Content Creator Multihoming and Attention Spillovers Across Platforms

Xinzhi Rao Jingyan Dai Yulin Hao Liangfei Qiu

Oct 2026

Abstract

Creator multihoming makes a competing platform both an alternative destination for attention and a potential source of discovery. This paper examines how second-platform adoption affects creators' audiences, engagement, and content supply on the original platform. We construct a unique dataset of creators on Douyin and RedNote, two leading Chinese content platforms, combining seven years of posting histories with daily follower counts and post-level engagement. Using a staggered difference-in-differences design, we show that RedNote adoption increases Douyin followers by 33%, engagement per post by 23%–26%, and content production by approximately three posts per month. Mechanism analyses suggest that highly engaging RedNote-exclusive posts are followed by increased Douyin follower growth, and that Douyin-exclusive posts receive more likes when they resemble the creator's RedNote offerings. This evidence is consistent with discovery extending beyond shared content. In addition, engagement gains exhibit an inverted-U association with creators' chosen cross-posting shares. Our findings highlight a trade-off between sharing content across platforms to reach new audiences and retaining exclusive content that distinguishes the original platform.

Keywords: creator multihoming, social media, content allocation, attention spillover

Xinzhi Rao (email: xinzhi.rao@warrington.ufl.edu), Warrington College of Business, University of Florida. Jingyan Dai (jingyan.dai@warrington.ufl.edu), Warrington College of Business, University of Florida. Yulin Hao (email: yulin.hao@phbs.pku.edu.cn), HSBC Business School, Peking University. Liangfei Qiu (email: liangfei.qiu@warrington.ufl.edu), Warrington College of Business, University of Florida. We thank Chiara Farronato, Anuj Kumar, Jingchuan Pu, and participants at the 2025 Conference on Information Systems and Technology (CIST 2025) and the 2026 Platform Strategy Research Symposium at Boston University for their valuable comments and suggestions. All errors are our own.

1 Introduction

Content platforms (e.g., Youtube, TikTok) compete for user attention through the content they offer, and creators shape these offerings through decisions about where to participate, what to publish, and how much effort to devote to each platform. An important feature of content markets is that many creators maintain active presences on multiple platforms, a practice we call creator multihoming.[1] In a survey reported in January 2026, two-thirds of creators planned to expand to new platforms that year (Lauron, 2026). Platforms differ in their response: Twitch began allowing creators to stream simultaneously on rival platforms in 2023 (Sung, 2023), whereas Kick reduces partner-program payments by 50% for streams who also broadcast on other platforms (Iyer, 2026). These policies illustrate the trade-off between preserving exclusive participation and enabling broader exposure. In this paper, we examine a central question underlying this trade-off: How does adopting an additional platform affect a creator's audience, engagement, and content supply on the original platform?

The answer depends in part on how creators organize their content across platforms. Digital content can be distributed at low marginal replication cost, while audience attention and production effort remain scarce. Creators may share the same posts, adapt content to different audiences, or maintain distinct offerings. These choices shape content overlap and the allocation of production effort across platforms. Shared posts give users another place to consume the same content and may facilitate attention substitution (Aridor, 2025). Operating on an additional platform may also divert production effort from the original platform (Kapacinskaite and Mostajabi, 2024). Outside exposure can, however, generate attention spillovers back to the original platform (Krijestorac et al., 2020). Such spillovers may extend beyond consumption of the same post. When users recognize a creator across platforms, discovery through one post may generate interest in related work elsewhere, including content exclusive to the original platform. The additional platform therefore serves as both a potential discovery channel and a competing destination for consumption. Understanding these effects requires distinguishing creators' participation across platforms from their duplication of content.

These competing channels leave the net effect of multihoming uncertain and suggest that its consequences may depend on how creators allocate content. We therefore address four research questions. First (RQ1), how prevalent is creator multihoming, and how do creators allocate content across platforms? Second, our central question (RQ2) is: How does adopting an additional platform affect a creator's audience, engagement, and content supply on the original platform? Third, our mechanism question (RQ3) is: What role does cross-platform discovery play in shaping the original-platform audience response? We investigate whether attention attracted by content on the additional platform extends to the creator's broader portfolio on the original platform, including platform-specific content. Finally (RQ4), how do original-platform performance changes following adoption vary with creators' mix of shared and platform-specific content? For creators, these questions inform whether to expand and how to allocate content across platforms. For platforms, they help assess the trade-off between preserving content exclusivity and enabling cross-platform discovery.

Our empirical setting is Douyin (Chinese version TikTok) and Xiaohongshu (hereafter, RedNote), two

[1] Multihoming is common across platform markets. Drivers work for both Uber and Lyft, app developers publish through multiple app stores, and merchants list products on multiple marketplaces. The literature refers to such behavior more generally as complementor multihoming (Abolfathi, 2026; Li and Zhu, 2021).

leading social media platforms in China that compete for users and creators (Suciu, 2025).[2] Douyin emphasizes short videos, while RedNote centers on lifestyle and product-related content. Both support video posts, making content sharing feasible, while their different content environments give creators reasons to maintain platform-specific offerings. We assemble creator-level data on 15,000 commercially active Douyin creators and collect detailed posting histories for a stratified subsample of 2,000, including both creators active on RedNote and those with no identified RedNote account. The posting histories span seven years and are linked to daily follower counts and post-level engagement measures. These data allow us to characterize platform adoption and content allocation, and to compare changes in focal-platform outcomes over time.

A key measurement challenge is linking creator identities and individual posts across platforms. We match Douyin creators to their corresponding RedNote profiles using detailed account information and manually check the resulting matches. We then classify cross-posted content by comparing the semantic similarity of titles and other text metadata within each creator's cross-platform portfolio. These linked records allow us to separate shared posts from platform-specific offerings. They also enable us to examine whether RedNote-exclusive content is followed by Douyin audience growth and how engagement relates to the creator's shared and exclusive content. Finally, we collect Douyin videos published around adoption and use multimodal models to measure their technical quality and presentation. These measures help assess whether changes in observable video characteristics provide an alternative explanation for the audience response.

We begin with descriptive evidence on creator participation and content allocation (RQ1). Multihoming is prevalent and approximately 42% of the 15,000 creators in our sampling frame have identified accounts on both platforms. Most multihoming creators begin publishing on the second platform within a year of first posting on the initial platform. Creators generally maintain a similar topical focus across platforms while selectively cross-posting about 30% of their Douyin posts. The cross-posting share increases gradually with time since adoption, consistent with creators learning how to distribute content across platforms while retaining substantial platform-specific offerings. Creators thus expand across platforms in stages while continuing to adjust what they share and what they keep exclusive.

We next address RQ2 using a difference-in-differences(DiD) design that compares changes in Douyin outcomes between creators who adopt RedNote and those who remain Douyin-only. Our estimates capture the bundled effect of adoption, including audience responses and adoption-induced changes in content production and allocation, averaged over adopters' subsequent posting and content-sharing choices. The causal interpretation relies on adopters and comparison creators having parallel trends. An adoption-timing model finds no statistically significant relationship between entry and growth rates in followers, likes, or posting activity. Event studies show no statistically detectable differential pre-trends. These analyses reduce concerns that the estimated effects merely reflect pre-existing differences in creators' performance trajectories.

The baseline estimates indicate that RedNote adoption increases creators' Douyin followers by approximately 33% and engagement per post by 23%–26% across likes, comments, and shares. Creators also publish about three additional Douyin posts per month, with the net increase driven by greater cross-posted output. Their Douyin content becomes modestly more aligned with the content prevailing on RedNote. Together,

[2]Douyin had approximately 766 million daily active users in 2024 (Curry, 2025), while RedNote had approximately 300 million monthly active users in January 2025 (Lu and Miao, 2025).

these findings suggest that creators can expand their original-platform audiences and content supply while increasing content overlap with a competing platform.

To address RQ3, we examine evidence for cross-platform attention spillovers and their relationship to creators' content portfolios. A daily event study around the publication of highly engaging RedNote-exclusive posts shows increased Douyin follower growth shortly afterward, controlling for other posting activity on both platforms. Because these posts are unavailable on Douyin, this pattern is consistent with discovery extending beyond the same piece of content to the creator's original-platform presence. Comparing Douyin-exclusive posts published by the same creator in the same month, we also find that posts resembling the creator's RedNote offerings receive more engagement. These findings complement prior evidence that cross-platform distribution can increase consumption of the same video on the original platform (Krijestorac et al., 2020). They suggest that discovery can also extend to a creator's broader portfolio, including related content that remains exclusive to the original platform.

To answer RQ4, we examine how creators' original-platform engagement changes vary with their mix of shared and exclusive content. This allocation involves a tension: sharing content may support cross-platform discovery, while exclusive content preserves additional reasons to engage on Douyin. Greater duplication therefore need not generate greater gains. We nonparametrically estimate how changes in Douyin engagement vary with creators' cross-posting shares. Engagement changes are measured relative to those of creators who remain Douyin-only. The observed returns to cross-posting have a boundary. Fitted Douyin engagement gains follow an inverted-U pattern, with the largest fitted gains among creators who cross-post roughly half their Douyin output. This pattern is consistent with the value of combining cross-platform content sharing with a distinct original-platform offering.

Finally, we assess an alternative explanation for the audience gains: creators may apply knowledge gained on RedNote to improve their Douyin content through cross-platform learning (Polidoro and Yang, 2024). Using multimodal measures of technical quality and presentation, we compare changes around adoption with those among paired Douyin-only creators. We find no statistically detectable improvements. These findings provide little support for improvements in measured production quality as an explanation for the audience gains.

Our paper contributes to research on multihoming and platform competition in two respects. First, we provide new empirical evidence on creator multihoming by assembling a linked dataset that follows the same creators and their posts across Douyin and RedNote. These records allow us to examine second-platform adoption together with creators' subsequent content-sharing choices and changes in their original-platform audiences and content supply. Second, we broaden the analysis of cross-platform attention spillovers from an individual video to a creator's broader body of work. Prior research shows that distributing a video on another platform can increase consumption of that same video on the original platform (Krijestorac et al., 2020). We find that highly engaging posts available only on RedNote are followed by Douyin follower growth, and that Douyin-only posts receive more likes when they resemble the creator's RedNote content. These patterns suggest that discovery through one post can generate interest in related work on another platform. The platform implication is that participation on a rival platform can coexist with stronger original-platform performance and substantial content exclusivity. The nonmonotonic relationship between sharing intensity and

engagement further highlights a strategic tension between enabling cross-platform discovery and preserving distinct offerings.

2 Related Literature

Multihoming strategies. Homing decisions shape competition and value creation in two-sided markets, and early research on multihoming largely developed theoretical models of platform competition (Armstrong, 2006; Belleflamme and Peitz, 2019). Empirical research faces a basic measurement challenge: identifying the same complementors across competing platforms and tracking their activity and performance over time. Nevertheless, empirical studies have established important consequences for both platforms and complementors. At the platform level, evidence from video games shows that nonexclusive software broadens the scope of indirect network effects across competing consoles (Corts and Lederman, 2009). Greater application multihoming can also reduce console sales, although this effect diminishes as platforms mature or gain market share (Landsman and Stremersch, 2011). At the complementor level, multihoming creates opportunities and costs that affect product quality and performance. Adapting products to different platform architectures involves trade-offs in quality and specialization (Cennamo et al., 2018). Participation on another platform may also generate learning benefits. Polidoro and Yang (2024) find improvements in software complementors' home-platform performance following multihoming, with evidence consistent with transferring knowledge across platforms. In creator markets, Abolfathi (2026) examines how single-homing and multihoming YouTube creators respond to increased exposure of the rival platform, Twitch. Following this increase, multihoming creators reduce content creation and underperform relative to single-homers. This study highlights how existing homing strategies shape creators' responses to changes in platform competition.

Our paper contributes in two respects. First, we examine how creators' second-platform adoption affects original-platform audiences, engagement, and content supply, extending research on complementors' home-platform performance and creators' responses to competition (Abolfathi, 2026). Second, our linked creator and post records distinguish platform participation from content duplication. The nonmonotonic relationship between engagement gains and cross-posting intensity highlights a tension between broader distribution and preserving distinct offerings, bringing content allocation into the analysis of creator multihoming.

Attention spillover in digital platforms. Krijestorac et al. (2020) provide the closest comparison to our study, showing that distributing a video on another platform increases consumption of the same video on the original platform. Bairathi et al. (2026) document similar spillovers in music consumption between TikTok and Spotify. These studies examine spillovers involving the *same* content across platforms. Within Twitch, Zhao et al. (2023) examine a different margin, finding that category entry benefits other streamers already in that category. Discovery can also benefit distinct offerings from the same producer: new album releases increase sales of an artist's earlier albums, with evidence consistent with consumers discovering the artist (Hendricks and Sorensen, 2009). This portfolio perspective motivates examining whether discovery on one platform extends to a creator's other work elsewhere.

Table 1 places our study alongside these papers and marks the features that distinguish it. We extend

cross-platform attention-spillover research from consumption of the same item to discovery across a creator's broader portfolio. Douyin follower growth following highly engaging RedNote-exclusive posts, together with higher like counts on related Douyin-exclusive posts, is consistent with this mechanism. The nonmonotonic association between cross-posting intensity and engagement gains suggests limits to the net benefits of broader sharing. For platforms, these findings suggest that preserving content exclusivity need not require restricting creators' participation on competing platforms.

Table 1: **Related spillover studies and the contribution of this paper**

Study	Cross-platform	Own expansion	Audience discovery	Exclusive content	Multihoming context	Implication
Hendricks and Sorensen (2009)			✓	✓		A new release lifts sales of earlier albums.
Wang and Goldfarb (2017)	✓	✓	✓			New channels add demand where the brand is unknown.
Krijestorac et al. (2020)	✓		✓		✓	Posting a video elsewhere raises its views at home.
Zhao et al. (2023)			✓			New streamers raise viewership for category incumbents.
Polidoro and Yang (2024)	✓	✓			✓	Multihoming creators learn and perform better at home.
Bairathi et al. (2026)	✓		✓			TikTok exposure raises the song's Spotify streams.
This paper	✓	✓	✓	✓	✓	**Platforms can keep content exclusive without keeping creators exclusive.**

Notes: Cross-platform means the spillover runs between platforms or channels. Own expansion means the focal seller or creator opens the second channel. Audience discovery means audiences who find the seller or creator in one place seek them out in another. Exclusive content means the benefit reaches content not available on the second platform. Multihoming context means the seller or creator operates on both platforms at the same time.

Creator economy. Our paper also contributes to the growing literature on the creator economy. Prior work studies the effectiveness of influencer marketing (e.g., Leung et al., 2022; Yang et al., 2025) and how firms should choose creators to improve campaign performance (e.g., Doosti et al., 2026; Mallipeddi et al., 2022). From the creator perspective, research examines strategic choices, such as topic specialization (Gong, 2021), collaboration and co-creation (Yang et al., 2023), and joining influencer agencies (Hao and Li, 2026), and how they shape engagement. From the platform perspective, studies show how platform design, including advertising policy (Ren, 2024), monetization programs (El-Komboz et al., 2023), recommendation (Qian and

Jain, 2024) affects creator participation and content supply. Much of this literature abstracts from creators adopting multiple platforms. Our paper documents creator multihoming as a first-order strategic choice and has implications for many standard questions, such as which creators' brands should partner with and how platforms should design incentives and governance when creators multihome.

3 Empirical Context and Data

In this section, we first introduce the research context. We then describe our data collection procedures. Next, we explain how we process the unstructured data and construct key measures. Finally, we describe how we identify cross-posting.

3.1 Two Leading Content Platforms in China: Douyin and RedNote

This study uses data from two major social media platforms in China: Douyin and RedNote. Douyin, launched in 2016, rapidly became China's dominant short-video platform. The platform has reached approximately 766 million monthly active users by the end of 2024 (Curry, 2025). RedNote, launched in 2013, has similarly grown into one of China's most active lifestyle and social-commerce communities, attracting over 300 million monthly active users (Lu and Miao, 2025). Both platforms are discovery-driven social media that rely on algorithmic recommendations (Casner and Teh, 2025).

Although Douyin and RedNote originated with distinct positioning, their differentiation has narrowed over time. Douyin initially focused on entertainment-driven short- and medium-form videos, while RedNote emphasized product reviews or lifestyle content in photo-based posts. Over time, however, each platform has moved into the other's format. In August 2020, RedNote launched a video account feature that raised the maximum length of video posts to 15 minutes, as part of its effort to compete with Douyin for video content and creators.[3] Douyin, in turn, has since added a photo-and-text posting option and increasingly promotes such posts.[4] Both platforms now accommodate a similar range of video and photo formats, and this convergence in formats and functionality has made the two platforms increasingly substitutable from both the creator and user perspectives. For creators, greater format compatibility reduces repurposing costs and facilitates multihoming. For users, the growing similarity in content offerings increases the likelihood of cross-platform usage or even migration. As a result, platform competition intensifies for user attention and creator participation.

Douyin and RedNote provide a well-suited empirical setting to study creator multihoming. They constitute two of China's most influential discovery-driven social media platforms and among the most active venues for influencer sponsorships. Their increasingly overlapping functionality and content formats makes them natural complements for creators to maintain parallel presences. While creators can also multihome

[3]Source: https://hoozi.com.cn/news/121.html (In Chinese. Accessed on September 3rd 2026)

[4]Source:https://www.21jingji.com/article/20240725/herald/2d2857d2b59e6ab6547897d3be3e6f1e.html (In Chinese. Accessed on September 3rd 2026)

on other platforms such as Kuaishou and Bilibili, [5] such behavior is less prevalent. Kuaishou caters to a distinct audience segment and creator ecosystem, and Bilibili is oriented toward longer form and more professionally produced video content. Accordingly, an important scope of our analysis is that we study creator multihoming between Douyin and RedNote. Finally, we focus on the Chinese market because the market offers unusually rich and granular creator-level and content-level data, enabling precise measurement of multihoming behavior and performance outcomes.

3.2 Source of Data

Creator-level data. We construct a sample of Douyin content creators and then identify which of these creators also adopt RedNote (i.e., multihome). For this step, we randomly sample 15,000 creators from Douyin's creator-advertiser transaction marketplace. Joining this marketplace requires accounts to meet minimum thresholds for follower counts and cumulative video likes. These eligibility requirements screen out small or inactive users, ensuring that our sample consists of commercially active creators. In our data, sampled creators average approximately 2 million followers (see Table 2).

For each sampled Douyin creator, we collect detailed creator-level information, including nickname, age, gender, and IP location. We also obtain their Douyin-defined content categories, daily follower counts, and advertising quotation rates for brand sponsorships. For creators identified as multihoming, we also collect the same set of measures from RedNote.

Next, we identify which creators multihome. We define multihoming as a content creator starting to post content on both Douyin and RedNote. Operationally, the onset of multihoming is identified as the month in which a creator publishes their first post on the new platform, regardless of whether the content shared across platforms is identical or distinct. So our definition captures cross-platform participation rather than the specific content-level replication decision. Conceptually, multihoming is related to, but broader than, cross-posting. Cross-posting refers to the simultaneous sharing of the same content across multiple platforms, whereas our definition of multihoming includes a broader range of behaviors, including posting entirely different or partially adapted content across platforms.

We use a two-step matching algorithm to match Douyin creator accounts with their corresponding RedNote accounts. In the first stage, we perform automated matching based on creator nicknames and geographic locations based on IP address. Because these attributes are not unique and may generate multiple plausible matches, we conduct a second-stage manual verification. We cross-check profile images, biography descriptions, and primary content categories across platforms. Through this process, we identify 6,313 creators with accounts on both platforms, indicating that approximately 42% of creators in our sample engage in multihoming. Figure 1 provides an example of a creator who maintains profiles on both Douyin (left) and RedNote (right).

A potential concern is that creators may use different identifiers across platforms, which could lead us to miss some true cross-platform matches. In practice, multihoming creators typically maintain consistent

[5] Kuaishou is a short video platform launched in 2011 that primarily serves users in lower-tier Chinese cities. It has approximately 400 million daily active users as of 2024 and is known for its emphasis on social connections and grassroots content. Bilibili is a video sharing platform launched in 2009 that originated as a community for anime and gaming content. It has since expanded into a broader range of long-form and mid-form video categories.

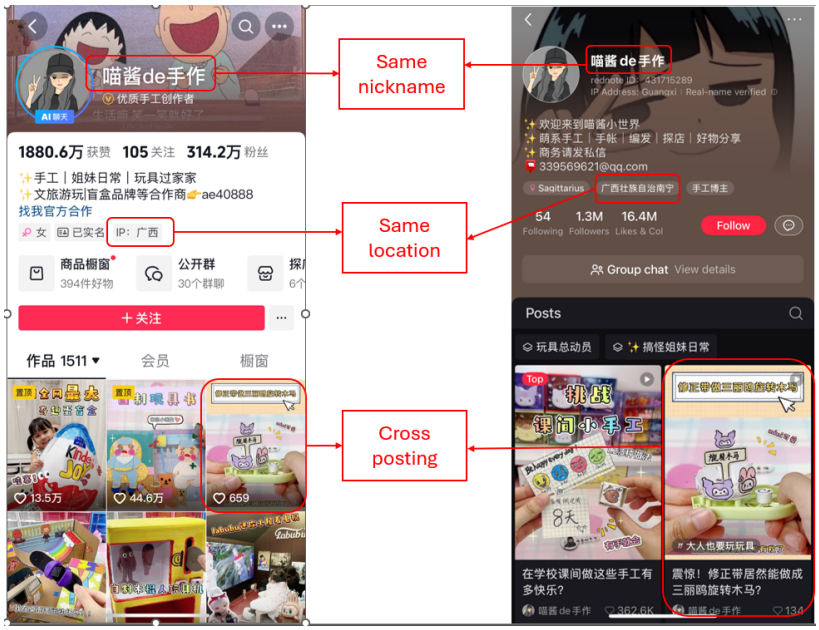


Figure 1: **The example of a multihoming content creator.**

Notes: The left screenshot shows the creator's Douyin account, while the right screenshot shows the corresponding RedNote account. The two profiles share the same nickname and geographic location, which allows us to identify this creator as a multihoming creator. In addition, the screenshots indicate that the creator has cross-posted at least one piece of content across the two platforms.

nicknames and content positioning across platforms to preserve creator brand recognition and facilitate audience expansion. In addition, creators may multihome across platforms outside our setting. Our analysis defines multihoming with respect to Douyin and RedNote and does not attempt to measure creators' full multihoming portfolio. Both concerns, if present, would mainly lead us to understate how prevalent multihoming is in the creator population, and our main findings on the impact of second-platform adoption on the focal platform are unlikely to be affected.

Post-level data. To ensure tractability, we collect complete posting histories for all 15,000 creators is infeasible, so we draw a stratified random subsample: 1,500 of the 6,313 multihoming creators and 500 of the creators for whom no RedNote account was found. For each sampled creator, we collect metadata for their complete posting history on the relevant platform or platforms through June 2024, including the posting time, title, description, hashtags, and engagement measures. These posting histories allow us to identify when each creator joined each platform and, for multihoming creators, when they began posting on the second platform. Based on the date of each creator's first post on each platform, 1,139 multihoming creators began posting on Douyin before joining RedNote, 326 began posting on RedNote before joining Douyin, and 35 began posting on both platforms on the same day.

A natural concern is that histories collected at a single point in time omit posts that creators later delete. Such deletions are a small fraction of all posts and concentrate on low-performing, outdated, or off-brand content (Branch, 2024). Because deleted posts are missing from both the pre- and the post-adoption periods, deletions bias our estimates only if creators prune posts differentially around the adoption date. Two features of our data limit this concern. First, our follower outcome comes from the daily follower counts in the creator-

level data and is unaffected by post deletion. Second, posting counts rise after adoption (Section 5.3, Table 3); selective deletion of post-adoption content would work against this pattern rather than produce it.

A second concern is that engagement counts collected at a single point in time may still be accumulating for recently published posts. Engagement on short-video platforms accrues quickly after publication, so this concern applies only to posts published shortly before collection. To verify this in our setting, we tracked 200 videos every two weeks for three months after publication. Their engagement counts are stable after the first two weeks, and the counts three months after publication differ from those two weeks after publication by about 1%. Because we collected the posting histories after the end of the sample window, the timing of collection does not affect the monthly engagement measures.

Panels A, B and C of Table 2 report summary statistics at the creator-month, creator and post levels for Douyin and RedNote, respectively. On average, creators post about eight times per month on each platform, indicating comparable levels of content production. Engagement per post and average follower size is higher on Douyin, consistent with its position as the more established platform with a larger user base.

Table 2: **Summary Statistics.**

	Mean	Std. Dev.	10th Percentile	90th Percentile
Panel A: creator-month level				
Number of posts on Douyin	8.78	12.06	0.00	19.00
Number of posts on RedNote	8.01	9.75	0.00	18.00
Panel B: creator level				
Follower size on Douyin	2,718,556	4,103,192	338,250	6,245,993
Follower size on RedNote	429,871	716,269	20,318	1,070,959
Panel C: post level				
Likes per post on Douyin	123,138	294,592	3,172	302,648
Comments per post on Douyin	5,777	20,097	99	12,635
Shares per post on Douyin	16,397	92,876	98	29,497
Likes per note on RedNote	12,213	44,381	122	25,281
Comments per note on RedNote	254	1,581	5	491
Collects per note on RedNote	2,764	13,928	17	4,860

Note: Statistics are computed over the 1,500 multi-homing creators and all months in the sample. Panel A reports creator-month observations. Panel B reports one observation per creator. Panel C reports individual posts: likes, comments and shares per Douyin post, and likes, comments, and collects (saves) per RedNote note.

Sample video-level data. The post-level data describe content through text metadata and engagement rather than through the videos themselves. To measure what creators actually produce, we also collect full Douyin video files for all creators in our study sample. Because downloading and processing videos from their entire posting histories would be time-consuming and computationally intensive, we restrict video collection to the two months before and the two months after each creator's event date. For treated creators, the event date is the date of their first RedNote post. Each control creator is assigned the adoption date of a treated creator as a pseudo-event date, so that treated and control creators are observed over comparable calendar windows. We collect every Douyin video published within these windows, yielding approximately 38,000 videos for the content-quality analysis. Section 3.3 describes the multimodal pipeline that converts

these videos into content and quality measures.

3.3 Data Processing and Measurement

Many of our analyses are content-based: whether a post is the same across platforms, what topic a video covers, and how much production effort it reflects. We construct these measures from two raw sources, the textual metadata attached to every post and, for the video-level sample, the video files themselves.

Text metadata. For the full post-level sample, we use titles, descriptions, and hashtags as a proxy for content. These fields typically summarize the core elements of a post and often carry explicit cues about topic and commercialization, and text is an established source of content measures in marketing and information systems research (Berger et al., 2020). We represent each post by a sentence embedding of its concatenated title, description, and hashtags. Sentence embeddings are produced by a BERT-type encoder fine-tuned so that texts with similar meaning map to nearby vectors (Reimers and Gurevych, 2019). Because the metadata are written in Chinese, we use bge-small-zh-v1.5, a Chinese sentence encoder from the BGE family trained on large-scale Chinese text-pair data (Xiao et al., 2024). We measure content similarity by the cosine similarity between embeddings. These similarities underlie the cross-posting identification in Section 3.4 and the cross-platform content similarity outcome in Section 5.2.

Video files. For the video-level sample, we measure content from the videos themselves, following work that extracts content features from video with machine learning (Yang et al., 2025) and uses LLM as standardized annotators (Li et al., 2024). Online Appendix Figure F.2 summarizes the pipeline. Each video is re-encoded into fixed-length analysis units so that measurement does not vary with video length. Technical quality is scored by FineVQ, a video-quality model calibrated to human ratings of user-generated content (Duan et al., 2025). Production and content attributes, such as lighting, speech clarity, staging, topic, and advertising, are coded by Qwen3-Omni, a multimodal model that reads sampled frames, audio, and on-screen text jointly (Xu et al., 2025), using fixed prompts with closed answer categories and deterministic decoding. Speech is transcribed with SenseVoice (An et al., 2024) and the transcripts coded with DeepSeek-V3 (DeepSeek-AI, 2024) for attributes best judged from the words themselves, such as scripting and verbal hooks. Online Appendix A describes the technical details.

3.4 Cross-Posting Identification

Among 1,500 content creators in our dataset that operate accounts on both Douyin and RedNote, the content they post is not always identical across platforms. In other words, not all posts are cross-posted. To identify cross-posting, we compare content across the two platforms using the text-embedding representations described above. Specifically, for each sampled multihoming creator, we perform pairwise comparisons between Douyin and RedNote posts that are published within the same calendar month starting from the month they establish multihoming. For each Douyin post t and RedNote post r created by the creator i

within the same calendar month m, we compute the cosine similarity as:

$$\text{CosSim}_{i,m,t,r} = \frac{\left(\mathbf{e}^{T}_{i,m,t}\right)^{\top} \mathbf{e}^{R}_{i,m,r}}{\left\|\mathbf{e}^{T}_{i,m,t}\right\|_2 \left\|\mathbf{e}^{R}_{i,m,r}\right\|_2},$$

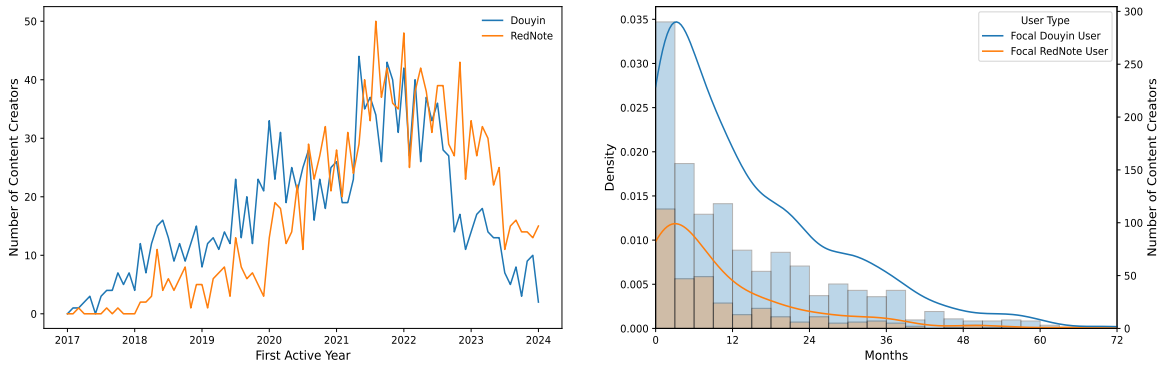
where $\mathbf{e}^{T}_{i,m,t}$ and $\mathbf{e}^{R}_{i,m,r}$ denote the K-dimensional embedding vectors for the Douyin post t and RedNote post r, respectively. For each post, we then identify the best-matching post on the other platform, if the maximum cosine similarity exceeds a threshold. Because titles, descriptions, and hashtags can differ across platforms even for the same underlying video (see Online Appendix Table E.1 for examples), we experiment with several similarity cutoffs (e.g., 0.75, 0.8, 0.85). For each threshold, we draw a subsample of matched pairs and manually verify whether they correspond to the same video across platforms. This validation exercise shows that 0.8 best separates genuinely cross-posted content from posts that are only topically related. Using this rule, we create a binary indicator equal to one if a post has a matched counterpart on the other platform (i.e., is cross-posted) and zero otherwise.

4 Facts About Creator Multihoming

More than four in ten creators multihome. To establish the empirical relevance of our setting, we first document the prevalence of creator multihoming. As described in Section 3.2, we randomly sample 15,000 Douyin creators and match them to RedNote accounts using nicknames and location, followed by manual verification. This procedure identifies 6,313 creators with accounts on both platforms, so about 42% of sampled Douyin creators multihome. As discussed there, this estimate is a lower bound, because creators who use different identifiers across platforms, or who multihome on platforms other than RedNote, are not counted. Multihoming is thus common but far from universal, and the creators who expand are not a random subset. A logit model of multihoming status (Table E.2 in the Online Appendix) shows that female creators, creators based in first-tier cities, and, most strongly, creators with higher quoted advertising prices are more likely to maintain a RedNote presence, consistent with RedNote's female-skewed, urban, commerce-oriented positioning. This selectivity indicates that multihoming is a deliberate decision, and it motivates the identification strategy in Section 5.1, which absorbs time-invariant creator characteristics through creator fixed effects and addresses selection on time-varying performance directly.

Most creators multihome within a year of entering the focal platform. We examine the timing of creators' platform adoption and show that multihoming typically occurs within one year of initial platform entry. Figure 2a displays the dates on which the 1,500 sampled multihoming creators first posted on Douyin and RedNote. Figure 2b presents the elapsed time between each creator's first posts on the two platforms. Most creators begin posting on the second platform within one year of their first post on the initial platform. The pattern points to a staged rather than simultaneous expansion: creators first establish themselves on one platform and add the second after some months. Prior research shows that agents often postpone expansion until they have accumulated sufficient experience (Jovanovic, 1982). In our context, creators may initially rely

on the focal platform to learn platform-specific norms, refine their content strategy, and build an audience before expanding to an additional platform.



(a) Joining Time (b) Time Difference

Figure 2: **Joining Time of content creators on Douyin and RedNote.**

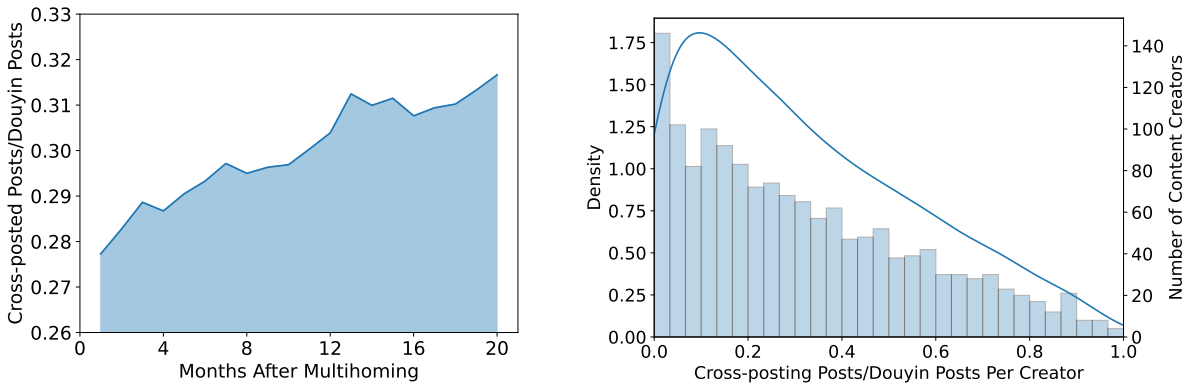
Notes: Panel (a) plots the number of new creators joining each platform monthly. Panel (b) plots the distribution of the time elapsed between a creator's first posting on their first and second platform. For focal Douyin users, Douyin is the first platform joined; for focal RedNote users, RedNote is the first platform joined. The left y-axis reports density, and the right y-axis reports the number of creators.

Creators keep a consistent topic focus while cross-posting selectively We examine how creators adjust their content strategies when operating across Douyin and RedNote. We show that creators maintain similar content topic across platforms, selectively cross-posting content, indicating that cross-posting is a deliberate and selective strategy rather than mechanical duplication.

We first categorize content creators using platform-specific user tags and compare content types across Douyin and RedNote. Because category labels are not defined identically, we manually calibrate tags across the two platforms (see Online Appendix Figure F.1 for details). Overall, approximately 68% share at least one common tag across platforms, while the remaining creators exhibit minor differences, indicating substantial overlap in content focus.

We next use post-level data to examine how creators allocate content across platforms after multihoming. Using the technique for cross-posting identification introduced in Section 3.4, we find that, on average, 30% of monthly Douyin posts are also shared on RedNote, while approximately 39% of RedNote posts are duplicated on Douyin, indicating a moderate reliance on content replication. We define the cross-posting ratio as the number of cross-posted posts divided by the total number of Douyin posts for each creator in a given month. Figure 3a plots the average cross-posting ratio by event time after multihoming. The figure shows that the cross-posting ratio increases gradually with time since multihoming, suggesting that creators rely more heavily on content replication as they gain experience operating on multiple platforms. While a small subset of creators fully replicate content across platforms, most creators only partially cross-post and continue to produce largely distinct content for each platform. Figure 3b presents the density of creators' average cross-posting ratios, highlighting this variation.

A natural question is whether creators selectively choose which posts to share across platforms. To provide suggestive evidence, we compare posts shared across platforms with posts shared on only one platform,



(a) (b)

Figure 3: **Monthly Changes of Cross-posting Ratio and Douyin Ratio.**

Notes: Panels (a) plot changes in the average cross-posting ratio. In Panel (b), the left y-axis reports density and the right y-axis reports the number of creators.

produced by the same creator within the same month. We find that posts selected for distribution across platforms subsequently receive approximately 33% more likes. On its own, this difference does not reveal when creators make the sharing decision. It shows, however, that the content they choose to distribute across platforms tends to perform better than the content they keep on only one platform.

Finally, to shed light on the timing of this decision, we examine when matched content is published across platforms. Among the 139,314 shared posts we identify, 76.1% are published on both platforms on the same day, and 95.6% appear within five days (see Online Appendix Figure F.3). The prevalence of same-day publication suggests that creators typically decide in advance which content to distribute across both platforms, rather than making this decision after observing its performance on one platform. These results together answer our RQ1.

5 Effects of Multihoming

We examine how adopting RedNote affects two dimensions of creators' activity on Douyin. The first is audience performance, measured by follower counts and engagement per post. The second is content supply, measured by posting volume and content positioning. These outcomes allow us to examine both how creators' audiences respond to expansion and how creators adjust their production and positioning.

5.1 Research Design

We exploit creators' staggered adoption of an additional platform to estimate its effects on audience outcomes and content supply on the focal platform. Our primary estimand is the average treatment effect on the treated (ATT), the average difference between adopters' outcomes following expansion and the outcomes they would have experienced without multihoming. Multihoming is a strategic decision, and a DiD design does not require treatment assignment or treatment timing to be as good as random (de Chaisemartin and D'Haultfœuille, 2025). The key identifying assumption is parallel trends, when absent multihoming, adopters

and comparison creators would have followed similar outcome trajectories on the focal platform. We discuss the identifying assumptions and their plausibility below.

Our treatment captures the extensive margin of multihoming: a creator's decision to begin publishing on an additional platform. Following adoption, creators choose how much to publish, which content to cross-post, and what to keep platform-specific. The ATT captures the total effect of initiating this additional presence, including audience responses and creators' behavioral adjustments induced by adoption. This bundled effect is relevant to focal platforms assessing the consequences of creators' expansion for audience activity and content supply. It averages over adopters' subsequent content allocation choices, allowing participation intensity and content overlap to differ across creators. Section 6.3 complements the adoption analysis by examining how focal-platform engagement gains vary with creators' chosen cross-posting shares. These analyses address the overall consequences of expansion and how the gains vary with subsequent content allocation.

We focus on creators whose initial presence is on Douyin, the more established platform and the first platform for most creators in our sample, and define treatment as the subsequent adoption of RedNote. All outcomes are measured on Douyin. We aggregate the data to the creator-month level and restrict the analysis window to July 2021 through June 2024, a period in which RedNote adoption became substantial (Ch, 2026). Treated creators adopt RedNote during this window; control creators remain Douyin-only throughout it. Our baseline specification is the two-way fixed-effects (TWFE) model:

$$Y_{it} = \beta D_{it} + \alpha_i + \delta_t + \epsilon_{it}, \tag{1}$$

where Y_{it} denotes the outcome for creator i in month t, as defined in Section 5.2. The treatment indicator D_{it} equals one from the month of a creator's first RedNote post onward and zero otherwise; it remains zero throughout the window for creators in the control group. Creator fixed effects α_i absorb persistent differences across creators, while month fixed effects δ_t absorb common shocks, such as platform-wide policy changes. We estimate Equation (1) by ordinary least squares and cluster standard errors at the creator level.

We complement the baseline estimates with dynamic event studies to examine how outcomes evolve around adoption. We replace D_{it} with indicators for months relative to adoption, omitting the month immediately before adoption as the reference period. With staggered adoption, conventional TWFE estimates may not recover the ATT when treatment effects vary across cohorts or over time (Goodman-Bacon, 2021). We therefore also use the interaction-weighted event-study estimator of Sun and Abraham (2021), which allows treatment effects to vary across adoption cohorts and over time.

Identification discussion. Identification requires that, absent RedNote adoption, adopters and comparison creators would have followed parallel trends in focal-platform outcomes. Strategic adoption does not itself invalidate the DiD design. The concern is whether adoption timing is related to changes in performance that would have occurred even without expansion. In particular, time-varying unobserved factors may influence both adoption and focal-platform performance. A creator's business may be growing along a distinct trajectory, making expansion more attractive while also increasing Douyin activity. Alternatively, a marketing campaign may coincide with RedNote adoption and independently improve Douyin outcomes. Such changes

differ from production adjustments induced by adoption, which are part of the total effect we seek to estimate.

We assess this concern using evidence on adoption timing, pre-adoption outcome dynamics, and sensitivity to creator-specific trends. First, a discrete-time hazard model relates adoption timing to the levels and growth rates (trends) of followers, posting activity, and likes (Online Appendix B). Outcome levels predict adoption, whereas the estimated relationships with growth rates are not statistically significant. Differences in outcome levels are compatible with DiD; creator fixed effects absorb persistent differences in these levels. Second, the event-study estimates (Figure 4) show no statistically detectable pre-trends. Finally, we allow for creator-specific linear time trends in a sensitivity analysis (Online Appendix C.3). This specification permits creators to follow different linear trajectories in their focal-platform outcomes. The results are largely consistent.

5.2 Outcome Variables

We examine both audience responses and creators' behavioral responses to multihoming, using four sets of outcomes measured on Douyin. Audience size and engagement capture whether expansion strengthens the creator's existing reach and the attention their content receives. Content production and positioning capture whether creators change how much they publish and how closely their content aligns with the additional platform.

Audience size. Expansion may introduce creators to new audiences. It also gives users another platform on which to follow them. We therefore examine whether adopting RedNote expands or weakens the creator's established Douyin following. We measure audience size as $\log(\text{Followers}_{it})$, where Followers_{it} is creator i's number of Douyin followers at the end of month t. Follower counts are widely used to measure social media influence (Mallipeddi et al., 2022; Qiu and Kumar, 2017).

Engagement. A larger following does not necessarily translate into greater interaction with content. We therefore examine engagement per post to assess whether changes in audience size are accompanied by changes in audience activity. We distinguish likes, which reflect audience approval; comments, which capture participation in discussion (Lee et al., 2018); and shares, which capture redistribution and social transmission (Berger and Milkman, 2012). For each creator–month, Likes_{it}, Comments_{it}, and Shares_{it} are the average observed numbers of likes, comments, and shares across Douyin posts published by creator i in month t. We use $\log(\text{Likes}_{it})$, $\log(\text{Comments}_{it})$, and $\log(\text{Shares}_{it})$. Engagement is measured at data collection. As documented in Section 3.2, engagement typically stabilizes within one to two weeks of publication, limiting the extent to which older posts have higher measured engagement simply because they have had more time to accumulate likes, comments, and shares.

Content production. Operating on an additional platform may divert production effort from Douyin, while opportunities to reuse content or reach a larger audience may encourage greater output. We examine the resulting change in focal-platform content supply using $\text{Number of posts}_{it}$, the number of Douyin posts creator

i publishes in month t. Posting volume captures production intensity and provides an indirect measure of effort on the platform (Cao et al., 2024).

Content positioning. Expansion may also change what creators produce. Creators may adapt their Douyin content toward the content prevailing on RedNote or maintain distinct offerings across platforms. We measure content positioning as Cross-platform content similarity$_{it} = \cos\left(\bar{\mathbf{e}}_{it}^{T}, \bar{\mathbf{e}}_{t}^{R}\right)$, where $\bar{\mathbf{e}}_{it}^{T}$ is the average caption embedding of creator i's Douyin posts in month t, and $\bar{\mathbf{e}}_{t}^{R}$ is the average caption embedding across all posts published that month by RedNote-first creators in our sample. A higher value indicates greater alignment between the creator's Douyin content and the content prevailing on RedNote.

5.3 Main Results

Adopting RedNote strengthens creators' existing presence on Douyin. Table 3 reports a follower coefficient of 0.285 log points, implying approximately 33% more followers relative to the counterfactual without adoption. Engagement per post also increases: likes rise by approximately 25.9%, comments by 24.2%, and shares by 22.5%. Expansion therefore enlarges creators' established audiences while increasing interaction with their content. Figure 4(a)–(d) reveals how these gains develop. Engagement rises around adoption and remains broadly elevated, while follower gains accumulate steadily over subsequent months. These patterns suggest an expansion of the creator's focal-platform presence that extends beyond the initial entry period. The TWFE and interaction-weighted estimates (see Online Appendix C.1) yield similar trajectories, although share estimates are less precise at some horizons.

Table 3: **Effects of Multihoming on Focal-Platform Audience Outcomes and Content Supply.**

	Audience size	**Engagement**			**Content production**	**Content positioning**
	(1) Log followers	(2) Log likes	(3) Log comments	(4) Log shares	(5) Number of posts	(6) Cross-platform content similarity
Multihoming	0.285***	0.230***	0.217***	0.203***	3.029***	0.0134***
	(0.036)	(0.050)	(0.055)	(0.059)	(1.128)	(0.0029)
Constant	13.64***	10.67***	7.572***	8.362***	8.905***	0.7486***
	(0.016)	(0.023)	(0.025)	(0.027)	(0.483)	(0.0013)
Creator fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,645	26,815	26,815	26,815	30,072	26,745
R-squared	0.887	0.628	0.562	0.614	0.568	0.544

Notes: All outcomes are measured on Douyin at the creator-month level. Standard errors clustered at the creator level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Creators' content offerings also become more closely connected across platforms. Monthly Douyin output increases by 3.029 posts, roughly one third of the sample mean of about nine posts (Table 2). The

decomposition by post type shows that this increase is concentrated in shared content. Cross-posted output increases by 3.895 posts per month, while the estimated change in Douyin-exclusive output is −0.865 posts and is not statistically significant (Online Appendix Table E.3). Cross-platform content similarity also rises by 0.0134, indicating that creators' Douyin offerings become more closely aligned with the content prevailing on RedNote. Thus, expansion is accompanied by both more content shared across platforms and closer alignment in content positioning. Figure 4(e)–(f) shows that the increases in posting volume and content similarity emerge around adoption and persist over the displayed post-adoption period.

Taken together, these results show that creators can strengthen their focal-platform audience outcomes while offering more content that overlaps with a competing platform. Where do these audience gains come from? RedNote may serve as a discovery channel, introducing users to creators whose work they subsequently seek out on Douyin. Creators may also learn from their experience on RedNote and improve their Douyin content. Section 6 examines evidence on these explanations, using the timing of audience responses, the relationship between engagement and content allocation, and changes in measured content quality to assess the sources of the gains.

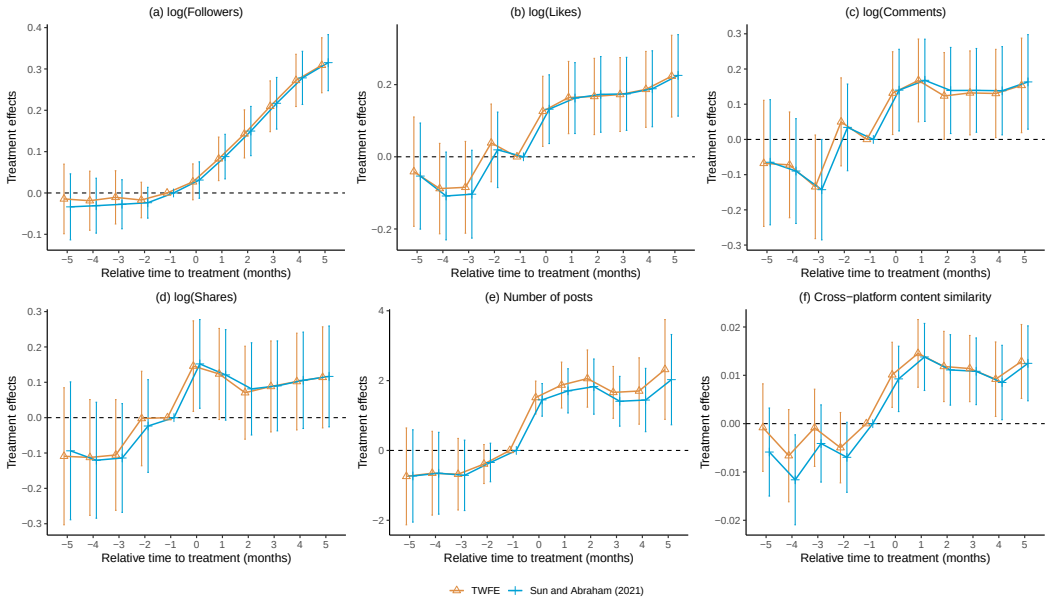


Figure 4: **Dynamic DiD Results.**

Notes: Estimated effects of RedNote adoption on each outcome by month relative to adoption, with the month before adoption ($k = -1$) as the reference period. Orange triangles are TWFE event-study estimates; blue crosses are the interaction-weighted estimates of Sun and Abraham (2021). Vertical bars are 95% confidence intervals with standard errors clustered at the creator level.

5.4 Robustness Checks

We conduct three additional robustness checks to validate our main findings. Table 4 summarizes these analyses. Online Appendix C provides the specifications and full results.

Table 4: **Additional Robustness Checks.**

Concern	Analysis	Finding
Spillovers onto comparison creators	Relate untreated creators' outcomes to adoption among same-category peers	No significant peer association
Sensitivity to treatment timing	Date treatment from the onset of cross-posting	Similar positive estimates
Dependence on the focal platform	Examine RedNote-first creators' subsequent adoption of Douyin	Similar patterns on RedNote

Spillovers onto comparison creators. The DiD estimates require that one creator's adoption does not affect the Douyin outcomes of creators who have not adopted, the stable unit treatment value assumption (SUTVA). This would fail if, for example, a peer's expansion to RedNote diverted attention away from or toward non-adopters in the same content niche, contaminating the comparison group. We assess this concern by regressing untreated creators' follower counts and likes on the share of their same-category peers who have adopted RedNote (Online Appendix C.2). The estimated associations are statistically indistinguishable from zero for both outcomes. This diagnostic provides no clear evidence of within-category spillovers onto comparison creators.

Alternative treatment timing. Our baseline definition dates treatment from the month a creator first becomes active on RedNote, regardless of what they post there. As a complementary timing check, we date treatment from the onset of cross-posting and re-estimate Equation (1) (Online Appendix C.4). The estimates remain positive, statistically significant, and similar in magnitude. The positive pattern therefore persists when outcomes are aligned around the start of content sharing rather than the first RedNote post.

Alternative focal platform. We also examine whether the findings extend to expansion in the opposite direction. Using creators who first became active on RedNote and subsequently adopted Douyin, we estimate changes in RedNote posting volume, likes per post, and content similarity to Douyin (Online Appendix C.5). Following Douyin adoption, creators publish more on RedNote, receive greater engagement per post, and produce content more closely aligned with Douyin. These results suggest that the combination of increased focal-platform activity and cross-platform content adaptation is not confined to creators expanding from Douyin to RedNote.

6 Content Portfolios and Cross-Platform Attention

Section 5 shows that multihoming increases creators' audience size, engagement, and content supply on the focal platform. We now examine what drives these gains and how they relate to creators' content choices (RQ3). Section 6.1 presents evidence that attention gained on RedNote can translate into audience growth on Douyin. We then ask how these gains relate to the content creators offer across the two platforms? Section 6.2 examines engagement across shared and exclusive content, while Section 6.3 considers how overall gains vary

with the cross-posting mix. Finally, Section 6.4 considers an alternative explanation: whether learning on the new platform improves focal-platform content quality.

6.1 Cross-Platform Attention Spillovers

The positive effects documented in Section 5 may arise through cross-platform attention spillovers or learning. Exposure on RedNote may introduce creators to users who subsequently follow or engage with them on Douyin (Krijestorac et al., 2020). Alternatively, creators may apply knowledge gained on RedNote to improve their Douyin content (Polidoro and Yang, 2024). We first examine the attention-spillover channel. Users who discover a creator on RedNote may seek out more of that creator's work on Douyin. Creators can also facilitate this process by listing their Douyin accounts on their RedNote profiles, reducing search costs and drawing attention to their presence on the other platform (see Online Appendix Figure F.4 for an example). This channel yields a testable prediction, attention attracted by content on RedNote should be followed by audience growth on Douyin.

We examine whether highly engaging RedNote-exclusive posts are followed by increases in the creator's Douyin follower growth. Approximately 66% of RedNote posts in our sample are classified as exclusive to that platform. These posts provide a useful setting because subsequent Douyin audience growth cannot be attributed to users encountering the same post on Douyin. If attention on RedNote spills back to Douyin, follower growth should increase shortly after the RedNote post is published.

We classify a RedNote post as viral if its observed likes fall in the top decile of the creator's RedNote posts published in the same calendar year. The event date is the post's publication date. We focus on days when a creator publishes at least one viral post on RedNote and none of their RedNote posts from that day are also shared on Douyin. To isolate the event, we exclude days with another viral RedNote post in the preceding or following seven days. Let E_{id} indicate such an isolated viral RedNote-exclusive day for creator i. We estimate

$$100\,\Delta\log(\text{Followers}_{id}) = \sum_{k=-5}^{5}\beta_k E_{i,d-k} + \gamma'\mathbf{X}_{id} + \alpha_{iw(d)} + \delta_d + \varepsilon_{id}, \tag{2}$$

where the dependent variable is daily Douyin follower growth expressed in percentage units. Creator-by-week fixed effects $\alpha_{iw(d)}$ absorb each creator's level of follower growth in the calendar week $w(d)$, including any momentum a creator has in the weeks around a viral note, while calendar-day fixed effects δ_d absorb common daily shocks. The coefficients β_k trace differences in follower growth from five days before to five days after an event, conditional on the included controls. The vector $\mathbf{X}_{id}$ includes five leads and lags of indicators for other posting activity on both platforms. We include indicators for non-isolated viral days and viral days involving cross-posted content to control for other bursts of attention on RedNote. We include indicators for ordinary RedNote posting days to control for routine activity on the additional platform, and indicators for the creator's Douyin posting and viral days to control for follower growth associated with activity on Douyin itself. Together, these controls help distinguish the pattern around isolated RedNote-exclusive events from that associated with other observed activity on the two platforms.

Figure 5 shows a concentrated increase in Douyin follower growth following these events. The five pre-event coefficients are small and jointly statistically indistinguishable from zero, providing no evidence of a

preceding increase in Douyin follower growth. Growth peaks during the following two days at approximately 0.10 percentage points above the baseline daily growth rate and then declines toward zero. This pattern is consistent with RedNote serving as a discovery channel for the creator's presence on Douyin. Discovery may extend beyond the content initially encountered. The RedNote post is unavailable on Douyin, yet the creator's Douyin audience grows shortly afterward. A matched-control version that compares each event creator with similar multihoming creators on the same calendar days yields a response of the same timing and size, which rules out topic-wide surges in attention as the source (see Online Appendix D). Section 6.2 next examines how engagement varies across the creator's shared and exclusive content.

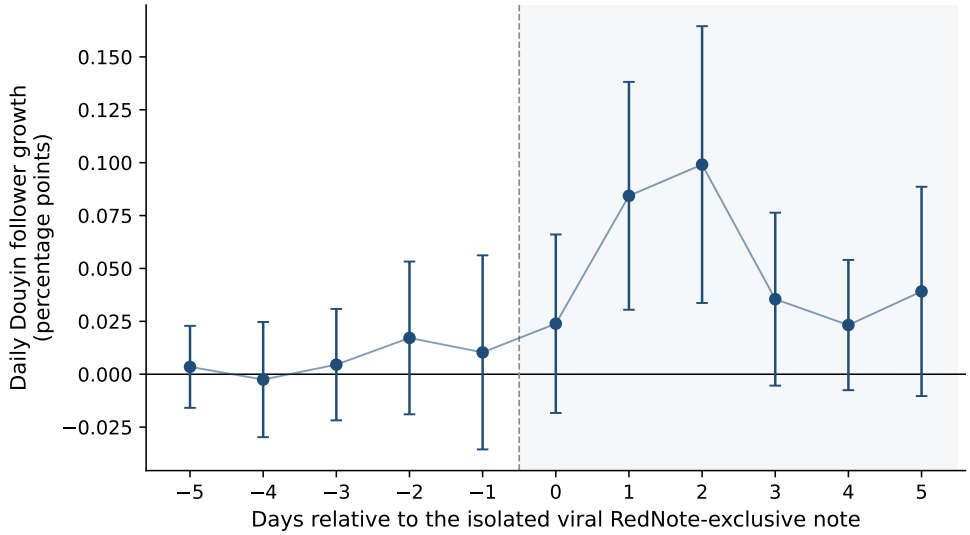


Figure 5: New Douyin followers arrive right after RedNote-exclusive content goes viral.
Notes: Coefficients β_k from Equation (2) with 95% confidence intervals based on standard errors clustered at the creator level. Growth is in % per day relative to the creator's ordinary post-adoption days.

6.2 Content Allocation and the Distribution of Engagement

Section 6.1 provides evidence consistent with attention on RedNote contributing to audience growth on Douyin. We now examine how engagement relates to the content connecting a creator's two platform presences. Krijestorac et al. (2020) show that introducing a video on an additional platform can increase consumption of that same video on the original platform. Creator multihoming also connects distinct pieces of content through a common creator. Users who discover a creator on RedNote may seek out additional content on Douyin that matches their interests, even when that content is unavailable on RedNote. This possibility motivates our analysis of whether Douyin-exclusive posts receive more engagement when they resemble the creator's RedNote content.

For each Douyin-exclusive post, we measure similarity to the creator's RedNote content as the maximum cosine similarity between its caption embedding and the embeddings of the same creator's RedNote notes published in the same calendar month, requiring at least three usable reference notes. This measure captures how closely a Douyin-exclusive post resembles what the creator offers on RedNote. Because the same-

month reference set may include notes published after the focal Douyin post, we also construct an alternative measure using only RedNote notes published during the 90 days preceding that post.

For Douyin-exclusive post p published by creator i in calendar month m, we estimate

$$\log\left(\text{Douyin Likes}_{ip}\right) = \beta \frac{\text{Similarity to Own RedNote Content}_{ip}}{0.1} + \gamma' \mathbf{X}_{ip} + \alpha_{im} + \varepsilon_{ip}, \tag{3}$$

where Douyin Likes_{ip} denotes the post's cumulative likes on Douyin. We scale similarity so that β represents the difference in log likes associated with a 0.1-point increase in similarity. Creator-by-publication-month fixed effects, α_{im}, ensure that the comparisons are between posts published by the same creator in the same month, absorbing conditions common to that creator's posts during the month. The vector $\mathbf{X}_{ip}$ includes publication-timing controls to account for differences in engagement associated with when a post is published. We also control for each post's similarity to the creator's other Douyin posts, excluding the focal post from this calculation. Posts that resemble the creator's RedNote content may also align more closely with the creator's usual Douyin offering and therefore appeal more to the existing Douyin audience. This control helps distinguish that relationship from the association with similarity to the creator's RedNote content.

Table 5: **Similarity to Creators' RedNote Content and Douyin Engagement.**

	Douyin-exclusive posts	
	(1)	(2)
RedNote reference posts	Same calendar month	Prior 90 days
Dependent variable: Douyin engagement, log(Likes)		
Similarity to creator's RedNote content	0.1384^{***}	0.0552^{**}
(per 0.1-point increase)	(0.019)	(0.027)
Creator-by-month fixed effects	Yes	Yes
Publication-timing controls	Yes	Yes
Similarity to creator's other Douyin posts	Yes	Yes
Observations	50,420	50,514

Notes: Both columns include only Douyin-exclusive posts. Similarity to the creator's RedNote content is the maximum cosine similarity between the focal Douyin post's caption embedding and the embeddings of the same creator's RedNote notes published in the same calendar month (column 1) or during the 90 days preceding the focal post (column 2). Coefficients represent a 0.1-point increase in similarity. Standard errors clustered at the creator level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 5 shows that Douyin-exclusive posts receive more likes when they resemble the creator's RedNote content. Using the same-month reference set, a 0.1-point increase in similarity is associated with 0.138 log points more likes, approximately 15% (column 1). Using the prior-90-day reference set, the association remains positive and statistically significant, at 0.055 log points, approximately 6% (column 2). Because the eligible sample also differs, the attenuation cannot be attributed to the reference window alone.

These findings extend the analysis of cross-platform engagement beyond shared content to related content that remains exclusive to the focal platform. Even among Douyin-exclusive posts, engagement is higher when the content resembles what the creator offers on RedNote. Together with Section 6.1, this pattern is consistent with users discovering a creator on RedNote and seeking out additional, related content on

Douyin. Under this interpretation, content on the additional platform introduces users to the creator, while related exclusive content gives them a reason to consume the creator's work on the focal platform. Cross-platform discovery and content exclusivity can therefore be complementary. This shifts attention from the distribution of an individual piece of content to the composition of a creator's portfolio across platforms. The next subsection examines how overall gains vary with creators' chosen mix of shared and platform-specific content.

6.3 Cross-Posting Intensity and Performance

Section 6.2 shows that the relationship between cross-platform content and engagement extends to related posts that remain exclusive to Douyin. This finding suggests that shared and exclusive content may play complementary roles in a creator's portfolio. We now move from comparisons across individual posts to comparisons across creators: how do focal-platform engagement gains vary with the share of content creators cross-post? Higher engagement on shared content does not necessarily mean that creators benefit from sharing a larger proportion of their output. Shared content may help attract attention across platforms, while Douyin-exclusive content may both appeal to the creator's existing Douyin audience and give users arriving from RedNote additional material to explore. The gains from multihoming may therefore depend on how creators combine shared content with content that preserves the distinct appeal of their focal-platform offering.

We focus on creators who begin publishing on RedNote during the study period, whom we refer to as adopters. For each adopter, we define the cross-posted share, s_i, as the number of Douyin posts cross-posted to RedNote divided by the total number of Douyin posts published during the first year after adoption:

$$s_i = \frac{\text{Number of first-year Douyin posts also posted on RedNote}}{\text{Total number of first-year Douyin posts}}.$$

$s_i = 0$ means none of the creator's Douyin posts are shared on RedNote, while $s_i = 1$ means all are shared. We construct the outcome using a DiD comparison. For each adopter, we calculate the change in average monthly log(Likes) on Douyin between the six months before adoption and the six months beginning with adoption. We then subtract the average change over the same calendar months among creators who remain Douyin-only throughout the study period, whom we refer to as never-adopters. The resulting outcome, $\tilde{y}_i$, measures the adopter's Douyin engagement change relative to this comparison group.

To allow this benchmark-adjusted change to vary nonlinearly with the cross-posted share, we estimate the following model:

$$\tilde{y}_i = f(s_i) + \beta T_i + \gamma_{g_i} + \varepsilon_i, \tag{4}$$

where $f(s_i)$ captures the relationship between the cross-posted share and the engagement change. T_i is log of the creator's total Douyin posts during the first year after adoption, and γ_{g_i} denotes adoption-month fixed effects. These controls account for differences in posting volume and conditions common to creators who adopt in the same month. An adjusted specification additionally controls for pre-adoption followers, likes, sponsored share, topic concentration, and content similarity to RedNote. Although the outcome follows a DiD construction, the estimated curve describes how engagement changes vary with creators' chosen cross-

posting shares; it does not identify the causal effect of changing those shares.

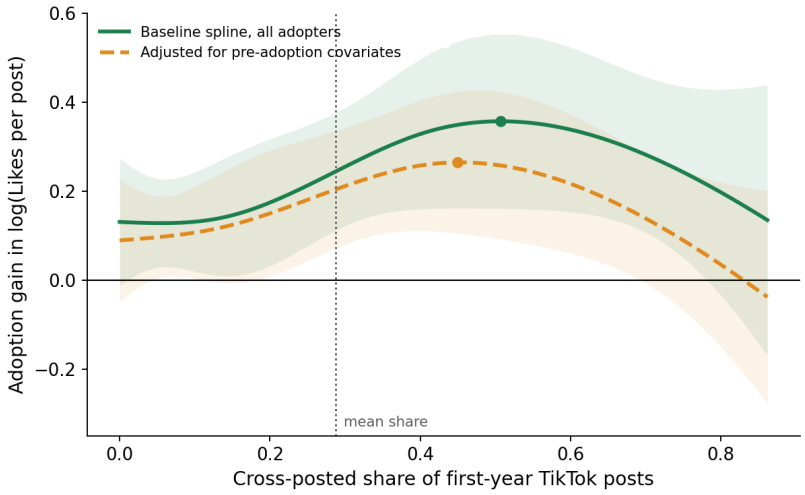


Figure 6: Cross-Posting Share and Douyin Engagement Gains.

Notes: The horizontal axis measures the share of Douyin posts also published on RedNote during the first year after adoption. The vertical axis shows the fitted change in average monthly log(Likes) from the six months before adoption to the six months beginning with adoption, minus the corresponding change among creators who remain Douyin-only. Zero therefore indicates the same engagement change as this comparison group. The solid curve reports the baseline specification in Equation (4); the dashed curve additionally controls for pre-adoption creator characteristics. Shaded bands indicate 95% confidence intervals.

Figure 6 shows an inverted-U relationship between the cross-posted share and Douyin engagement gains. The baseline curve peaks when creators cross-post approximately 51% of their Douyin content; controlling for pre-adoption creator characteristics yields a similar pattern, with a peak at 45%. Thus, the largest fitted gains occur among creators who share roughly half of their Douyin content while keeping the remainder exclusive. These peaks describe the observed relationship rather than a causal optimum. This pattern complements the preceding subsections. Section 6.1 provides evidence consistent with audience discovery across platforms, while Section 6.2 shows that higher engagement extends to related content that remains exclusive to Douyin. The inverted-U relationship is consistent with shared content supporting cross-platform discovery while exclusive content preserves an additional reason to engage on Douyin. Together, these findings suggest that the gains from multihoming depend on how creators combine shared and exclusive offerings, rather than simply on how much content they duplicate.

6.4 Alternative explanation: Changes in Measured Content Quality

The preceding analyses emphasize horizontal differentiation: creators adjust what they offer across platforms, and engagement varies with how closely their content matches their RedNote offering. We now examine vertical quality, whether creators improve the technical quality and presentation of their Douyin videos after expansion. Learning from RedNote could contribute to the audience gains by helping creators produce better videos, beyond changing their content's positioning or distribution (Polidoro and Yang, 2024). Examining measured quality therefore helps distinguish improvements in production from changes in how content connects with audiences across platforms.

We use the video-level sample described in Section 3.2. Each control creator is paired with a treated creator and assigned the treated creator's RedNote adoption date. For videos published within 60 days before or after that date, we estimate

$$y_{iv} = \beta \left(\text{Multihoming}_i \times \text{Post}_v\right) + \gamma\, \text{Multihoming}_i + \delta\, \text{Post}_v + \lambda_{p(i)} + \varepsilon_{iv}, \tag{5}$$

where y_{iv} measures a quality or presentation attribute of video v published by creator i. The indicator Multihoming_i equals one for creators who adopt RedNote, and Post_v indicates videos published after the pair's assigned adoption date. Pair fixed effects, $\lambda_{p(i)}$, absorb characteristics common to each treated–control pair. The coefficient β measures the change among treated creators relative to the change among their paired controls.

We examine eight attributes using the measurement tools described in Section 3.3. FineVQ measures overall technical quality. The Qwen3-Omni audio–visual coder measures lighting, visual stability, and speech clarity, while its presentation task assesses background staging and subtitle readability. The DeepSeek-V3 text coder identifies whether the transcript contains clearly scripted speech and a verbal opening hook. Together, these measures capture technical quality and observable presentation choices through which learning might affect audience engagement.

Table 6: **Multihoming and Measured Douyin Video Quality and Presentation.**

	Technical quality				Visual presentation		Narrative features	
	(1) Overall quality (0–100)	(2) Lighting (1–3)	(3) Visual stability (1–3)	(4) Speech clarity (1–3)	(5) Background staging (1–3)	(6) Subtitle readability (1–3)	(7) Scripted speech (0/1)	(8) Verbal opening hook (0/1)
Multihoming × Post	0.067	−0.012	−0.039	−0.029	0.009	−0.025	−0.017	−0.025
	(0.546)	(0.035)	(0.036)	(0.047)	(0.017)	(0.036)	(0.040)	(0.032)
Constant	56.012***	2.496***	2.836***	2.384***	2.028***	2.640***	0.248***	0.455***
	(0.490)	(0.031)	(0.019)	(0.032)	(0.014)	(0.025)	(0.033)	(0.031)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	37,870	34,650	34,650	30,842	35,763	28,434	23,352	29,624
R^2	0.467	0.347	0.194	0.378	0.206	0.268	0.497	0.281

Notes: Each column estimates Equation (5) using Douyin videos published within 60 days before or after the treated creator's RedNote adoption date. Each control creator is assigned the adoption date of the paired treated creator. Multihoming identifies treated creators, and Post identifies videos published after the assigned adoption date. All specifications include the main effects of Multihoming and Post and pair fixed effects. Standard errors clustered at the creator level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 6 shows no statistically detectable changes in these attributes following adoption. The interaction coefficients are statistically indistinguishable from zero across all eight outcomes. Online Appendix Table E.4 and Figure F.5 report a complementary analysis separating post-adoption videos into cross-posted and Douyin-exclusive content. Neither category exhibits statistically significant changes relative to creators'

pre-adoption videos, net of the change among controls. Thus, the aggregate results do not conceal detectable improvements concentrated in either category.

The results reveal a distinction between horizontal content adjustments and vertical quality improvements. Creators' Douyin content becomes more closely aligned with RedNote, and related exclusive posts receive more engagement, yet we find no detectable improvement in measured technical quality and presentation after adoption. Taken together, these findings are consistent with multihoming gains arising through audience discovery and the alignment of content with audience interests, without a broad increase in measured vertical quality.

7 Conclusion

This paper examines how creator multihoming affects audience performance and content supply on the platform where creators began. Using linked panel data from Douyin and RedNote, we document that multihoming is common and involves selective content sharing. Our staggered DiD analysis indicates that adopting RedNote increases creators' Douyin followers, engagement per post, and posting volume. Content-level evidence supports cross-platform discovery as one contributor to these gains. Highly engaging posts classified as exclusive to RedNote are followed by increased Douyin follower growth, and Douyin-exclusive posts receive more engagement when they resemble the creator's RedNote offerings. These patterns occur without statistically detectable improvements in measured technical quality and presentation. Our central contribution is to connect creators' platform-entry decisions to audience formation across portfolios of shared and exclusive content. The findings suggest that discovery can extend beyond the content initially encountered to the creator's other offerings, allowing participation on a competing platform to strengthen an established account. This insight highlights how cross-platform discovery and differentiated content portfolios can jointly shape the returns to creator multihoming.

Managerial Implications For creators, our findings suggest evaluating platform expansion through its effects on the audience gained on the additional platform and the performance of the established account. In our setting, expansion is accompanied by increases in focal-platform followers, engagement, and content supply, making outside exposure a potentially valuable component of audience development. Content allocation is also relevant. The evidence on related exclusive posts suggests that maintaining a coherent presence across platforms can connect discovery on one service to engagement with additional content on another. Creators may therefore benefit from experimenting with combinations of shared and platform-specific content. The observed allocation curve does not establish a common optimal sharing ratio; creators should assess these combinations against their own audience responses and the costs of production, adaptation, and account management.

For platform managers, the findings identify a discovery benefit to consider when evaluating creators' participation on competing services. A creator's outside presence can be followed by greater audience activity and content supply within the creator's established account. Assessments of exclusivity should therefore weigh these potential returns alongside the value of differentiated content and the risk of audience substitution.

Platforms could test whether features that help users recognize creators across services or facilitate selective content sharing strengthen these connections. Such evaluations should measure overall platform usage and revenue as well as participating creators' outcomes, because creator-level gains alone do not establish the net benefit of facilitating multihoming. Our evidence thus informs the trade-off surrounding exclusivity while leaving the preferred governance policy dependent on platform-specific benefits and costs.

Limitations and Future Research Directions. Our study has two limitations that suggest directions for future research. First, we do not trace individual users across platforms. We therefore cannot distinguish new Douyin users from existing users who begin following a creator, or determine whether creator-level gains expand total platform attention or redistribute attention from other creators. Future research linking user-level activity across platforms could distinguish platform adoption from creator discovery and assess how these audience flows affect overall platform usage. Second, our analysis focuses on outcomes shortly after multihoming adoption. The longer-run effects may differ as creators adjust their content portfolios, users change their consumption patterns, and platforms respond through governance policies. Longer panels, combined with experimental variation in cross-platform discovery or content sharing, could establish whether the observed gains persist and how they depend on these adjustments. Such evidence would help identify the conditions under which creator multihoming generates sustained benefits for creators and competing platforms.

References

Abolfathi, N. (2026). Multihoming and single-homing complementors' responses to intensified between-platform competition: Evidence from the YouTube–Twitch rivalry. *Strategic Management J.*, 47(1):141–188.

An, K., Chen, Q., …, and Zheng, S. (2024). FunAudioLLM: Voice understanding and generation foundation models for natural interaction between humans and LLMs. Technical report arXiv:2407.04051, Alibaba Group.

Aridor, G. (2025). Measuring substitution patterns in the attention economy: An experimental approach. *RAND J. Econom.*, 56(3):302–324.

Armstrong, M. (2006). Competition in two-sided markets. *RAND J. Econom.*, 37(3):668–691.

Bairathi, M., Lambrecht, A., and Rao, A. (2026). Social media, music consumption, and cross-platform spillover effects. *Management Sci. Forthcoming*.

Belleflamme, P. and Peitz, M. (2019). Platform competition: Who benefits from multihoming? *Internat. J. Indust. Organ.*, 64:1–26.

Berger, J., Humphreys, A., Ludwig, S., Moe, W. W., Netzer, O., and Schweidel, D. A. (2020). Uniting the tribes: Using text for marketing insight. *J. Marketing*, 84(1):1–25.

Berger, J. and Milkman, K. L. (2012). What makes online content viral? *J. Marketing Res.*, 49(2):192–205.

Branch, G. (2024). How to more effectively select influencers for your next campaign. Forbes, March 7, 2024. https://www.forbes.com/councils/forbesagencycouncil/2024/03/07/how-to-more-effectively-select-influencers-for-your-next-campaign/. Accessed September 29, 2026.

Cao, Z., Zhu, Y., Li, G., and Qiu, L. (2024). Consequences of information feed integration on user engagement and contribution: A natural experiment in an online knowledge-sharing community. *Inform. Systems Res.*, 35(3):1114–1136.

Casner, B. and Teh, T.-H. (2025). Content-hosting platforms: Discovery, membership, or both? *RAND J. Econom.*, 56(4):668–688.

Cennamo, C., Ozalp, H., and Kretschmer, T. (2018). Platform architecture and quality trade-offs of multihoming complements. *Inform. Systems Res.*, 29(2):461–478.

Ch, D. (2026). RedNote: User growth and revenue statistics (2026). SendShort, January 11, 2026. https://sendshort.ai/statistics/rednote/. Accessed September 29, 2026.

Corts, K. S. and Lederman, M. (2009). Software exclusivity and the scope of indirect network effects in the U.S. home video game market. *Internat. J. Indust. Organ.*, 27(2):121–136.

Curry, D. (2025). TikTok revenue and usage statistics (2025). Business of Apps, November 20, 2025. https://www.businessofapps.com/data/tik-tok-statistics/. Accessed September 29, 2026.

de Chaisemartin, C. and D'Haultfœuille, X. (2025). Credible answers to hard questions: Differences-in-differences for natural experiments. Working paper, Sciences Po, Paris. SSRN 4487202.

DeepSeek-AI (2024). DeepSeek-V3 technical report. Technical report arXiv:2412.19437, DeepSeek.

Doosti, S., Lee, S., and Tan, Y. (2026). Social media sponsorship: Metrics and strategies for finding the right content creator-sponsor matches. *Management Sci.*, 72(7):6043–6062.

Duan, H., Hu, Q., Wang, J., Yang, L., Xu, Z., Liu, L., Min, X., Cai, C., Ye, T., Zhang, X., and Zhai, G. (2025). FineVQ: Fine-grained user generated content video quality assessment. In *Proc. CVPR*, pages 3206–3217, Nashville, TN. IEEE.

El-Komboz, L. A., Kerkhof, A., and Loh, J. (2023). Platform partnership programs and content supply: Evidence from the YouTube "Adpocalypse". CESifo Working Paper 10363, CESifo, Munich.

Gong, Z. (2021). Growing influence. Working paper, University of Toronto, Toronto.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *J. Econometrics*, 225(2):254–277.

Hao, Y. and Li, X. (2026). Shaping the influencers: The role of multi-channel networks. Working paper, University of Rochester, Rochester, NY. SSRN 4720070.

Hendricks, K. and Sorensen, A. (2009). Information and the skewness of music sales. *J. Political Econom.*, 117(2):324–369.

Higgins, S. (2024). Financial technology adoption: Network externalities of cashless payments in Mexico. *Amer. Econom. Rev.*, 114(11):3469–3512.

Iyer, R. (2026). Multistreaming guide (2026): Twitch, Kick and YouTube rules explained. Streams Charts, July 2, 2026. https://streamscharts.com/news/multistreaming-guide-2026-rules-explained. Accessed September 29, 2026.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50(3):649–670.

Kapacinskaite, A. and Mostajabi, A. (2024). Competing with the platform: Complementor positioning and cross-platform response to entry. *Strategic Management J.*, 45(12):2577–2607.

Krijestorac, H., Garg, R., and Mahajan, V. (2020). Cross-platform spillover effects in consumption of viral content: A quasi-experimental analysis using synthetic controls. *Inform. Systems Res.*, 31(2):449–472.

Landsman, V. and Stremersch, S. (2011). Multihoming in two-sided markets: An empirical inquiry in the video game console industry. *J. Marketing*, 75(6):39–54.

Lauron, S. (2026). Multi-platform creator strategies: Why diversification matters in 2026. Later, January 12, 2026. https://later.com/blog/how-diversification-strategies-create-resilient-creator-businesses/. Accessed September 29, 2026.

Lee, D., Hosanagar, K., and Nair, H. S. (2018). Advertising content and consumer engagement on social media: Evidence from Facebook. *Management Sci.*, 64(11):5105–5131.

Leung, F. F., Gu, F. F., Li, Y., Zhang, J. Z., and Palmatier, R. W. (2022). Influencer marketing effectiveness. *J. Marketing*, 86(6):93–115.

Li, H. and Zhu, F. (2021). Information transparency, multihoming, and platform competition: A natural experiment in the daily deals market. *Management Sci.*, 67(7):4384–4407.

Li, P., Castelo, N., Katona, Z., and Sarvary, M. (2024). Frontiers: Determining the validity of large language models for automated perceptual analysis. *Marketing Sci.*, 43(2):254–266.

Lu, S. and Miao, H. (2025). TikTok refugees find an alternative in China. Wall Street Journal, January 13, 2025. https://www.wsj.com/tech/tiktok-refugees-find-an-alternativein-china-0b0418f6. Accessed September 29, 2026.

Mallipeddi, R. R., Kumar, S., Sriskandarajah, C., and Zhu, Y. (2022). A framework for analyzing influencer marketing in social networks: Selection and scheduling of influencers. *Management Sci.*, 68(1):75–104.

Polidoro, Jr., F. and Yang, W. (2024). Porting learning from interdependencies back home: Performance implications of multihoming for complementors in platform ecosystems. *Strategic Management J.*, 45(9):1791–1821.

Qian, K. and Jain, S. (2024). Digital content creation: An analysis of the impact of recommendation systems. *Management Sci.*, 70(12):8668–8684.

Qiu, L. and Kumar, S. (2017). Understanding voluntary knowledge provision and content contribution through a social-media-based prediction market: A field experiment. *Inform. Systems Res.*, 28(3):529–546.

Reimers, N. and Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proc. EMNLP-IJCNLP*, pages 3982–3992, Hong Kong. Association for Computational Linguistics.

Ren, Q. (2024). Advertising and content creation on digital content platforms. *Marketing Sci.*, 43(4):734–750.

Suciu, P. (2025). TikTok users heading to RedNote: It looks similar, has the same issues. Forbes, January 15, 2025. https://www.forbes.com/sites/petersuciu/2025/01/15/tiktok-users-heading-to-rednote--it-looks-similar-has-the-same-issues/. Accessed September 29, 2026.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometrics*, 225(2):175–199.

Sung, M. (2023). Twitch's CMO says competition from rival platforms is good for creators. TechCrunch, October 27, 2023. https://techcrunch.com/?p=2621378. Accessed September 29, 2026.

Wang, K. and Goldfarb, A. (2017). Can offline stores drive online sales? *J. Marketing Res.*, 54(5):706–719.

Xiao, S., Liu, Z., Zhang, P., Muennighoff, N., Lian, D., and Nie, J.-Y. (2024). C-Pack: Packed resources for general Chinese embeddings. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Washington, DC, USA. Association for Computing Machinery.

Xu, J., Guo, Z., …, and Lin, J. (2025). Qwen3-Omni technical report. Technical report arXiv:2509.17765, Alibaba Group.

Yang, J., Han, Q., and Jerath, K. (2023). Collaboration among content creators. Working paper, Columbia University, New York. SSRN 4538856.

Yang, J., Zhang, J., and Zhang, Y. (2025). Engagement that sells: Influencer video advertising on TikTok. *Marketing Sci.*, 44(2):247–267.

Zhang, X., Gu, Y., Zhang, Y., and Manchanda, P. (2026). When professionals become influencers: The impact of live streaming on service demand. Working paper, London Business School, London. SSRN 5915365.

Zhao, K., Lu, Y., Hu, Y., and Hong, Y. (2023). Direct and indirect spillovers from content providers' switching: Evidence from online livestreaming. *Inform. Systems Res.*, 34(3):847–866.

Online Appendix:

Content Creator Multihoming and Attention Spillovers Across Platforms

A Technical Details of Video and Text Analysis

This appendix documents how the video-level sample described in Section 3.3 is converted into the content measures used in Section 6.4, following the pipeline in Figure F.2. We first separate each video into the modalities we analyze, the audio-visual stream and the speech transcript, and then run modality-specific analyses. Below we describe, for each measure, the tool used, what it returns, and how the variable is constructed.

A.1 Video Preparation

Using FFmpeg, we derive three clips from each original video. The *semantic clip* is the centered window of at most 58 seconds; the multimodal coder receives 32 frames sampled from this window together with its full audio. The *quality clip* is the centered 12-second window, re-encoded losslessly at a constant 30 frames per second, because the quality model expects a constant frame rate and any lossy re-encoding would inject compression artifacts into the construct being measured. The *opening clip* covers the first six seconds and is used only to code opening devices, since a centered window does not contain the opening of videos longer than the window. Fixing the analysis window keeps videos of different lengths comparable: without it, a three-minute video would supply several times more evidence to the coder than a twenty-second video. Videos shorter than a window are analyzed in full. Clips carry blinded identifiers, so no model observes a creator's identity, treatment status, or posting period.

A.2 Technical Quality

We score technical quality with FineVQ (Duan et al., 2025), a fine-grained video quality assessment model for user-generated content trained on the FineVD database of 6,104 videos with human quality ratings. The model takes the quality clip as input and returns an overall quality score on a 0 to 100 scale together with five component scores for blur, color, noise, artifact, and temporal stability. The overall score is the *FineVQ overall* outcome in Table 6 (column 1). FineVQ is used only for scoring technical quality; it is never asked to interpret content.

A.3 Production Attributes

Production attributes are coded by Qwen3-Omni (Xu et al., 2025), an open-weights multimodal large language model that reads sampled frames, the audio track, and on-screen text jointly. The model receives

the semantic clip and a fixed prompt that defines each attribute and lists its allowed values, and it returns a single JSON object; decoding is deterministic (temperature zero), and a response that does not conform to the allowed values is retried up to three times and otherwise recorded as missing. Two coding tasks supply the Table 6 measures. The production task codes *lighting quality* (poor, basic, polished), *visual stability* (unstable, mixed, stable), and *speech audio clarity* (poor, basic, polished), which we map to 1 to 3 scales (columns 2 to 4). The presentation task codes *background environment* (cluttered, ordinary tidy, deliberately staged) and *subtitle readability* (hard to read, adequate, highly readable, among videos with subtitles), again mapped to 1 to 3 scales (columns 5 and 6). The same coder also returns the video's topic and format, a set of further production items (multi-scene structure, B-roll, transitions, graphics, audio layering), and whether advertising appears; these attributes are coded under the same protocol and are available in the replication package. The prompt for the production task reads as follows (allowed values abridged to the attributes used in the paper):

> *You are a blinded research coder measuring observable production attributes of one Chinese short-form video. Use both the video and its soundtrack. Judge only evidence that can be directly seen or heard. Do not infer actual labor hours, budget, team size, monetary cost, creator fame, treatment status, engagement, popularity, or likely success. Ignore topic appeal, creator attractiveness, platform watermarks, and user-interface elements. Return exactly one JSON object and no markdown. Use only the allowed values below.* […] *"lighting_quality": one of ["poor", "basic", "polished", "unknown"]; "visual_stability": one of ["unstable", "mixed", "stable", "unknown"]; "speech_audio_clarity": one of ["poor", "basic", "polished", "no_speech", "unknown"].* […] *A model rating is an observable proxy, never an estimate of actual hours or cost.*

The presentation task defines its two attributes as follows:

> *background_environment: tidiness and staging of the visible shooting environment. cluttered: visibly messy or distracting background in a meaningful share of the clip. ordinary_tidy: a normal, reasonably tidy everyday environment without deliberate decoration for the camera. deliberately_staged: the environment is visibly arranged for the camera, such as a styled backdrop, studio wall, ring-light-lit dressing area, or branded set. subtitle_readability: readability of the primary content subtitles added by the creator, ignoring platform interface elements and watermarks. no_subtitles; hard_to_read: small, low-contrast, cut off, overlapping visuals, or too fast to read; adequate: readable with normal attention; highly_readable: large, high-contrast, well-timed subtitles styled for effortless reading.*

A.4 Transcript-Based Attributes

Following Zhang et al. (2026), we analyze the speech track separately from the visuals. We extract the full audio of each video, resample it to 16 kHz mono, and transcribe it with SenseVoice (An et al., 2024), an open-weights multilingual speech recognition model that returns time-stamped utterances. The transcript covers the entire video rather than the centered window, and we mark its opening, middle, and closing thirds so that the position of speech within the video is known. Transcripts shorter than 30 characters are classified as containing insufficient speech and are excluded from transcript-based measures.

Each transcript is then coded by DeepSeek-V3 (DeepSeek-AI, 2024) under the same protocol as the video coder: a fixed prompt with closed answer categories, deterministic decoding, and validation of the

returned JSON object. Two attributes enter Table 6. *Scripted speech* (column 7) equals one when the coder classifies the speech structure as clearly scripted rather than spontaneous or lightly structured. *Verbal opening hook* (column 8) equals one when the opening third of the transcript begins with a deliberate attention hook. The text coder also returns the topic of the speech, using the same taxonomy as the video coder, together with sentiment and spoken advertising, which we use as transcript-based checks on the video-based codings. The relevant part of the prompt reads:

> *You are coding the speech transcript of one Chinese short-form video for a scientific study. The transcript was produced by automatic speech recognition: expect occasional homophone errors and missing punctuation; read through such noise charitably, but judge only what the transcribed speech says. You cannot see the video. Do not infer visuals, the creator's identity, treatment status, popularity, engagement, income, effort, or success. The transcript is divided by time markers into the first, middle, and final third of the video.* [...] *"speech_structure" judges from wording alone how scripted the speech reads: "clearly_scripted" = polished written-style phrasing, planned rhetorical structure, dense information, virtually no fillers; "spontaneous" = fillers, self-corrections, loose rambling, reactive in-the-moment talk; "lightly_structured" = an evident plan but casual delivery. "opening_hook_speech" is "yes" when the opening third opens with a deliberate attention hook in the words themselves: a direct question to the viewer, a promise or payoff tease, suspense, or a shocking claim. A plain greeting or diving straight into content is "no".*

B Discrete-Time Hazard Model Results

We formally test whether the timing of multihoming is correlated with the levels or trends of observable characteristics. Following Higgins (2024), we estimate a discrete-time hazard model, where the dependent variable indicates whether a creator adopts multihoming in a given month. The model includes time-varying covariates, including follower size, number of posts, number of likes, and their corresponding growth rates. As reported in Table B.1, the levels of followers, posts, and likes are significantly associated with the timing of multihoming, whereas their growth rates are not statistically significant. This distinction is important for identification. In a DiD framework, selection on outcome levels is permissible and absorbed by creator fixed effects. In contrast, selection on short-run trends would threaten the parallel trends assumption. The absence of a systematic relationship between adoption timing and growth rates suggests that differential pre-trends are unlikely to drive our results.

Table B.1: **Discrete-Time Hazard Model Results.**

Variable	(1) Mean	(2) SD	(3) Discrete-time hazard
Female	0.475	0.499	0.003 (0.004)
Big city	0.089	0.285	0.017** (0.007)
log(*Followers*)	13.733	1.180	−0.007*** (0.002)
Growth rate of followers	0.092	0.146	0.030 (0.019)
log(*Likes*)	10.792	1.236	0.009*** (0.002)
Growth rate of likes	−0.512	3.683	−0.0001 (0.0003)
log(*Number of posts*)	2.105	0.850	0.005** (0.003)
Growth rate of post counts	−0.248	1.242	−0.0006 (0.0016)
Month FE			Yes
Category dummies			Yes
Observations			9,653
R-squared			0.020

Notes: Columns (1) and (2) show the mean and standard deviation of characteristics in the estimation sample. Column (3) tests whether these characteristics predict the timing of multihoming in a single regression, using a linear probability discrete-time hazard with month fixed effects and content-category dummies; big city is relative to all other cities. The dependent variable is a dummy variable indicating whether creator i adopts RedNote at time t. Standard errors are clustered at the creator level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Robustness Checks

This appendix gives the specifications and full results for the robustness checks summarized in Table 4. It covers the two event-study estimators behind Figure 4, the test for interference across creators, the alternative definition of multihoming, and the symmetric analysis with RedNote as the focal platform.

C.1 Relative Time Model and Staggered DiD

TWFE event study. We replace the treatment indicator D_{it} in Equation (1) with indicators for each month from five months before to five months after RedNote adoption, binning earlier and later months at the two ends of the window. The month before adoption is omitted, so each coefficient measures the change relative to that month. All indicators are zero for control creators. The specification keeps the creator and month fixed effects and clusters standard errors by creator. The pre-adoption coefficients test for differential trends between adopters and controls, and the post-adoption coefficients trace the effect month by month. The orange triangles in Figure 4 report these estimates with 95% confidence intervals.

Staggered DiD estimator. Under staggered adoption, the TWFE estimator averages all two-by-two comparisons in the data, including ones in which earlier adopters serve as controls for later adopters. These comparisons can bias the estimates when effects differ across cohorts or change with time since adoption (Goodman-Bacon, 2021). We therefore re-estimate the dynamic effects with the interaction-weighted estimator of Sun and Abraham (2021). It estimates the effect for each adoption cohort and event time separately, with never-treated creators as the only controls and the month before adoption as the reference period, and then averages the cohort estimates at each event time by their share of treated observations. The blue crosses in Figure 4 report these estimates on the same sample, with standard errors clustered by creator.

C.2 Interference Across Creators

Our DiD design assumes that one creator's RedNote adoption does not affect untreated creators' Douyin outcomes (SUTVA). Competition for attention within content categories could violate this assumption. We examine this possibility by regressing outcomes for never-treated and not-yet-treated creators on the monthly share of their same-category peers already multihoming, excluding the focal creator. The regressions include creator and month fixed effects, with standard errors clustered by content category (22 categories). The exposure coefficients in Table C.1 are statistically insignificant and small relative to the main effects in Table 3. We therefore find no evidence of within-category interference.

C.3 Adding Linear Time Trends

A remaining concern is that creators may adopt RedNote while already on a distinct trajectory on Douyin, which creator fixed effects do not absorb. We therefore add a creator-specific linear time trend to the baseline specification, so that the multihoming coefficient is identified from deviations of each creator's outcome from their own linear path. Table C.2 reports the results. All effects retain their sign and remain significant, but

Table C.1: **SUTVA Test: Leave-One-Out Exposure to Same-Category Peers' Multihoming.**

VARIABLES	(1) log (Followers)	(2) log (Likes)	(3) log (Comments)	(4) log (Shares)
Peer exposure	−0.066	−0.031	0.537	0.074
	(0.469)	(0.457)	(0.517)	(0.447)
Constant	13.71***	10.65***	7.291***	8.205***
	(0.254)	(0.245)	(0.278)	(0.240)
Creator fixed effects	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes
Observations	15,205	14,496	14,496	14,496
R-squared	0.910	0.637	0.584	0.633

Notes: Peer exposure is the leave-one-out share of the creator's same-content-category peers already multihoming in that month. An insignificant exposure coefficient indicates no interference from peers' adoption onto untreated creators (SUTVA). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

smaller. Because the trends are fitted over the full window, an effect that builds after adoption, as the follower effect does in the event study, is partly absorbed into the trend, so these estimates are conservative.

Table C.2: **Effects of Multihoming with Creator-Specific Linear Trends.**

VARIABLES	(1) log (Followers)	(2) log (Likes)	(3) log (Comments)	(4) log (Shares)	(5) Number of posts	(6) Cross-platform content similarity
Multihoming	0.138***	0.159***	0.117**	0.091*	2.347***	0.0131***
	(0.022)	(0.045)	(0.051)	(0.054)	(0.579)	(0.0028)
Constant	13.71***	10.70***	7.617***	8.412***	9.198***	0.7487***
	(0.010)	(0.020)	(0.023)	(0.024)	(0.248)	(0.0013)
Creator fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Creator linear trends	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,645	26,815	26,815	26,815	30,072	26,745
R-squared	0.970	0.732	0.664	0.698	0.815	0.641

Notes: Each regression adds a creator-specific linear time trend to the creator and month fixed effects of the baseline specification. Likes, comments and shares are the log of the creator's mean per post in the month. Robust standard errors are clustered at the creator level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.4 Alternative Definition of Multihoming

In the main analysis, treatment begins in the month a creator first becomes active on RedNote, regardless of what they post there. A stricter definition dates multihoming from the point at which the creator actually uses both platforms for the same content. Using the cross-posted posts identified in Section 3.4, we redefine treatment onset as the creator's first month with at least two cross-posted posts, floored at the RedNote adoption month, and re-estimate Equation (1) on the same sample, with the same fixed effects and clustering, as

Table 3. Table C.3 reports the results. The estimates are slightly larger than the baseline estimates for all six outcomes and equally significant, so the main results do not depend on the timing convention.

Table C.3: **Effects of Multihoming (Robustness Check).**

VARIABLES	(1) log (Followers)	(2) log (Likes)	(3) log (Comments)	(4) log (Shares)	(5) Number of posts	(6) Cross-platform content similarity
Multihoming	0.327***	0.266***	0.274***	0.276***	3.530***	0.0169***
	(0.035)	(0.048)	(0.055)	(0.058)	(1.088)	(0.0030)
Constant	13.66***	10.68***	7.577***	8.360***	9.101***	0.7489***
	(0.012)	(0.017)	(0.019)	(0.020)	(0.340)	(0.0010)
Creator fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,645	26,815	26,815	26,815	30,072	26,745
R-squared	0.889	0.629	0.563	0.615	0.569	0.545

Notes: Robustness to a more conservative definition of multihoming: treatment begins in the creator's first month with at least one cross-posted posts (floored at the RedNote adoption month), rather than at RedNote adoption. Likes, comments and shares are the log of the creator's mean per post in the month. Robust standard errors are clustered at the creator level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.5 Symmetric Analysis with RedNote as the Focal Platform

To check that the results are not specific to Douyin as the focal platform, we repeat the analysis for creators whose initial presence is on RedNote. The sample consists of the 194 RedNote-first creators who later begin posting on Douyin; treatment starts in the month of their first Douyin post. Outcomes are measured on RedNote: the logarithm of likes, the number of posts, and cross-platform content similarity, defined as the cosine similarity between the creator's RedNote content in the month and the contemporaneous content of Douyin-native creators, the mirror image of the measure used for Douyin-first creators. The sample window and cohort rules mirror the main specification. Table C.4 reports TWFE estimates with creator and month fixed effects and standard errors clustered at the creator level. After joining Douyin, creators post more on RedNote, receive more likes, and move their RedNote content closer to Douyin content, mirroring the Douyin results.

Table C.4: **Impact of Multihoming on Creator Outcomes: RedNote as the Focal Platform.**

VARIABLES	(1) log (Likes)	(2) log (Comments)	(3) log (Collects)	(4) Number of posts	(5) Cross-platform content similarity
Multihoming	0.564***	0.361***	0.569***	2.974***	0.0175**
	(0.116)	(0.090)	(0.128)	(0.717)	(0.0070)
Constant	7.332***	4.145***	5.680***	6.309***	0.7345***
	(0.097)	(0.075)	(0.107)	(0.566)	(0.0059)
Creator fixed effects	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	4,958	4,958	4,958	5,916	4,932
R-squared	0.707	0.656	0.691	0.523	0.502

Notes: RedNote-first creators who later begin posting on Douyin; treatment starts in the month of their first Douyin post. Outcomes are measured on RedNote: likes, comments and collects are the log of the creator's mean per note in the month, and posts is the number of notes. RedNote reports no share count. Robust standard errors are clustered at the creator level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D Matched-Control Event Study

The event study in Section 6.1 compares each creator's Douyin follower growth after a viral RedNote-exclusive note with the same creator's ordinary days. If a note goes viral because its topic is surging across platforms, Douyin follower growth would have risen on those days for any creator who posts on that topic. We therefore compare each event creator with control creators observed on the same calendar days. For each isolated viral RedNote-exclusive note, controls are the three to five multihoming creators closest to the event creator in the topic mix of their Douyin posts, their follower count and follower growth, and their posting activity on both platforms, all measured before the event.

We estimate

$$\Delta \log(\text{Followers}_{ied}) = \sum_{\substack{k=-14 \\ k\neq -1}}^{5} \beta_k \, \mathbf{1}\{\text{Event creator}_{ie}\} \times \mathbf{1}\{d - d_e = k\} + \gamma X_{id} + \alpha_{ie} + \delta_{ed} + \varepsilon_{ied}, \tag{D.1}$$

where $\Delta \log(\text{Followers}_{ied})$ is creator i's Douyin follower growth on calendar day d in the window of event e, d_e is the day of the note, α_{ie} are creator-by-event fixed effects, and δ_{ed} are event-by-calendar-day fixed effects, which absorb anything common to the event creator and the controls on a given day. The vector X_{id} contains the activity controls of Equation (2), the leads and lags of the creator's posting and viral days on both platforms. Events are weighted equally, controls are weighted equally within each event, and standard errors are clustered by creator.

Figure D.1 shows the result. Event creators grow faster than their controls throughout the window (Panel A), so the coefficients β_k measure the change in this gap around the note. The pre-event coefficients are small and jointly insignificant (Panel B, $p = 0.65$ for days -5 to -2). Growth rises relative to the controls by 0.04 percentage points on the day of the note and by 0.07 percentage points the following day and returns to zero within five days, which adds up to 0.21 percentage points more Douyin follower growth over days 0 to 5 (95% confidence interval 0.12 to 0.30). The response matches Figure 5 in timing and is of comparable size, so it is not an artifact of topic-wide surges in attention.

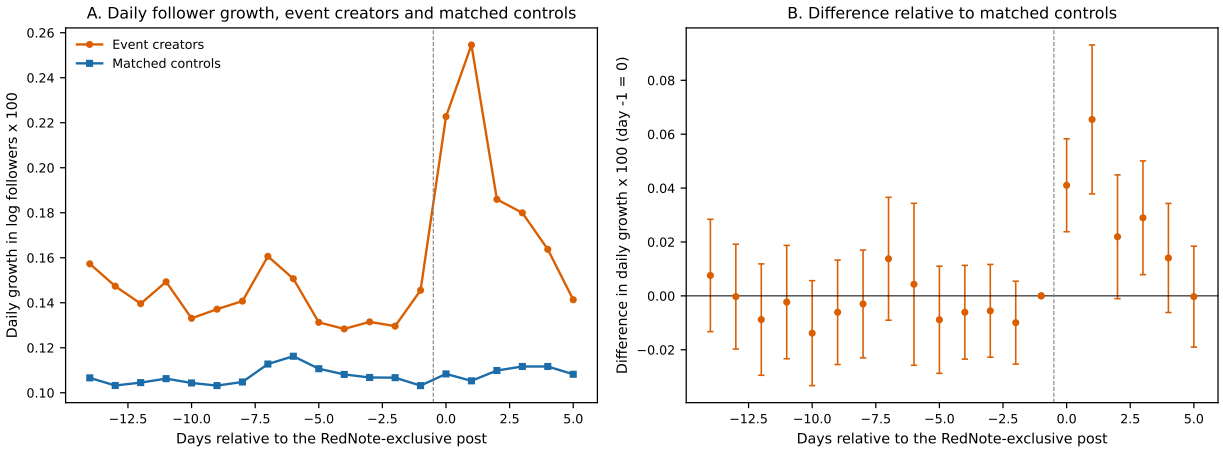


Figure D.1: Follower response to RedNote-exclusive viral notes, event creators versus matched controls.

Notes: Panel A shows average daily Douyin follower growth (in %) for event creators and their matched controls. Panel B shows the coefficients β_k from Equation (D.1) with 95% confidence intervals based on standard errors clustered by creator.

E Additional Tables

Table E.1: **Examples of Identified Cross-platform Posts.**

Posts from Douyin	**Posts from RedNote**
冷都女烟熏黑道千金风 # 美妆百万新星计划 # 单眼皮	冷都女烟熏黑道千金风 # 单眼皮眼妆 # 今日妆容 # 妆前妆后
Cold city girl smoky dark gangster chic look	Cold city girl smoky dark gangster chic look #MonolidMakeup #TodaysMakeup #BeforeAndAfterMakeup
当方圆脸尝试网红芭比妆，走进现实能看吗? # 方脸 # 方圆脸 # 美妆博主回归现实	当方圆脸尝试网红感芭比妆,现实中能看吗? # 方圆脸 # 方脸 # 芭比 # 网红妆容回归现实 # 新手化妆技巧 # 完美底妆这样画 #UD 定妆喷雾
When a square/round face tries the trendy Barbie makeup, does it look good in reality?	When a square/round face tries the trendy Barbie-style makeup, does it look good in real life?
无睫毛　超详细跟练版韩系姐姐气质单眼皮 # 单眼皮 # 新手化妆教程	无睫毛　超详细跟练版气质单眼皆真的超级无敌简单易学! # 单眼皮眼妆 # 新手化妆 # 妆前妆后大对比
No false lashes—super detailed follow-along Korean-style elegant monolid look	No false lashes—super detailed follow-along elegant monolid look, truly super easy to learn!

Notes: In each pair of rows, the first row presents the Chinese title, and the row below provides its English translation.

Table E.2: **Who Multihomes: Creator Characteristics and Multihoming Propensity.**

VARIABLES	(1) Logit	(2) Average marginal effect
Big city	0.244***	0.057***
	(0.055)	(0.013)
Female	0.329***	0.077***
	(0.035)	(0.008)
log(*Advertising price*)	0.373***	0.087***
	(0.017)	(0.004)
Constant	−4.191***	
	(0.174)	
Observations	14,290	14,290
Pseudo R^2	0.033	

Note: The dependent variable equals one if the creator is matched to a RedNote account. Column (1) reports logit coefficients and column (2) the corresponding average marginal effects. Big city indicates Beijing, Shanghai, Guangzhou, or Shenzhen; advertising price is the creator's quoted price for a 1–20 second sponsored video. Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.3: **Decomposition of the Posting Effect.**

VARIABLES	(1) All posts	(2) Cross-posted	(3) Douyin- exclusive
Multihoming	3.029***	3.895***	−0.865
	(1.128)	(0.483)	(0.795)
Constant	8.905***	−0.051	8.956***
	(0.483)	(0.207)	(0.341)
Creator fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Observations	30,072	30,072	30,072
R-squared	0.5683	0.4623	0.5491

Notes: Standard errors clustered at the creator level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.4: **Multihoming and Video Quality by Post Type.**

VARIABLES	(1) FineVQ overall (0–100)	(2) Lighting (1–3)	(3) Visual stability (1–3)	(4) Speech clarity (1–3)	(5) Background staging (1–3)	(6) Subtitle readability (1–3)	(7) Scripted speech (0/1)	(8) Verbal opening hook (0/1)
Cross-posted vs. pre-adoption	0.297	0.047	−0.010	0.039	0.011	−0.035	−0.035	0.057
	(0.613)	(0.039)	(0.048)	(0.054)	(0.022)	(0.040)	(0.044)	(0.036)
Exclusive vs. pre-adoption	−0.046	−0.041	−0.054	−0.064	0.008	−0.020	−0.006	−0.071*
	(0.595)	(0.039)	(0.038)	(0.050)	(0.017)	(0.038)	(0.045)	(0.037)
Multihoming	0.805	−0.081	−0.019	−0.055	0.029	0.082**	−0.034	−0.056
	(0.894)	(0.061)	(0.037)	(0.057)	(0.030)	(0.042)	(0.071)	(0.061)
Post	0.487	0.048**	−0.006	0.042	−0.008	0.042	−0.032**	0.017
	(0.390)	(0.021)	(0.025)	(0.039)	(0.011)	(0.028)	(0.015)	(0.020)
Constant	56.012***	2.496***	2.836***	2.384***	2.028***	2.640***	0.248***	0.455***
	(0.491)	(0.031)	(0.019)	(0.032)	(0.014)	(0.025)	(0.033)	(0.030)
Pair fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	37,870	34,650	34,650	30,842	35,763	28,434	23,352	29,624
R-squared	0.467	0.349	0.194	0.380	0.206	0.268	0.497	0.284

Note: The first two rows are the coefficients on the two type-specific interactions, each comparing that type with the creators' pre-adoption videos net of the matched control creators' change over the same period. Standard errors clustered at the creator level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F Additional Figures

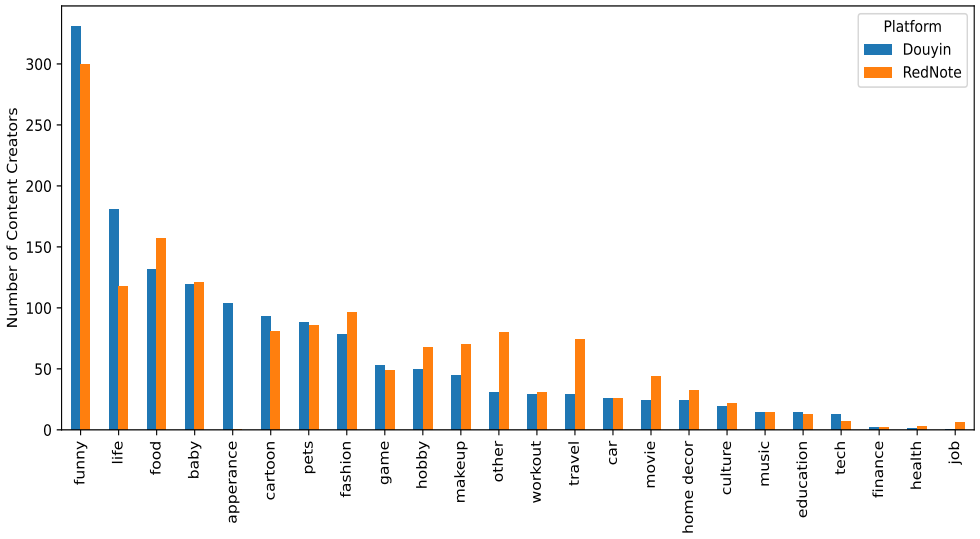


Figure F.1: **Categories of Content Creators.**

Notes: We manually reviewed and matched the tag systems of the two platforms, identifying 22 common content types, along with one unique category on Douyin ("appearance") and one unique category on RedNote ("job"). Among the 1,500 content creators in our dataset, the category distributions appear highly similar between the two platforms. This similarity suggests that creators often maintain consistent content themes across platforms.

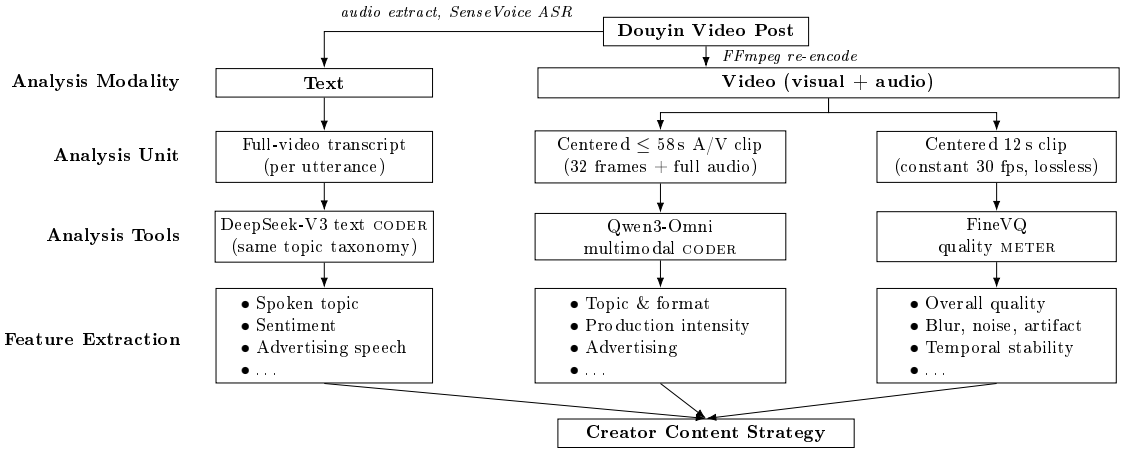


Figure F.2: **Multimodal content measurement pipeline.**

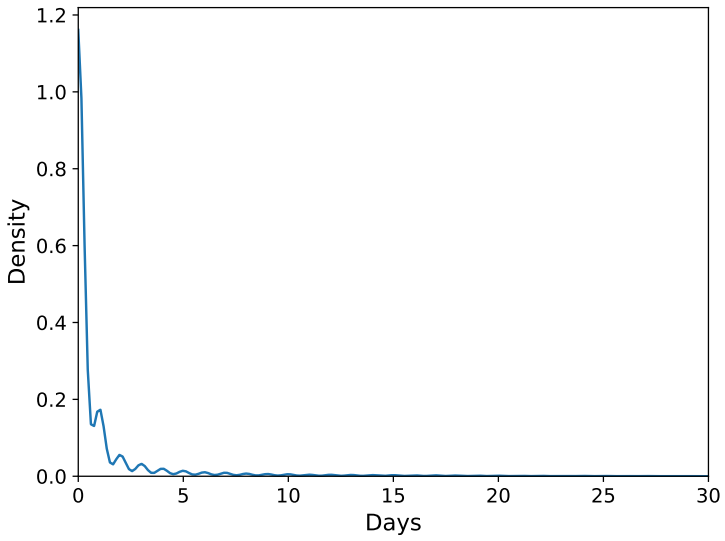


Figure F.3: **Time Differences between Cross-Posted Posts on the Two Platforms.**

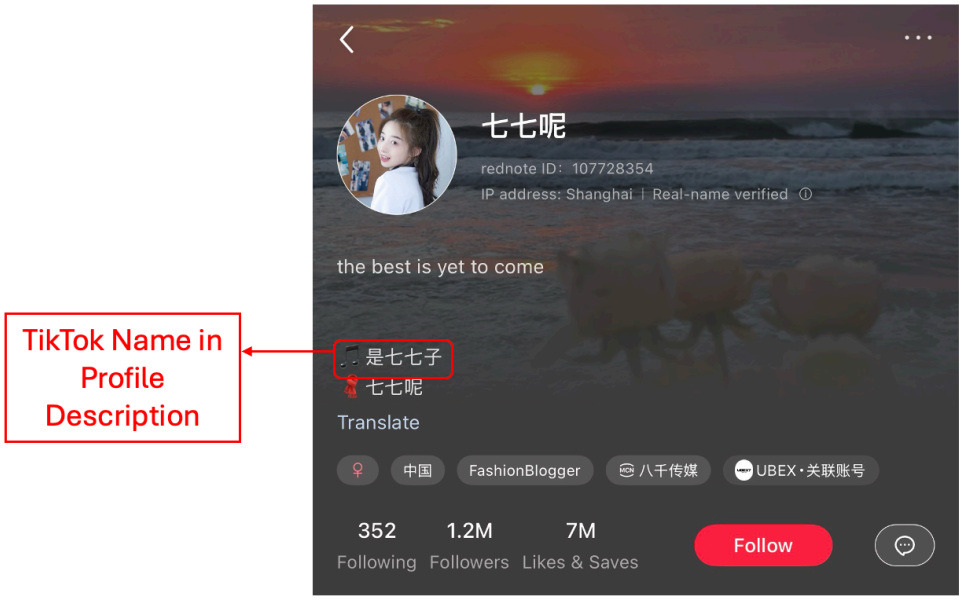


Figure F.4: **Screenshot of a RedNote Account.**

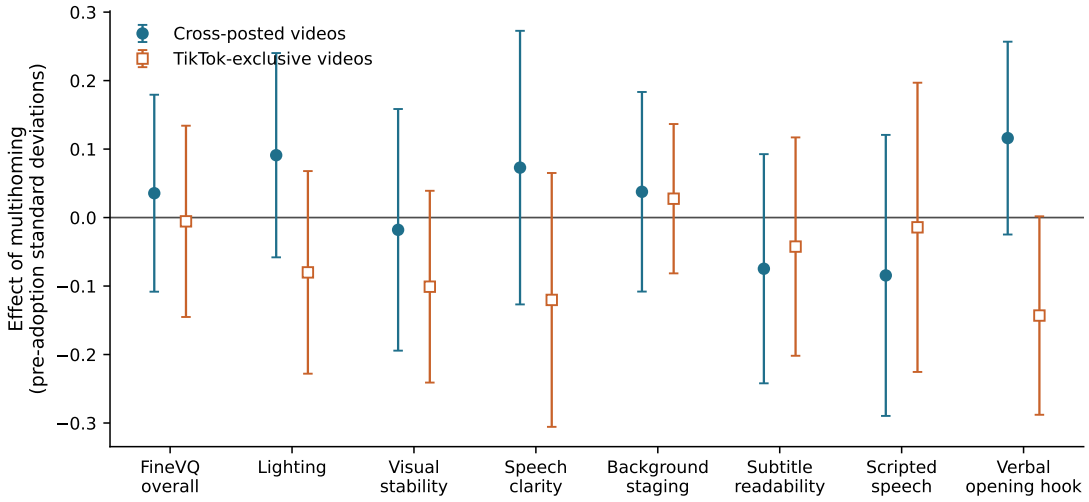


Figure F.5: Multihoming and video quality by post type.

Notes: Coefficients β_C (cross-posted) and β_E (Douyin-exclusive) for the eight facets of Table E.4, divided by the pre-adoption standard deviation of each outcome, with 95% confidence intervals based on standard errors clustered at the creator level.