: Sponsored posting after the co-authored post

	Dependent variable:
	Sponsored Posting Rate
	(4)
Treat	0.0210
	(0.0319)
Treat×BrandSize	-0.0036
	(0.0227)
Treat×BrandEng	-0.0342
	(0.0255)
Treat×PostBrandDiff	0.0200
	(0.0271)
Treat×PostCreatorDiff	-0.0315
	(0.0251)
Treat×CreatorSize	-0.0251
	(0.0271)
Treat×CreatorEng	-0.0394
	(0.0318)
θ	5.2459
	(9.5104)
IG Account FE	✓
Date FE	✓
Time-To-Post FE	✓
Cross-Posting Controls	✓
Avg. Value of Dep. Var.	0.0161
Observations	49,097
R^2	0.1897

Note: standard errors are clustered at the IG account level. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

depends on the change in sponsored-post-price a creator receives as a result in the increase in followers multiplied by that number of sponsored posts that creator makes.