The Impact of Platform Commission Design on Creators' Pricing Strategy and Productivity

Pu Zhao* Georgios Zervas* Xintong Han†

Latest version: March 14, 2024

Abstract

We investigate the effects of digital platform commission change on creators' economic behaviors, with a focus on pricing strategies and productivity. Utilizing a comprehensive dataset from a leading creator platform in China with over 7,186 creators in 51 months, our research employs two difference-in-differences frameworks to empirically assess the impact of a staggered commission reduction. The focal platform increased commission from 5% to 20% in August 2019 uniformly for all creators. Concurrently, a policy allowing creators to revert to the original 5% commission upon meeting specific criteria in differential timing was introduced. Our findings reveal that the adoption of the reduced commission policy led to an 8%-13% increase in subscription prices, a 31%-70% rise in original content production, and an 8%-26% growth in creator-subscriber engagement. We test two mechanisms for the price increase given reduced commission, one through market structure change due to the initial uniform commission increase policy, and another through the selective nature of staggered commission reduction to subsidize creators who were more willing to produce content in greater number and higher quality and therefore incurred higher cost of content production. These results underline the potential of tailored commission rates to encourage more economically beneficial content production, thereby enhancing the platform's vibrancy and sustainability.

Keywords: Platform Commission Design, Third-Degree Price Discrimination, Interrupted Time Series, Staggered Difference-in-Differences

*Boston University, Questrom School of Business. Boston, MA, USA. Contact: `puzhao@bu.edu`
†Université Laval, Faculty of Business Administration. Quebec City, Quebec, Canada.

1 Introduction

Digital marketplaces such as Apple's App Store, Amazon, and YouTube catalyze business growth by connecting them to huge, global audiences. Yet, these platforms are in themselves substantial commercial entities that seek to maximize their own profits by taking commission from the participants' revenue. This leads to a constant tension: independent developers, content creators, and online retailers owe their success to platforms like the App Store, creator platforms, and e-commerce apps, but they must continually compensate platforms for the privilege of infrastructure, exposure, and distribution through arguably high platform commission.

The recent court case, *Epic Games vs. Apple*, has heightened public discussions on the necessity and fairness of 30% platform commission, affecting businesses and customers beyond these two companies. Apple maintained that its 30% commission on In-App Purchases was fair, citing its superior marketing efforts, customer service, and distribution services as benefits. In contrast, Tim Sweeney, Epic Games' CEO, argued that an 8% commission would suffice for profitable operations in digital platforms. To demonstrate this argument, Sweeney launched the Epic Games Store with only 12% commission.

The heated debate in the field led to two streams of research that investigates the effect of the platform commission on various outcomes of all players in the platform ecosystem. First, recent studies using structural models demonstrate that the impact of a simultaneous fixed commission *reduction* or *cap* is negative on the welfare of at least party of the platform ecosystem (e.g., Barwick and Pathak 2015, Robles-Garcia 2019, Sullivan 2022, Lu, Goldfarb, and Mehta 2023). Our work using an *interrupted time series* design provides novel empirical evidence of the effect of a simultaneous fixed *increase* in commission by a subscription-based creator platform in a natural experiment setting. Second, among research studies on the influence of commission *cap*s, very few adopt a casual inference framework. To the best our knowledge, it is limited to Li and Wang (2024) using the canonical Two-Way Fixed Effects method. We are the first to conduct a series of more robust causal inference analysis of commission reduction that features multiple periods and variation in treatment timing on a creator platform. Another major distinction of our paper and all others is that the platform commission reduction in this paper can be considered as a subsidy in favor of big contributors on the platform, while the platform commission cap in papers like Li and Wang (2024) works as a heavier platform tax on big businesses.

Understanding the platform commission design is crucial to identifying the optimal business strategy of the whole digital platform ecosystem. In this paper, we focus on the impact of platform commission design on creators' pricing strategy and their productivity. Supply side factors such as price and productivity are first order issues that drive the welfare analysis of our focal platform which links creators and their followers through yearly subscription. Note that most platforms charge a commission proportional to the transaction value between creators and their subscribers, and therefore the price set by a creator not only represents the revenue per subscriber but also affects the platform profitability. This alignment of creators and the platform's revenue is especially salient in our case because the creator platform we utilize in this paper is advertisement-free, and commissions alone contributes to 100% of platform revenue. In addition, the creators' productivity with no doubt directly influences customer experience and the residual traffic to the platform, which ultimately decides the user base on this platform. In sum, the platform commission change, regardless affecting all creators or some of them, may move the equilibrium for all players on this platform.

To empirically investigate the impact of platform commission design on creators' pricing and productivity, we leverage a comprehensive panel data set of a leading Chinese creator platforms called "Knowledge Planet"[1] which accommodates thousands of creators publishing highly personalized content and charging yearly subscription fee to their audience in virtual space called "planet". Our data consists of 7,186 planets and more than 111,000 observations for price information as well as more than 224,000 observations for productivity metrics over

[1]The Chinese pinyin for Knowledge Planet is *zhi shi xing qiu*. Readers may check out its official website at `https://www.zsxq.com/`. Even though Knowledge Planet provides web login-in, its major traffic goes through the app. Our data records information from all channels.

a period of 51 months from February 2018 to April 2022.

Knowledge Planet collects a fixed proportion of creator subscriptions as its platform commission. Before August 2019, the focal platform applied a fixed commission at 5%. In August 2019, the platform decided to raise the commission from 5% to 20% to all creators ("Policy 1" henceforth) and at the same time allow eligible creators to enjoy the original 5% commission if they got government-issued business license and passed the platform's policy compliance check in differential timing after August 2019 ("Policy 2" henceforth). By implementing Policy 2, the platform gave up the uniform pricing scheme where it charged the same fee to all creators, regardless of fixed 5% or fixed 20%, and started to conduct a third-degree price discrimination as it set different fees (i.e., different commissions between 5% and 20%) to different types of creators based on their eligibility.

In this paper, we focus on measuring the impact of the commission reduction through adopting Policy 2 using two causal inference methods. For the main analysis, we adapt the Staggered Difference-in-Differences (Staggered DID) method of Callaway and Sant'Anna (2021) and the Panel Data Difference-in-Differences with Matching (PanelMatch) method of Imai, Kim, and Wang (2023). Both Staggered DID and PanelMatch methods accommodate (1) potential outcomes by constructing appropriate control observations for each focal treated observation and (2) varying (staggered) treatment times. Further, because planets chose to adopt Policy 2 by incurring some time and monetary costs for the reduced commission, selection into treatment may limit the interpretation generalization of the effects we estimate. To overcome such selection issue, we utilize various ways of matching in PanelMatch method to rule out the noise from time-varying observables. We also carefully confirm the robustness of our results by checking necessary identification assumptions and utilizing different sets of model parameters.

Two major finds emerge from our analysis regarding (1) the overall effect of reduced commission through adopting Policy 2 and (2) the mechanisms through which the platform commission change can affect creators' pricing decision by a less competitive platform market structure and improved creators' quality provision. First, we find causal evidence for positive main effects of reduced commission on creators' prices. In our sample, adoption of Policy 2 is associated with an increase of 8%-13% in price, 31%-70% in the number of post, and 8-26% in the number of Q&A. The results are robust to using multiple methods including Staggered DID, PanelMatch with matching only on outcomes, and PanelMatch with Propensity Score Matching.

Second, we are able to test two mechanisms through which the platform commission reduction caused soaring subscription price set by creators. We first examine the the change of market structure due to a sudden platform commission increase through Policy 1 and find a significant reduction of the share of active planets. This result suggests that remaining active creators may gain market power to charge higher prices in a less competitive environment. We then examine the mediation role of content provision to creators' pricing decision. We propose that Policy 2 works as a tool to select creators who are more willing to provide content in larger number and of higher quality, so platform operator can subsidize these creators by reducing their commission. The creators of these treated planets then raised their prices to make up the cost from proliferating content production and/or higher-quality content provision. We observe that upon adopting Policy 2, creators who published their own original content and responded to subscribers' questions set up higher prices. We find similar results for treated creators who produced at least one content on average regardless of an original post or a personalized Q&A.

Our study contributes to a large growing body of research studying the effect of platform commission design. Current research studying commission rate primarily constructed sophisticated structural models and focused on how *reducing* the platform commission may affect demand and/or supply side entry of the market and mixed welfare changes in different empirical setting. For example, Barwick and Pathak (2015) run a counterfactual study in the real estate market which keeps agent commission rate fixed, but at lower levels. They show that under a 50% cut in commissions, there would be 40% fewer agents, social savings amount to 23% of industry revenue, and each agent sells 73% more houses (i.e., productivity increase). Robles-Garcia (2019) utilizes data from the mortgage market where brokers get a per-sale commission from lenders. She finds in a counterfactual of banning broker commission that such a ban leads to a decrease in consumer surplus because the ban reduces broker market power at the expense of increasing lender market power. Sullivan (2022) runs a counterfactual study with and without 15% commission caps on food delivery platforms in various metro areas and finds that the sum of caps's effects on the total welfare of restaurants, customers, and the platform is negative. On gaming platform Steam, the simulation analysis of Lu, Goldfarb, and Mehta (2023) implies that lowering the commission from the current 30% to 20% and to 12% will lead to more updates while fewer games will be released which results in an over-investment in Steam Early Access Program. Our paper provides new evidence of significantly increased subscription price and reduced content provision to consumers on a creator platform in a natural experiment setting where the focal platform *increased*[2] the commission from 5% to 20%.

There are very few literature utilizing causal inference models to study the effect of platform commission designs. In fact, research on platform commissions is limited to Li and Wang (2024). They use the data of commission caps for independent restaurants on food delivery platform and show that even though commission fee caps were intended to support independent restaurants, delivery platforms become less likely to recommend independent restaurants to consumers compared to chain restaurants after cities enact policies to cap platforms' commission fees, which leads to a decline in independent restaurants' orders and revenue. Li and Wang (2024) and our paper both study the third degree price discrimination in the platform commission design context, but our studies differ in the directions of which group of players are favored in the price discrimination. Specifically, in Li and Wang (2024), the price discrimination is in favor of the small businesses (independent restaurants)

[2]In an auction platform for vintage wines, Marra (2021) finds that doubling commission on sellers' side may not benefit bidders as entry and exit of sellers affect bidder's willingness to use the platform. Her research, again, utilizes a structural model and runs a counterfactual of increasing the commission, compared to our direct observation of the focal platform changing commission in the field.

against big businesses (chain restaurants) while in our case the price discrimination is in favor of big creators (business planets) against small creators (personal planets). In terms of methodology, Li and Wang (2024) adopt the canonical two-way fixed effects model with heteorogeneous covariates. Our study, benefting from a large and comprehensive panel data set, utilizes the most recent development in difference-in-differences models (e.g., De Chaisemartin and d'Haultfoeuille (2020), Goodman-Bacon (2021), Sun and Abraham (2021), Callaway and Sant'Anna (2021), etc.) and provides pioneering works investigating the effect of the staggered adoption of platform commission changes.

The rest of the paper proceeds as follows. In Section 2, we review the popular platform commission designs and highlight the commission policy changes on our focal platform. Section 3 then describes our data. In Section 4, we introduce our empirical approaches and show results. In Section 5, we discuss mechanisms through which a reduced commission caused increased price. Section 6 concludes.

2 Institutional Background

2.1 Platform Commission

Online marketplaces allow small businesses and independent creators to flourish. But these leading platforms do not provide their distribution for free. Instead, they charge various levels or combinations of commission to make their own profits.

Even though it is widely believed that Apple App Store chose to apply 30% commission that later became the the industry standard, the adoption of 30% commission can be traced back to the early 80s when two video game companies, Namco and Hudson Soft, attempted to have their games distributed on Nintendo's console. The 30% commission in fact combined 10% licensing fee and 20% manufacturing cost of video game cartridges[3], and this commission has then been used by numerous platforms afterwards until today.

[3]Source: *Epic's Battle With Apple and Google Actually Dates Back to Pac-Man.*

Given the industry standard commission at 30%, the commission designs by online platforms vary significantly. In Table 1, we list the commission design adopted by popular platforms in the United States and in China in five categories: app stores, live streaming, gaming, ride-hailing, and e-commerce. The vast majority of platforms decide a fixed, uniform commission for all players. For example, Twitch takes 50% of subscription fees and ad revenue from the residing creators. YouTube takes 45% from ad revenue and 30% from memberships. E-Commerce platforms like Amazon and TMALL specifies distinct uniform commission for certain category of goods. Another common practice adopted by these platforms is applying varying commissions based on certain criteria (e.g., cumulative revenue, different product types, etc.). This means a player on these platforms may be charged different amounts of commission throughout the time. For example, due to regulatory scrutiny, many platforms announced the beneficial program which lowered the commission as much as one half to alleviate the financial burden borne by small developers and creators. The most famous examples of this beneficial practice are varying commissions for lower income developers on Apple and Google app stores. In general, the commission of these small developers is 15% for their first million dollar revenue, and it jumps to 30% once the revenue exceeds one million dollars.

Steam, the leading PC gaming platforms adopted a quite different commission design from all others: starting from October 1, 2018, it has charged lower commission for game developers with higher revenue. In the announcement[4] introducing this unusual commission plan, the Steam Team highlighted the contribution from the big game studios which brings in the majority of the revenues and players to the platform: "*The value of a large network like Steam has many benefits that are contributed to and shared by all the participants. [...] It's always been apparent that successful games and their large audiences have a material impact on those network effects so making sure Steam recognizes and continues to be an attractive platform for those games is an important goal for all participants in the network.*"

[4]Steam's announcement on commission adjustment can be accessed at `https://steamcommunity.com/groups/steamworks/announcements/detail/1697191267930157838?content_only=true`

Table 1: Commission Design of Major Platforms across the U.S. and China

Type	Platform	Commission	
		Uniform	Varying
App Store	Apple App Store		15%, 30%
	Google Play		15%, 30%
Live Streaming	Twitch	50%	
	Douyin	50%	
Gaming	Epic	12%	
	Steam		20%, 25%, 30%
Ride-hailing	Uber	20%	
	CaoCao	23%	
	DiDi	20%	
E-Commerce	eBay	12.5%	
	Amazon	8∼20%	
	TMALL	0.5%∼10%	
	JD	2%∼10%	

Note: The commission interval for an e-commerce platform indicate the range of *uniform* commission for a category of goods sold on that platform. In other words, one specific category of goods is charged the same uniform commission, but commissions may be different across categories.

We pay special attention to Steam's commission design because, as we demonstrate in the following subsection, our focal platform Knowledge Planet adopted a similar commission design following the same logic of subsidizing big creators who paid more attention to their subscribers and produced more and higher-quality content. The platform operators believed a strong positive correlation between content production and subscriber base, and they were more than willing to reward these creators who brought more traffic to the platform. In an interview with the platform executives, we are informed that, in the past couple of years, top 10 superstar creators on Knowledge Planet contributed nearly one third of the total platform revenue through commissions, and top 100 creators contributed almost all of the total platform revenue. As discussed below about Knowledge Planet's business model, Knowledge Planet resembles many other creator platforms as a typical superstar-driven digital platform.

2.2 Knowledge Planet: the Focal Platform

Knowledge Planet is the leading digital platform in China which provides the virtual space for people who are specialized in certain fields to deliver their knowledge and expertise, to engage a community of fans, and to monetize through a simple yearly subscription-based business model. The biggest difference between Knowledge Planet and other content creation or MOOC platforms lies in its effort in launching, engaging, growing, and monetizing online communities by more than 8,500 creators who publish exclusive and highly personalized content in communities called "planet".

Figure 1: Q&A in Knowledge Planet, adapted from Han, Ryoo, and Zhao (2019)

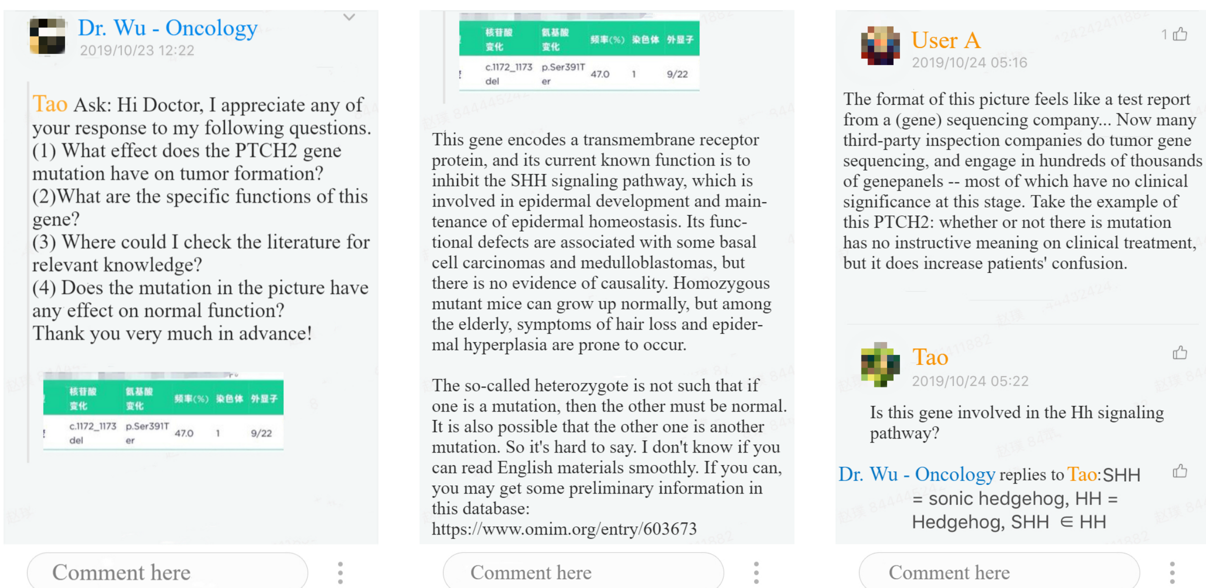

(a) Questions from a User *Tao* to Creator *Dr. Wu*

(b) Example of an Answer from Creator *Dr. Wu*

(c) Comments of Another User *A*, *Tao*, and *Dr. Wu*

A typical example which reflects the exclusive and personalized content creation in Knowledge Planet is Q&A. Any subscriber to a planet can ask questions directly in the main console to the creator who will then decide to cite the question and answer accordingly or not. Q&As enable the interactions between creators and their subscribers. As illustrated in Figure 1, the content creator Dr. Wu, an oncology specialist, wrote a post in details to answer four questions about gene mutation from a subscriber named Tao. In the comment section of this

Q&A post, another subscriber (i.e., "User A") got involved into the discussion between the creator Dr. Wu and the subscriber Tao, and provided his own viewpoint. A creator might receive many in-depth questions from all his or her subscribers, but he or she might not have time to respond to all of the questions. But as shown in Figure 1, a Q&A is a special type of post of higher quality since creators will spend particular time and effort to please their subscribers who may both promote the creators' knowledge or expertise delivery and advertise the planet.

The platform creates an advertisement-free environment and gives creators full pricing power to decide their subscription price charged to subscribers. As mentioned in last section, the platform charges a commission upon each transaction from either new subscriber joining or existing subscriber extending the ongoing subscription.

The analysis in this paper focuses on two platform commission policy changes on Knowledge Planet. From 2015 to August 2019, the platform collected a fixed 5% commission from all planets. In our previous study about the referral program on the same platform in Han, Ryoo, and Zhao (2019), we found out the optimal platform commission which maximized its profit targeted at 20%. Within a week after communicating this finding to the platform operators, the executives and the engineer team made the adjustment following our advice in August 2019. We call this simultaneous platform commission increase from 5% to 20% which applies to all planets as "Policy 1".

As with many other digital content creation platforms, Knowledge Planet relies on a limited number of superstar creators to generate the majority of its revenue. The platform executives follow the same logic as Steam to prioritize the superstar creators and their large groups of subscribers. Since August 2019 (the same month of Policy 1), any content creator who ran a planet with a government-issued business license and passed the policy compliance validation by platform operators could declare themselves as a "business" planet and enjoy the benefit of the original 5% commission. We call this commission decrease from 20% to 5% which only applies to part of the planets as "Policy 2". Note that planets were treated by

Policy 2 in a staggered manner up to the time when they submitted the business license to platform and up to the duration that platform completed the policy compliance validation. We further specify that all planets running as "personal" planets (even if they might have been running the planet with the business license) before August 2019, and those planets which either could not get business license from government or could not pass the platform policy compliance checks also as "personal" planets since August 2019.

3 Data and Sample Construction

To study the effects of platform commission design on creators' pricing strategy and productivity, we collect monthly panel data on Knowledge Planet from July 2016 to April 2022. Policy 1 took place at month 38 and the first creator adopted Policy 2 at month 39. We subset the data since February 2018 (i.e., the 20^{th} month in the original data) due to the data sparsity before this month. In fact, the overwhelming majority (74%) of all planets were founded since February 2018. This yields the data in total 51 months from February 2018 to April 2022. We end up with 7,186 distinct planets and 111,756 distinct planet-month observations in pricing related data and 224,618 distinct planet-month observations for productivity related data. In fact, our final sample of pricing variables is an unbalanced monthly panel of planet-month observations because the platform data warehouse stored the pricing-related variables only when transactions took place. We cannot rule out the possibility of one creator running multiple planets, but we assume that one creator makes pricing and productivity decisions independently across planets if he or she owns multiple ones.

When registering their planets on the platform, creators can utilize nine different tags to signal which area they are specialized. We then have nine types of knowledge or expertise, namely, Not Specified, Art, Economics, Education, Entertainment, Fashion, Health, Life, Science. There were quite a few planets with each of the last four types, and therefore we group them into "Others". This results in five types of content, and the distribution is reported in Table 2.

Table 2: Distribution of Planet Types

Type	Num.	Pct. (%)
Not Specified	1981	27.57
Art	1461	20.33
Economics	1957	27.23
Education	1071	14.90
Others	716	9.96

Table 3 provides key descriptive statistics for pricing related variables and productivity related variables. The revenue and price of each planet-month observation is recorded in the unit of Chinese Yuan. Since the exchange rate between Chinese Yuan (CNY) and U.S. Dollar (USD) in the data period from 2018 to early 2022 fluctuates a lot, we decide to stick with the original currency in Chinese Yuan. In the original data, we cannot directly observe the price change within each month, if any. We assume that not many creators would adjust their price in high frequencies within a month, so we calculate the average monthly price by dividing the total revenue in a month over the number of subscribers joining that planet in the same month.

Table 3: Descriptive Statistics of Planet-Month Observations

Variable	Obs.	Mean	St. Dev.	Min.	Med.	Max.
Pricing-related variables						
Revenue	111,756	9983.65	131250.24	149	1280	15551.5
Num. of Subscribers	111,756	33.11	154.20	1	7	66.0
Price	111,756	329.22	544.71	50	188	666.0
Productivity-related variables						
Num. of Post	224,618	62.25	348.93	0	6	115.0
Num. of Q&A	224,618	9.80	85.83	0	0	10.0

Notes: (i) The currency for revenue and price is Chinese Yuan (CNY). (ii) Price equals to Revenue divided by Num. of subscribers in each month for all planets. (iii) Num. of Q&A only records the number of questions that were ultimately answered by creators, but not includes the number of questions that were not replied by creators.

On average, a creator attracted 33 subscribers and generated 9,983 CNY revenue in a month. The average subscription price is 329.22 CNY, 1.07% of the average disposable income[5] of Chinese in 2019 (i.e., the year when both platform commission policy changes took place or started to take place) and 13.10% of the average expenditure on "Education, Culture, and Entertainment" category in the same year. This means that from the subscribers' perspective, they did invest a considerable amount of money on Knowledge Planet, assuming that they subscribed to only one planet. The productivity-related variables include the number of post which the creators published and the number of Q&A in response to subscribers. Note that the latter productivity variable excludes the questions left unanswered. We argue that the number of Q&A is not only a measure of productivity but also a measure of content quality, since creators have to apply their own expertise to some highly personalized questions from their subscribers, as shown in Figure 1. Creators on average produce much fewer Q&As than their own original posts when we compare the mean of these variables in Table 3, but any response to subscribers' questions marks creators' attention to their subscribers and entails considerable effort and cost.

4 Empirical Strategy and Results

The observational nature of our data have several identification challenges of the treatment effects. First, every planet is either treated or untreated in each time period, and we do not observe its counterfactual outcome in the unobserved condition. Second, planets adopt the Policy 2 in a staggered pattern with differential adoption timing, which requires careful computation and aggregation of treatment effects across planets and across time period (De Chaisemartin and d'Haultfoeuille 2020, Goodman-Bacon 2021, Sun and Abraham 2021, Callaway and Sant'Anna 2021, Imai, Kim, and Wang 2023, etc.). Third, the decision of when to adopt Policy 2 is endogenous and may be correlated with creators' differential expectations of benefit from Policy 2. Fourth, the creator of a given planet could have implemented other unobserved policies (e.g., discounts) concurrently with adopting Policy 2, and the observed

[5] According to National Bureau of Statistics of China, the average disposable income of Chinese in 2019 was 30,733 CNY, and the average expenditure on education, culture, and entertainment category was 2,513 CNY. Source: `https://www.gov.cn/xinwen/2020-01/17/content_5470095.htm`

effect cannot be separately attributed to the Policy 2.

To address these identification challenges, we utilize two empirical methods: Staggered DID based on Callaway and Sant'Anna (2021) and PanelMatch based on Imai, Kim, and Wang (2023). The technical details of these two methods are provided in Subsection 4.1, but we provide here an overview of how these methods can address the above-mentioned identification challenges.

The first challenge of unobserved potential outcomes is addressed by constructing a control group to predict the counterfactual potential outcomes of an adopting planet as if it did not adopt Policy 2. To do so, we treat the cohorts of planets that adopted Policy 2 in our data collection period as our focal treatment group. The control group for each adopting cohort comprises data from planets that either haven't adopted and never adopt Policy 2 for Staggered DID estimation and from planets that are matched based on observed characteristics in lagged time periods for PanelMatch estimation.

The second challenge of staggered adoption is addressed by using both Staggered DID and PanelMatch methods to compute treatment effects for each cohort of adopting planets separately and then aggregating them in to an overall average treatment effect on the treated (ATT), resolving issues such as negative weighting of treatment effects identified by Goodman-Bacon 2021, De Chaisemartin and d'Haultfoeuille 2020, Callaway and Sant'Anna 2021, etc. We also leverage two diagnostics to check for negative weights in treated units and checked for treatment effect homogeneity, as developed in Jakiela (2021). We show that the proportion of treated planets with negative weights is considerably low and the treatment effects are indeed nearly homogeneous across planets and time periods.

The third challenge of endogenous adoption timing is addressed primarily through the robust matching in PanelMatch estimation, where we construct the matched set either based purely on the lagged outcome variables or based on propensity score matching of lagged observed characteristics. The last challenge of unobserved confounders is addressed by controlling both planet fixed effects and time fixed effects.

4.1 Effect of Policy 2 on Creator's Pricing and Productivity

Before we extend our detailed discussion on Staggered DID method and PanelMatch method, we consider a canonical two-way fixed effects (TWFE) model as our benchmark. We let $i = 1, 2, ..., N$ denote planet and $t = 1, 2, ..., T$ time period (natural month). The TWFE model is specified as:

$$Y_{it} = \eta_i + \tau_t + \alpha G_{it} + \nu_{it}, \tag{1}$$

where Y_{it} denotes the observed outcome of interest, which is either log price of a planet i in month t, or the log number of post the creator of planet i published in month t, or the log number of questions that the creator of planet i answered in month t[6]. For each planet i, let G_{it}, a binary variable, equals to one if the creator of planet i adopted Policy 2 at month t. η_i is the unit fixed effect, τ_t is the time fixed effect, and ν_{it} are idiosyncratic, time-varying unobservables.

Table 4: Two-Way Fixed Effects Model Results

	(1) ln(Price)	(2) ln(Post + 1)	(3) ln(Q&A+1)
Policy 2	0.1532^{***}	0.6075^{***}	0.1176^{**}
	(0.0247)	(0.0567)	(0.0370)
Num. obs.	83782	182693	182693
R^2	0.8408	0.7346	0.7321
Planet FE	YES	YES	YES
Time FE	YES	YES	YES

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Standard errors are clustered at the planet level.

TWFE model does not guarantee a generally robust casual effect of participating in the treatment unless under the parallel trends assumption and under the homogeneity assumption (See Callaway (2023) for extensive discussion.). If we can still loosely interpret α as an overall average treatment effect in case with treatment effect heterogeneity (as in our case of dynamic treatment adoption), then we have preliminary evidence from Table 4 that Policy

[6]When taking the log of a count variable, we add a factor of one to avoid taking logs of zero for number of post and the number of Q&A.

2 led positive significant increase in all of our outcome variables.

Now in order to tackle with the four challenges stated at the start of this section, we introduce Staggered DID and PanelMatch methods and discuss how they can properly deal with each of the identification challenges.

4.1.1 Staggered DID Method

Policy 2 was in fact a third-degree price discrimination if we think of the platform commission as a fee charged to creators, because the platform sets different fees/commissions to personal and business planets. We estimate the treatment effects of Policy 2 using the Staggered DID framework proposed by Callaway and Sant'Anna (2021). This group-time treatment effect framework fits well in our setting because we have more than 1,000 units of planets that are eventually treated (i.e., adopting Policy 2) over a period of more than 30 months. Following the terminology of TWFE model, we let $Y_{it}(0)$ denote i's potential outcome at month t if planet i is untreated at month t and $Y_{it}(g)$ denote planet i's potential outcome at month t if planet i adopted Policy 2 at month g. This framework allows us for heterogeneous treatment effects with respect to specific treatment time.

We assume that once a planet had been treated, it remained treated and that the intensity of treatment was the same for all units[7]. We assume that there was no anticipation of treatment, so $Y_{it}(0) = Y_{it}(g)$ for all $t < g$. The relationship between observed and potential outcomes is as follows:

$$Y_{it} = Y_{it}(0) + \sum_{g=1}^{T}[Y_{it}(g) - Y_{it}(0)]G_{ig} \tag{2}$$

We then define the group-time treatment effect of interest as

$$ATT(g,t) = \mathrm{E}[Y_{it}(g) - Y_{it}(0)|G_{ig} = 1], \tag{3}$$

where $ATT(g,t)$ measures the average treatment effect at month t for the group of planets

[7] In Subsection 4.2.2, we confirm the homogeneous treatment effect across planets and time.

that adopted Policy 2 at month g. The group-time treatment effect framework is flexible enough to accommodate for the estimation of many different types of aggregated treatment effects. We estimate a weighted average $ATT(g,t)$ conditional on $e = t - g$ where e is the desired length of exposure, and in this paper, we are particularly interested in estimating the dynamic event study average treatment effect when e is up to 3.

Callaway and Sant'Anna (2021) show that the group-time treatment effects are identified from data on Y_{it} and G_{ig} as long as the parallel trend condition holds, in addition to standard independence and support conditions. Formally, the parallel trend assumption for our setting is stated as follows:

$$\begin{aligned}&\text{For each } g,\ h, \text{ and } t \text{ such that } g \leqslant t \leqslant h,\\ &\mathrm{E}[Y_{it}(0) - Y_{i,t-1}(0)|G_{ig} = 1] = \mathrm{E}[Y_{it}(0) - Y_{i,t-1}(0)|G_{ih} = 1].\end{aligned} \tag{4}$$

This parallel trend assumption requires that if treated groups had instead not been treated, then their outcome would follow the same pattern as groups that have not yet been treated. Note that this assumption cannot be directly tested because the left-hand side is not observed as we are not able to observe the counterfactual outcome for treated planets. However, we are able to test whether groups of planets have a different trend with respect to the amount of time left until the month these planets were treated in the pre-treatment period. We will return to this point in Subsection 4.2.1.

Each group-time treatment effect $ATT(g,t)$ is estimated by computing a weighted difference-in-differences estimate where the reference month is $g - 1$. The treated group includes the planets with $G_{ig} = 1$, and the control group includes the planets where $G_{ig} = 0$, **have not yet been treated** by month t, and **were never treated**. We refer to Callaway and Sant'Anna (2021) for technical details on estimation and inference. Here we implement the estimator in R using the package *did* (version 2.1.2) developed along the same reference paper.

We compute the treatment effect by length of exposure to the treatment. We define $\theta(e)$

as the weighted average of $ATT(g,t)$ for all t and g such that $e = t - g$:

$$\theta(e) = \frac{1}{\kappa_e} \sum_g \omega_g ATT(g, g+e), \tag{5}$$

where ω_g is proportional to the number of planets with treatment time g and κ_e normalizes the weights so they sum up to one. The parameter $\theta(e)$ can be interpreted as a dynamic event study. For $e < 0$, $\theta(e)$ captures the trend in outcomes for groups that are e periods away from the being treated by Policy 2 relative to the control groups of planets that are *not yet treated and never treated*. For $e > 0$, $\theta(e)$ captures the trend in outcomes of interest for groups of planets that are e months since being treated by Policy 2 relative to groups that are *not yet treated and never treated*.

One challenge in interpreting $\theta(e)$ is that the composition of planets varied with e. This is because not all planets are observed for e periods post-treatment. For example, we can only estimate the case of $\theta(e)$ with $e \leqslant 0$ for the (group of) planets which were treated by Policy 2 in the very last month, but cannot estimate $\theta(e)$ for any $e > 0$. To correct for composition changes of planets according to different e's and to rule out confounders over a longer time horizon, we estimate $\theta(e)$ by using only the planets that are observed 3-month post-treatment.

Column (1) in Table 5 presents the estimates of the 4-month ATT's of Policy 2 based on dynamic event study aggregation on price, the number of post, and the number of Q&A, which are valued at 12.69%, 30.72%, and 8.27%, respectively. Note that all these ATT's are all statistically significant at 5%. Figure 2 shows the treatment effects of Policy 2 on price, the number of post, and the number of Q&A over a full 3-month post-treatment period. We observe that most of the pre-treatment estimates contain zero in their confidence intervals, with only one violation in period $t = -1$ for the number of post estimation[8]. Another observation is that the treatment effect estimates for price are positive and keep climbing

[8]We give explanation of how this violation happens for this specific pre-treatment period and this specific outcome variable in Subsection 4.2.1.

Table 5: Treatment Effect of Policy 2 on Pricing and Productivity

Methodology	SDID	PanelMatch	
Matching	N/A (1)	Lag of Outcomes (2)	Propensity Score (3)
ln(Price)	0.1195* (0.0319)	0.0894* (0.0308)	0.0796* (0.0309)
ln(Post + 1)	0.2679* (0.0432)	0.3675* (0.0507)	0.5295* (0.0533)
ln(Q&A + 1)	0.0795* (0.0333)	0.1640* (0.0383)	0.2275* (0.0382)

Notes: (i) $^*p < 0.05$. (ii) Control group includes never-treated and not-yet-treated planets. (iii) For Staggered DID, we constraint the event-study aggregation of ATT's according to planets' length of exposure to Policy 2. We further balances the sample with respect to event time. This ensures that the composition of groups does not change when event time changes. (iv) In Staggered DID estimation, we calculate doubly robust standard errors, and in PanelMatch estimation, we use block bootstrapping procedure to calculate standard errors.

as length of exposure to Policy 2 extends, but the estimates for the number of post and the number of Q&A are only positive for the immediate few months and decline toward zero in later periods.

Combining the results in Table 5 and Figure 2, one of our starkest findings is that given a platform commission reduction by 15% (i.e., from 20% to 5%), the creators did not choose to cut their prices and passed through[9] the profit margin to their subscribers. Instead, they decided to increase the price by 12.69%, almost as much as they benefit from the reduced commission by 15%. In Section 5, we discuss two potential mechanisms through which this observation could possibly occur.

[9] According to Weyl and Fabinger (2013), the pass-through rate is the rise in price to consumers for each infinitesimal unit of specific tax imposed. In our case, we treat commission as the specific tax that the platform imposed on planets.

Figure 2: Treatment Effects of Policy 2 using Staggered DID Method

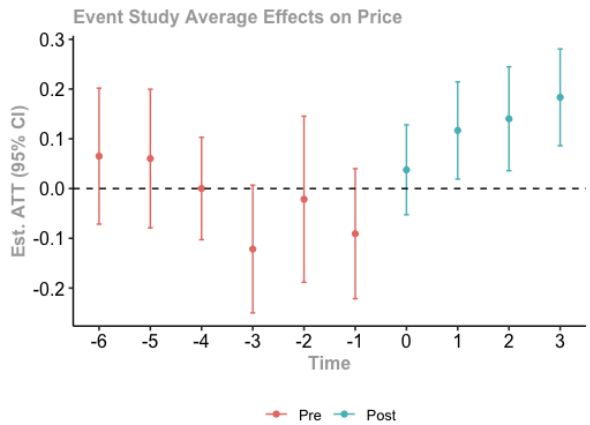

(a) ln(Price)

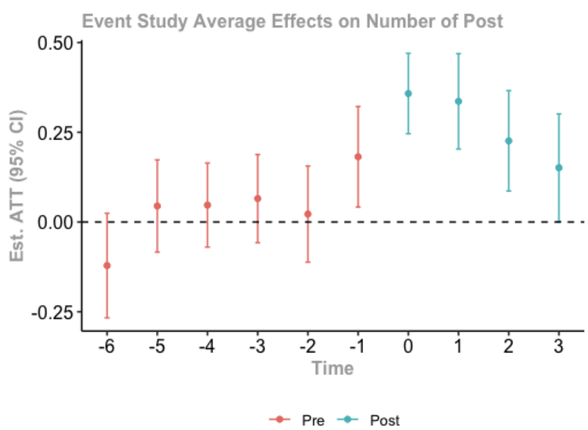

(b) ln(Post+1)

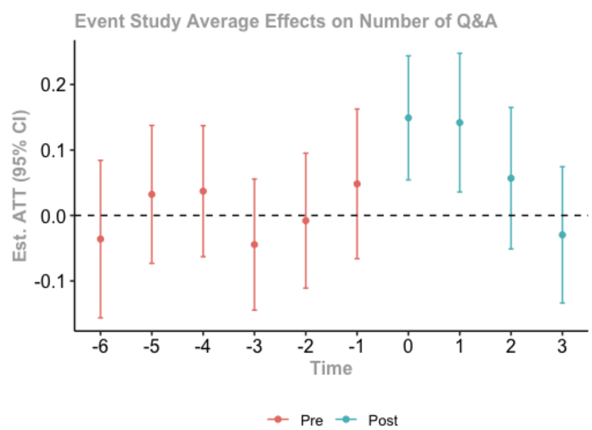

(c) ln(Questions Answered+1)

4.1.2 PanelMatch Method

Imai, Kim, and Wang (2023) propose a general matching method for causal inference with panel data where for each treated observation. We first find a set of control observations that have the identical treatment history up to the pre-specified number of time periods, and then refine the matched set of treated and control observations by adjusting for observed confounding via standard matching techniques so that the treated and matched control observations have similar covariate values. Then we apply the DID estimator to account for an underlying time trend.

Following the same terminology used to calibrate TWFE model and Staggered DID method, we let i denote planet, t denote time period (natural month), Y_{it} denote potential outcome and G_{it} denote Policy 2 adoption. In addition, let Z_{it} denote a vector of time-varying covariates, and let F and L denote nonnegative integer of leads and lags. The number of F represents the outcome Y_{it} measured at F time periods after the adoption of treatment, and it plays a similar role as $e = t - g$ in Staggered DID method. Specifying $F = 3$ allows us to examine a cumulative treatment effects up to 3 months since planets being treated by Policy 2. The number of L is to adjust for the extent to which past treatment status could be a confounder affecting the current outcome as well as the current treatment.

Once the two parameters L and F are specified, we define $\delta(F, L)$ as the ATT:

$$\begin{aligned}\delta(F, L) = \mathrm{E}[&Y_{i,t+F}(G_{it} = 1, G_{i,t-1} = 0, \{G_{i,t-l}\}_{l=2}^{L}) \\ &- Y_{i,t+F}(G_{it} = 0, G_{i,t-1} = 0, \{G_{i,t-l}\}_{l=2}^{L})|G_{it} = 1, G_{i,t-1} = 0],\end{aligned} \tag{6}$$

where the treated observations are those who adopt Policy 2, i.e., $G_{it} = 1$ and $G_{i,t-1} = 0$. The key identification assumption, i.e, the parallel trend assumption, using PanelMatch method is stated as follows:

$$\begin{aligned}&\mathrm{E}[Y_{i,t+F}(G_{it} = 0, G_{i,t-1} = 0, \{G_{i,t-l}\}_{l=2}^{L}) - Y_{i,t-1}|G_{it} = 1, G_{i,t-1} = 0, \{G_{i,t-l}, Y_{i,t-l}\}_{l=2}^{L}, \{\mathbf{Z}_{i,t-l}\}_{l=0}^{L}] \\ =&\mathrm{E}[Y_{i,t+F}(G_{it} = 0, G_{i,t-1} = 0, \{G_{i,t-l}\}_{l=2}^{L}) - Y_{i,t-1}|G_{it} = 0, G_{i,t-1} = 0, \{G_{i,t-l}, Y_{i,t-l}\}_{l=2}^{L}, \{\mathbf{Z}_{i,t-l}\}_{l=0}^{L}],\end{aligned} \tag{7}$$

where the conditioning set includes the treatment history $G_{i,t-l}$, the lagged outcomes $Y_{i,t-l}$ (except $Y_{i,t-1}$), and the covariate history $\mathbf{Z}_{i,t-l}$.

We then define the way to construct the matched set of control observations that share the identical treatment history from time $t - L$ to $t - 1$ for each treated observation indexed by planet-month pair (i, t). The matched set is defined mathematically as

$$\mathrm{M}_{it} = \{i' : i' \neq i, G_{i't} = 0, G_{i't'} = G_{i,t'} \text{ for all } t' = t-1, ..., t-L\} \tag{8}$$

for the treated observation (i, t) with $G_{it} = 1$ and $G_{i,t-1} = 0$. There are several caveats for practical purposes. First, the matching happens exactly on the treatment history in order to control for the carryover effects. Second, only observations from the same time period are included into the matched set in order to adjust for time-specific unobserved confounders. We further refine the matched set by using the distance measure based on the estimated propensity score. The propensity score is defined as the conditional probability of treatment assignment given pre-treatment covariates. To estimate the propensity score, for each time period, we consider all treated observations and their matched control observations, and then utilize the logistic regression model to fit the treatment assignment:

$$e_{it}(\{\mathbf{U}_{i,t-l}\}_{l=1}^{L}) = \Pr(G_{it} = 1 | \mathbf{U}_{i,t-1}, ..., \mathbf{U}_{i,t-L}) = \frac{1}{1 + \exp(-\sum_{l=1}^{L} \beta_l^T \mathbf{U}_{i,t-L})}, \tag{9}$$

where $\mathbf{U}_{it'} = (G_{it'}, \mathbf{Z}_{it'}^T)^T$. We also adjust for the lagged covariates by matching on the estimated estimated propensity score, which gives the following distance measure:

$$S_{it}(i') = |logit\{\hat{e}_{it}(\{\mathbf{U}_{i,t-l}\}_{l=1}^{L})\} - logit\{\hat{e}_{i't}(\{\mathbf{U}_{i',t-l}\}_{l=1}^{L})\}| \tag{10}$$

for each matched control observation $i' \in \mathrm{M}_{it}$, where $\hat{e}_{it}(\{\mathbf{U}_{i,t-l}\}_{l=1}^{L})$ is the estimated propensity score. Once the distance measure $S_{it}(i')$ is computed for all control observations in the matched set, we refine the matched set by selecting up to J most similar control observations and the refined matched set for the treated observation (i, t) is defined as:

$$\mathrm{M}_{it}^* = \{i' : i' \in \mathrm{M}_{it}, S_{it}(i') \leqslant S_{it}^{(J)}\}, \tag{11}$$

where $S_{it}^{(J)}$ is the J^{th}-order statistic of $S_{it}(i')$ among the control observations in the original matched set M_{it}. Put it simply, we choose a subset of control observations within the original matched set that are most similiar to the treated observation (i, t) in terms of both treatment history $G_{it'}$ and the observed confounders $\mathbf{Z}_{it'}$.

PanelMatch Results. To perform the PanelMatch analysis, for each adoption group g, we construct a matched set of treated planets and their associated control planets. When estimating the ATT's of Policy 2 on the outcome variables, we condition on 6 months of lag, that is, $L = 6$, and estimate the ATT up to 3 months after a planet adopts Policy 2, that is, $F = 1, 2, 3$. We further refine the matched set by selecting up to 10 most similar control observations to each focal treated observation.

Column (2) in Table 5 presents the estimated ATT's from PanelMatch method using only treatment history plus the lag of the outcome variable to construct the matched set. The 4-month ATT's of Policy 2 on price, the number of post, and the number of Q&A are valued at 9.35%, 44.41%, and 17.82%, respectively. Column (3) presents the estimated ATT's using propensity score matching based on treatment history, outcome variable, and time-varying covariates to construct the matched set. The 4-month ATT's of Policy 2 on price, the number of post, and the number of Q&A are valued at 8.20%, 69.81%, and 25.55%, respectively. Note that all these estimated ATT's are statistically significant at 5%, implying a rise in both pricing and content production.

Figure 3 illustrates the estimated treatment effects of Policy 2 on outcome variables at each relative length of exposure to Policy 2. The left three subfigures depict the treatment effects based on matching only by treatment history and the 6-month lags of outcome variables, while the right three subfigures show the treatment effects based on propensity score matching that also accommodates 6-month lags of time-varying covariates. The effects are positive and significant in the immediate month after adoption, and then attenuate in small magnitude in the later periods.

Figure 3: Treatment Effects of Policy 2 using PanelMatch Method

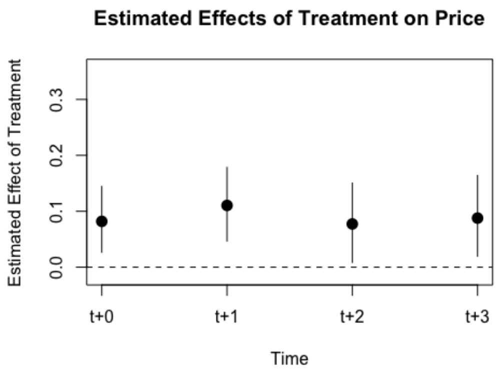

(a) Matching on Lag of Outcomes

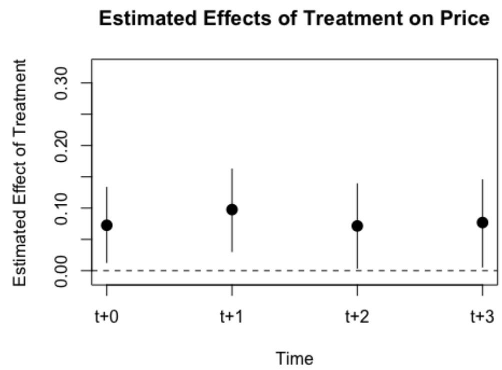

(b) Propensity Score Matching

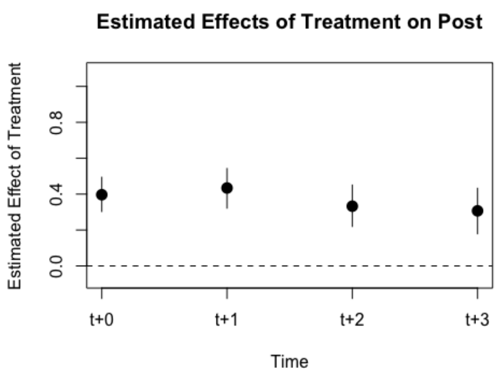

(c) Matching on Lag of Outcomes

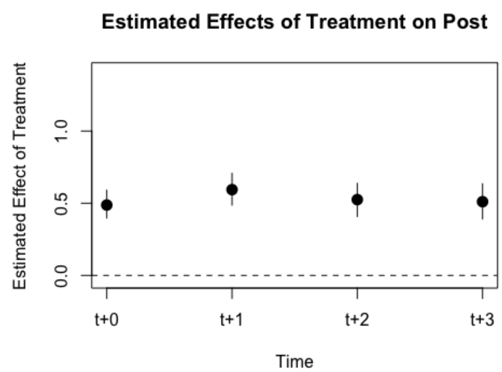

(d) Propensity Score Matching

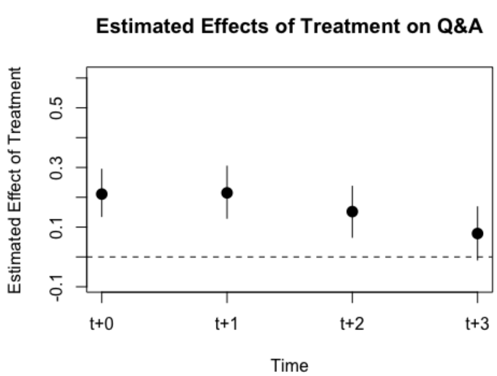

(e) Matching on Lag of Outcomes

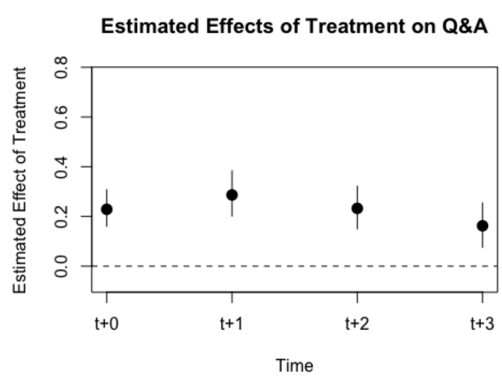

(f) Propensity Score Matching

4.2 Robustness Checks

4.2.1 Parallel Trend Assumption

Staggered DID. Figure 2 shows several important implications of parallel trend for Staggered DID regressions. First and foremost, for the treatment effect of Policy 2 on price and the number of Q&A, there are no differential trends in the pre-Policy 2 period between groups of planets that are six months away from being treated and other groups of planets that are not yet treated or never treated. This provides strong evidence in support of the (unconditional) parallel trend assumption discussed in Equation 4.

The pre-Policy 2 estimate at $t = -1$ for the number of post deviates significantly from 0. The creators who applied for the business license in order to enjoy the commission reduction from Policy 2 might have a strong belief that the final approval of Policy 2 was contingent on their labor outputs and/or content quality. In addition, creators could be uncertain about the processing time from applying to the local government for the business license to being eventually treated by Policy 2. During this processing time period, they might proliferate their content provision to persuade the platform operator to grant the reduced commission for Policy 2. In the real world, however, the platform operator, upon receiving their official business license, only conducted a policy compliance check to make sure these creators had not involved in any illegal practices such as politically sensitive content, gambling, uncertified medical devices, etc.

As a robustness check, we estimate the treatment effects by length of exposure allowing for one month anticipation. In other words, we allow creators to anticipate participating in Policy 2 one month prior to becoming a business planet which could affect their untreated potential outcomes. The results are reported in Appendix A. At the first glance of Figure 7, there is still no differential pretrends as in the case when we specify no anticipation for price and the number of Q&A, so the parallel trend assumption holds. Qualitatively, the results and patterns are similar to those without anticipation. But the treatment effect on price are smaller in magnitude and are estimated with less precision. The treatment effects on the number of post and the number of Q&A become more salient but still follows the decreasing pattern.

PanelMatch. Figure 4 illustrates the outcome and time-varying covariate balance in terms of the standardized mean differences due to matching over the pre-treatment time period. The solid black lines represent the balance of the lagged outcome variables whereas other colored lines in red, green, and blue show the balance of covariates. Again, the left three subfigures depict the pre-treatment matching of outcome variables based on matching only by treatment history and the 6-month lags of outcome variables, while the right three subfigures show the pre-treatment matching of both outcome variables and time-varying covariates based on propensity score matching that also accommodates 6-month lags of time-varying covariates. In particular, we specify different set of time-varying covariates when conducting Propensity Score Matching. For PanelMatch analysis on price, we control for 6-month lags of the number of subscribers, the number of post, and the number of Q&A. For PanelMatch analyses on the number of post and on the number of Q&A, we control for 6-month lags of price as well as the number of subscribers.

In all these cases corresponding to three outcome variables, we find that some degree of imbalance remains for the constructed matched sets only based on the treatment history and lagged outcome variables, even though the standardized mean differences are already as low as 0.5. In contrast, the improvement of balance due to Propensity Score Matching is substantial. In particular, Propensity Score Matching eliminates almost all imbalance in the number of post and the number of Q&A in pre-treatment period.

We also observe in Figure 4 that the standardized mean difference for the lagged outcomes, regardless of the matching methods, stays relatively constant over the entire pre-treatment period. This suggests that the parallel trend assumption for the PanelMatch estimations may be appropriate.

Figure 4: Outcome and Covariate Balance due to Matching over the Pre-Treatment Time Period using PanelMatch Method

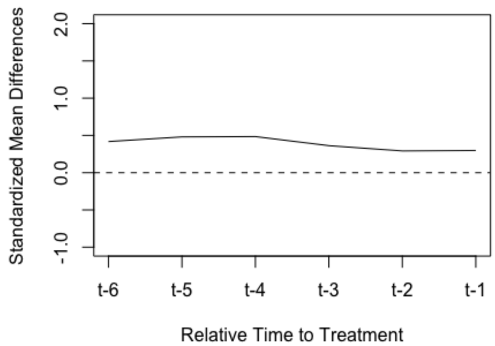

(a) ln(Price), Matching on Lag of Outcomes

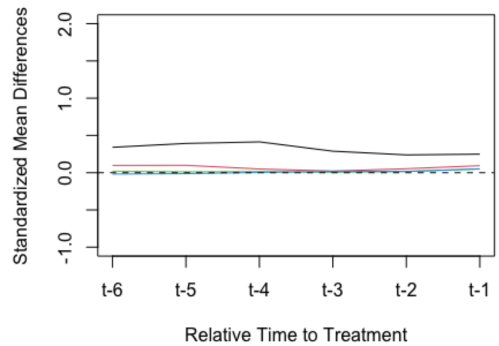

(b) ln(Price), Propensity Score Matching

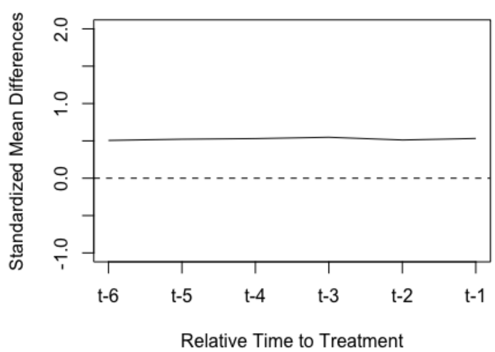

(c) ln(Post+1), Matching on Lag of Outcomes

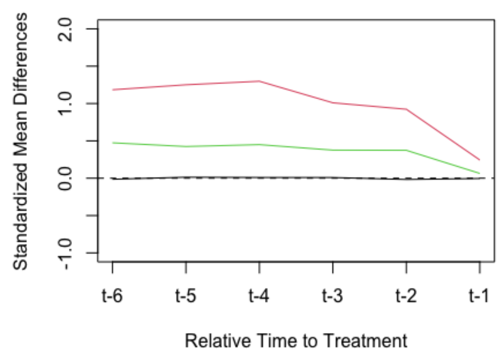

(d) ln(Post+1), Propensity Score Matching

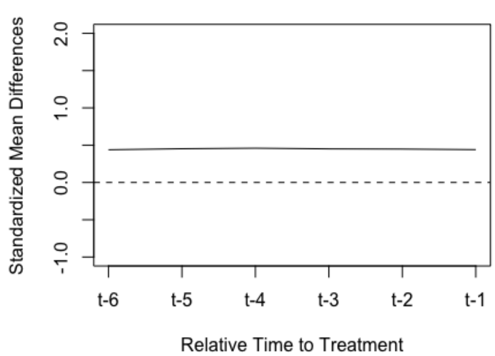

(e) ln(Q&A+1), Matching on Lag of Outcomes

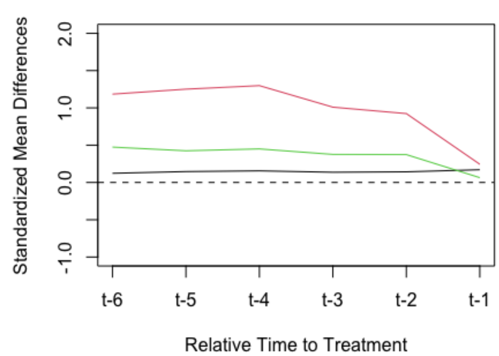

(f) ln(Q&A+1), Propensity Score Matching

4.2.2 More Evidences against Negative Weighting Concerns

A typical feature of Policy 2 is the differential timing of adoption for different planet. This may raise some concerns with the negative weighting in treated planets, particularly in (1) early adopter planets, since the planet-level treatment mean may be high, and (2) later months, since the month-level treatment mean may be high. According to Jakiela (2021), having negative weights on treated planets is not a big problem if we have (1) enough never-treated planets, and (2) enough pre-treatment data, and (3) the treatment effects are homogeneous across all planets.

In our panel dataset, more than 80% of planets (more than 75% of observations) that were never treated by Policy 2. For those planets that are treated by Policy 2, more than 27% of observations took place before the planet had been treated by Policy 2. Therefore, the first two criteria of diagnostics in Jakiela (2021) are met. We are now left to test the homogeneous treatment effects hypothesis. This hypothesis testing is particularly important when we consider price as our outcome variable, since we are about to explore mechanisms for pricing decision following Policy 2 using two-way fixed effects model in Section 5.

The diagnostic tests for homogeneous treatment effect in Jakiela (2021) is based on the mathematical relationship between the residuals of the outcome variable ($\tilde{Y}_{it}$) and the residuals of the treatment variable ($\tilde{G}_{it}$). Essentially, if there is no difference in slopes across treated and untreated observations in a regression of $\tilde{Y}_{it}$ on $\tilde{G}_{it}$, there is no evidence for heterogeneity and all is well. Mathematically, we have

$$\tilde{Y}_{it} = \delta_0 + \delta_1 \tilde{G}_{it} + \delta_2 G_{it} + \delta_3 \tilde{G}_{it} \times G_{it} + \varepsilon_{it}. \tag{12}$$

If we specify $log(Price)$ as our outcome variable, then the p-value for δ_3 (i.e., the interaction term between the treatment variable and its residual) is 0.805, which shows the change in slope between the treated group and the untreated group is statistically insignificant. So for regressions using $log(Price)$ as outcome variable, there is not enough evidence to reject the hypothesis that the slope between the treated groups and untreated groups are the same. In other words, we fail to reject the hypothesis that the treatment effects of Policy 2 on price is homogeneous. This also means we do not need to worry too much about the negatively weighted treated planet-month observations.

5 Mechanisms

The most counter-intuitive finding from our analysis of the effect of adopting Policy 2 is the increased price of treated planets given their reduced platform commission. To uncover the driving forces of this unusual price change, we consider two mechanisms. The first mechanism comes from the fact that at the same month when Policy 2 was announced (i.e., month 38), the platform raised its commission uniformly for all planets from 5% to 20%. We argue that this one-shot, uniform commission increase (i.e., Policy 1) drastically changed the market structure on the platform, making the number of active planets decline and enlarging market power for those creators of the remaining active planets to set up higher prices. The second mechanism comes from the fact that Policy 2 works as a platform tool to select those who are more willing to provide more and high-quality content, and the platform subsidizes these planets by reducing their commission. For these treated planets, the creators raised their prices to cover the cost from proliferating content production and/or higher-quality content provision. We illustrate the two mechanisms in Figure 5.

Figure 5: Mechanisms of Policy 1 and Policy 2 on Price Change

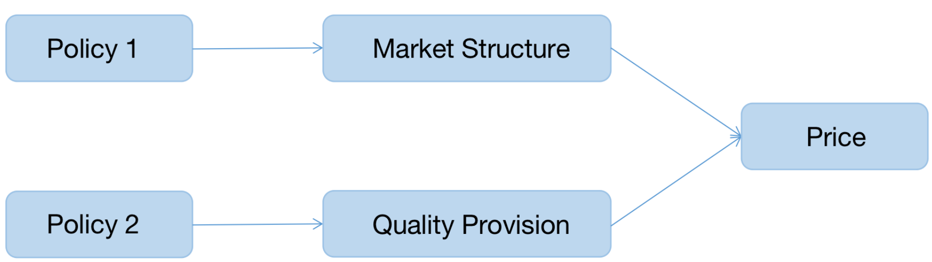

In this section, we first examine the effect of Policy 1 on the market structure and on creators' pricing decision, and we further examine whether those who were more willing to provide content in greater number and of higher quality did increase their prices following Policy 2.

5.1 Effect of Policy 1 on Price

We estimate the effect of Policy 1 using a interrupted time series model which accommodates panel data and captures the immediate level change of outcome variables due to Policy 1. Our model allows more data granularity than the classic interrupted time series model since we can directly observe the outcomes of interest for each planet in each month instead of the average of outcomes of all planets per month. Note that (1) we have longitudinal data of outcomes of interest for a considerably long time period over 51 months before and after Policy 1, and (2) all planets were affected at one specific time in August 2019. We include the planet fixed effects to adjust for unobserved, planet-specific and time-invariant confounders.

Let Y_{it} be the same set of observed outcome of interest as in Section 4.1. For each planet i, we let $time_{it}$ be the linear time trend that a specific month t passed from the start of the observation period. We let $policy1_{it}$ be the binary treatment variable which equals to 0 if planet i was before the implementation of Policy 1 at month t or equals to 1 if planet i was after the implementation of Policy 1 at month t. Finally, we let $since\_policy1_{it}$ indicate the number of month passed since the implementation of Policy 1 to month t, and we set its value to 0 if month t was before the implementation of Policy 1. Since Policy 1 was applied to all planets at the same time in August 2019 and no comparison group existed, the model takes the following form:

$$Y_{it} = \alpha_i + \beta_1 time_{it} + \beta_2 policy1_{it} + \beta_3 since\_policy1_{it} + \varepsilon_{it}, \tag{13}$$

where α_i denotes the planet fixed effects and ε_{it} is the error term. The coefficient of $time_{it}$ (i.e., β_1) indicates the trend of outcomes of interest before the Policy 1, or the pre-Policy 1 time trend. The coefficient of $policy1_{it}$ (i.e., β_2) indicates the the immediate level change of outcomes of interest after Policy 1, or the immediate effect of Policy 1. The coefficient of $since\_policy1_{it}$ (i.e., β_3) indicates the difference in linear trends after Policy 1, or the sustained effect of Policy 1.

To estimate the model specified in Equation 13, we only consider planets with data both before and after August 2019. Since Policy 2 started in the same month as Policy 1 took place, some planets took the staggered adoption of Policy 2 and enjoyed the platform commission reduction. We eliminate the observations of these planets since the month when they were treated by Policy 2. To eliminate the seasonality concerns, we restrict our model estimation within 6 months before and after August 2019 when Policy 1 applied to all planets and report the results based on this 6-month window as the main results in Table 6.

Table 6: Interrupted Time Series Model Results with 6-Month Window around Policy 1

	(1) ln(Price)	(2) ln(Post + 1)	(3) ln(Q&A + 1)
Time	0.0218^{***}	-0.0111^{*}	-0.0204^{***}
	(0.0020)	(0.0053)	(0.0034)
Policy 1	0.0251^{***}	-0.1439^{***}	0.0018
	(0.0073)	(0.0174)	(0.0109)
Since Policy 1	-0.0168^{***}	-0.0762^{***}	-0.0248^{***}
	(0.0027)	(0.0071)	(0.0045)
Num. obs.	23625	40807	40807
R^2	0.9160	0.7842	0.8266
Planet FE	YES	YES	YES

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Standard errors are clustered at the planet level.

Table 6 reports the main effect of Policy 1 on price, number of post, and number of Q&A. Here, we are mainly interested in the immediate effect of Policy 1 on price. Column (1) in Table 6 confirms that at the month of Policy 1 implementation, there was a significant level increase in price by 2.54%. In other words, creators pass through 16.95% of the 15% commission increase to their subscribers. We also examine the impact of a sudden commission increase due to Policy 1 on creators' content provision. Column (2) in Table 6 shows a significant immediate level decrease in the number of post by 13.40%. Given a sudden, uniform commission increase, creators made less money from subscription, and reduced their labor output as a response to this commission change. Column (3) in Table 6 implies that there was no significant immediate level change in the number of Q&A, because this outcome variable only records the number of questions that have been answered by creators in a specific month, which are presumably held stable given no obvious change in subscribers' behavior of asking questions.

As robustness checks, we run the same analysis in Equation 13 but (1) restrict the estimation within 3 months before and after Policy 1 took place, and (2) relax the estimation within 9 months before and after Policy 1 took place. Table 11 reports the coefficients of immediate level change due to Policy 1 for price, the number of post, and the number of Q&A under these two time windows around Policy 1. We obtain qualitatively and quantitatively similar results for the immediate effect of Policy 1 on price. When we only consider 3 months before and after Policy 1, there was a significant level increase in price by 3.98%. When we extend the time window to 9 months before and after Policy 1, there was a significant level change in price by 2.64%.

5.2 Effect of Policy 1 on Market Structure

To explore the effect of Policy 1 on market structure, we need to define some key measurement of planets' activity and their exit decisions. We first define a planet's activity (i.e., the condition in which planets were producing content on the platform) by checking whether or not the sum of the number of post and the number of Q&A is larger than 1. The state of being active or not for a planet is measured in a given month t. In other words, in a certain month t, if the creator of a planet produced at least one content regardless of a post or a Q&A, we think of this planet as being active.

Alike many other creator platforms, our focal platform Knowledge Planet does not have a clear way of determining whether or not a creator permanently exits. We thus define a planet's exit by observing its creator not producing any content for three consecutive months after its last month being active. For example, if a planet had been detected to be active in month 50 and its creator did not post anything or answer any questions from his or her subscribers since then, then this planet is considered to exit the platform in month 53. Since we have data of the exact time when a planet was registered on the platform, we are able to calculate the life cycle of each planet and their existence in each month. In fact, 41.32% of the planets ever existed in our dataset exit the platform by the end of the data collection period (i.e., month 70).

Now for each month t, we have both the number of active planets and the number of existing planets, we are able to construct the share of active planets by dividing the former measure over the latter one. Subfigure 6a illustrates the aggregate share of active planets over time before and after Policy 1 at month 38. In subfigure 6b, we plot the same measure separately according to their types of content. The share of active planets oscillates around 84% before Policy 1, but it plunges sharply right after Policy 1 to 72% at month 44, and climbs back to the pre-Policy 1 level nearly at the end of data collection period.

Figure 6: Share of Active Planets over Time

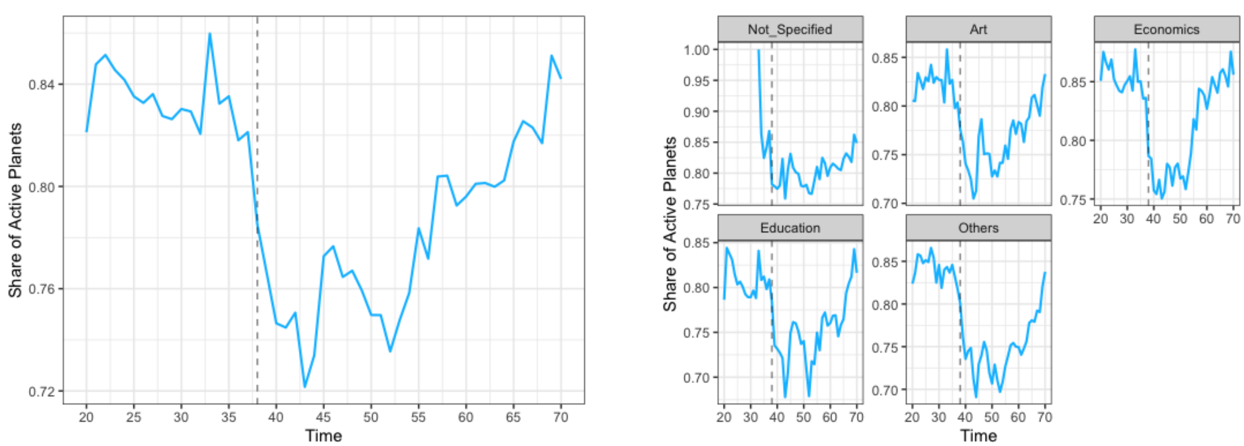

(a) at Aggregate Level

(b) w.r.t. Planet Type

Table 7 reports the OLS coefficients for the effect of Policy 1 on the share of active planets. In column (1), we consider the full 51 time periods from month 20 to month 70, and the share of active planets plummeted by 5.19% due to Policy 1. In column (2), we zoom into a shorter time window by considering only half year before and after Policy 1, and the share of active planets dropped by 7.82% over this period of time.

Table 7: Aggregate Effect of Policy 1 on Share of Active Planets

	(1) Full Time	(2) Half Year around Policy 1
(Intercept)	0.8340^{***}	0.8312^{***}
	(0.0064)	(0.0076)
Policy 1	-0.0533^{***}	-0.0814^{***}
	(0.0080)	(0.0104)
Num. obs.	51	13
R^2	0.4757	0.8489

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

Combining our analyses of the effect of Policy 1 on both price and the market structure, we conclude that the commission decrease due to Policy 1 is associated with a jump of creators' prices by 2.54% and a fall of the share of active planets by 7.82%. One explanation is that the drastically soaring platform commission by 15% significantly reduced the share of planets whose creators were still willing to provide content. For those who remained active on the platform after Policy 1, especially over the immediate couple of months since Policy 1, they were competing with much fewer creators for attracting and acquiring new subscribers, which gave them market power to charge higher prices without producing more content than before.

5.3 Effect of Policy 2 with Heterogeneous Content Provision

We expect that the focal platform has utilized Policy 2 as a selection process for those who were willing to produce content in greater number and/or higher quality. In Section 4.1, we have shown robust results of how the staggered adoption of Policy 2 led to significant increase in content provision of both post and Q&A. We now examine how different quantities and qualities of content provision can mediate the effect of Policy 2 on price.

We use the following two binary variables to measure the quantity and quality of content provision for planet i in month t: (1) $Post\_and\_Q\&A_{it}$ that equals to 1 if the product of the number of post and the number of Q&A in that specific month t is larger than 0; (2) $Content\_30$ that equals to 1 if the sum of the number of post and the number of Q&A is larger than 30. $Post\_and\_Q\&A_{it}$ shows whether or not a creator of planet i indeed produced his or her own post *AND* answer questions from his or her subscribers. $Content\_30$ measures whether or not a creator could, on average, produce at least one content regardless of post or Q&A on a daily basis.

Table 8: Effects of Policy 2 on Price w.r.t. Heterogeneous Content Provision

	(1)	(2)
Policy 2	0.1260***	0.1274***
	(0.0277)	(0.0283)
Post and Q&A	0.0112	
	(0.0076)	
Policy 2×Post and Q&A	0.0437*	
	(0.0196)	
Content 30		0.0143
		(0.0096)
Policy 2×Content 30		0.0383†
		(0.0214)
Num. obs.	83782	83782
R^2	0.8409	0.8408
Planet FE	YES	YES
Time FE	YES	YES

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$; †$p < 0.1$. Standard errors are clustered at the planet level.

Table 8 reports the TWFE estimation results of the effect of Policy 2 on price in which these three measures work as mediators. Column (1) in Table 8 suggests that upon adopting Policy 2, if the creator of a planet provided both original post and personalized response to specific questions from subscribers, his or her price was 4.47% higher compared to those who either failed to provide one type of content or did not produce content at all.

From Column (2) in Table 8, we find evidence that for those adopting Policy 2, if a creator of a treated planet produced content on a daily basis so his or her average number of monthly content was more than 30, this creator charged 3.90% higher than those who failed to reach this monthly production threshold.

These results are consistent with the expectation of general subscribers on a creator platform. Subscribers pay for a creator's own ideas, thoughts, and experiences and they value the personalized responses that a creator give to their unique questions.

5.4 Effect of Policy 2 on Price given Competition among Planets

In Subsection 5.2, we associate Policy 1 with the market structure and provide evidence that Policy 1 intensified the competition among existing planets by significantly reducing the share of active planets, which ultimately led to higher price. One may wonder whether or not Policy 2 is associated with any potential market structure change that caused price fluctuations for certain creators. In this subsection, we examine the potential mediation effect of Policy 2 on price through various measures of market competitiveness.

Due to differential timing of Policy 2 treatment, we need to define properly the market and related competition metrics. We define a market $mkt_t^{\mathcal{T}}$ by both time t and type of content $\mathcal{T}$, where $t \in \{38, 39, ..., 69, 70\}$. For example, creators in market $mkt_{t=50}^{\mathcal{T}=\text{Economics}}$ only compete with others in the same month $t = 50$ and the same type $\mathcal{T}$ = Economics. It is very likely to be time- and resource- infeasible for a focal creator to explore all other planets and keep track of their content provision to figure out who the potential competitors are. We thus assume that creators leverages the prices of other planets in the same market as a tool to determine potential competitors, because price in creator platforms often signals quality provision.

We use the following two variables as proxies for the competition among planets: (1) the number of planets in the same market with similar price, and (2) the number of subscribers in planets in the same market with similar price. The similarity of prices between a focal planet and other planets is defined by a radius of 5 CNY. That is, if the difference between the price $P_i^{mkt_t^{\mathcal{T}}}$ of a focal planet i and the price $P_j^{mkt_t^{\mathcal{T}}}$ of another planet j is no larger than 5, then we consider planet j as a competitor of planet i in market $mkt_t^{\mathcal{T}}$.

Table 9 reports the TWFE estimation results when using two competition proxies and

Table 9: Heterogeneous Effects of Policy 2 on Price w.r.t. Competition among Planets

	(1)	(2)
Policy 2	0.1711***	0.1216**
	(0.0407)	(0.0446)
ln(Num. Competitor + 1)	−0.1007***	
	(0.0077)	
Policy2×ln(Num. Competitor + 1)	−0.0124	
	(0.0132)	
ln(Subscriber of Competitor + 1)		−0.0871***
		(0.0042)
Policy2×ln(Subscriber of Competitor + 1)		0.0026
		(0.0078)
Num. obs.	83782	83782
R^2	0.8453	0.8501
Planet FE	YES	YES
Time FE	YES	YES

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Standard errors are clustered at the planet level.

their interaction terms with Policy 2 as independent variables. The raw number of competitors and the number of subscribers in competitor planets significantly reduced the price by 9.58% and 8.34%, respectively. However, we do not find evidence of mediation effect of either competition proxies on the effect of Policy 2 on price according to the insignificant coefficients for the interaction terms. Table 12 in Appendix C reports consistent results of the insignificant coefficients of the interaction terms between Policy 2 and our competition proxies, when $|P_i^{mkt_t^{\mathcal{T}}} - P_j^{mkt_t^{\mathcal{T}}}|$ is no larger than 10 or 20 CNY for $i \neq j$. We therefore rule out the potential mechanism that the competition among planets mediates the effect of Policy 2 on price.

6 Conclusion

This research analyzes the impact of a staggered platform commission reduction on the pricing strategy and productivity of creators, using the data of a leading Chinese creator platform, Knowledge Planet. This commission change policy reflected platform executives' intention to maximize the platform profits while keeping large subscriber base brought by the superstar creators.

By employing two robust methodological frameworks of Staggered DID and PanelMatch, our investigation has illuminated the nuanced interplay between platform commissions and the economic behaviors of content creators, specifically in terms of pricing strategies and productivity. The results of our analysis show that the adoption of a reduced commission significantly increased creator price, the number of post, and the number of Q&A by an average of 8%-13%, 31%-70%, and 8-26% in the four months after adoption, respectively. Although these ranges are wide and result from using multiple estimations, we demonstrate that the positive effects are substantial and robust using multiple methods and analyses. We further demonstrate two mechanisms for a price increase given a commission reduction policy through (1) decreasing competition among planets due to market structure change and (2) the selective nature of this policy that subsidized those who were more willing to provide content in larger number and higher quality.

Framing these stark results in the third-degree price discrimination literature in two-sided market, we find no theoretical models that accommodate the case where the platform sets up higher fees/commissions to small businesses (in terms of revenue, participant, or market share) while lowers the fees/commissions to big businesses. Instead, recent industrial organization literature have put effort to solve for the welfare analysis of the case where platform favors the small businesses by lower their fees/commission. For example, Wang and Wright (2017) argue that ad valorem fees[10] are a way for platforms to efficiently price discriminate heterogeneous merchants. Motivated by the app store controversies, Bhargava, Wang, and Zhang (2022) show that a small-business oriented (SBO) differential revenue sharing design can increase total welfare and outputs on the platform. In addition, although smaller producers almost always benefit from the shift in revenue sharing design, spillover effects can

[10] Ad valorem fees increase proportionally with transaction prices. In the platform commission context, this means the platform charges higher commissions for higher transaction values.

also make large producers better off under some conditions. de Cornière, Mantovani, and Shekhar (2023) find out that, driven by the existence of network effects across buyers and sellers, third-degree price discrimination set by the platform on the seller side can increase participation on both sides, increase total welfare, and may result in a Pareto improvement with all sellers being better-off (regardless of the different levels of commission charged on them) than under uniform pricing. Our paper calls for theoretical modelling which fills the gap in the literature to consider the case where the price discrimination works in favor of bigger producers/sellers/creators and against the smaller counterparts.

Throughout this paper, we also ignore the network effects across creators and subscribers as described in Bhargava, Wang, and Zhang (2022) and de Cornière, Mantovani, and Shekhar (2023). We also only emphasize short-run changes due to platform commission changes and preclude any long-term market structure conducts such as competition among planets. In turn, this research suggests future avenues for researchers to create better modelling and analytical solutions to solve these challenges for all players in the platform ecosystem.

In sum, our findings demonstrate that the introduction of a differentiated commission structure, moving from a uniform rate to a varying rates, significantly enhanced the creators' ability to monetize their content more effectively if they were more willing to provide more content and content of higher-quality. This was evidenced by an uptick in content production volumes and a strategic elevation in subscription prices, contributing to a healthier platform ecosystem where quality content generation is incentivized, and financial rewards are more directly aligned with creators' contributions to the platform.

This research contributes to the burgeoning discourse on the economic implications of platform policies, providing empirical evidence that challenges prevailing assumptions about the one-size-fits-all approach to platform commission rates. Our study highlights the potential of tailored commission structures to serve as a catalyst for both enhanced creator productivity and increased platform revenue.

References

Barwick, Panle Jia and Parag A Pathak (2015), "The costs of free entry: an empirical study of real estate agents in Greater Boston," *The RAND Journal of Economics*, 46 (1), 103–145.

Bhargava, Hemant K, Kitty Wang, and Xingyue Zhang (2022), "Fending off critics of platform power with differential revenue sharing: Doing well by doing good?," *Management Science*, 68 (11), 8249–8260.

Callaway, Brantly (2023), "Difference-in-differences for policy evaluation," *Handbook of Labor, Human Resources and Population Economics*, pages 1–61.

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

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2020), "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 110 (9), 2964–2996.

de Cornière, Alexandre, Andrea Mantovani, and Shiva Shekhar (2023), "Third-degree price discrimination in two-sided markets," *Working paper.*

Goodman-Bacon, Andrew (2021), "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225 (2), 254–277.

Han, Xintong, Jun Hyun (Joseph) Ryoo, and Pu Zhao (2019), "Referral marketing in the creator economy: an empirical study of referral programs by content creators," *NET Institute Working Paper*, 19 (03).

Imai, Kosuke, In Song Kim, and Erik H Wang (2023), "Matching methods for causal inference with time-series cross-sectional data," *American Journal of Political Science*, 67 (3), 587–605.

Jakiela, Pamela (2021), "Simple diagnostics for two-way fixed effects," *arXiv preprint arXiv:2103.13229.*

Li, Zhuoxin and Gang Wang (2024), "Regulating Powerful Platforms: Evidence from Commission Fee Caps," *Information Systems Research.*

Lu, Wei, Avi Goldfarb, and Nitin Mehta (2023), "Product Development and Platform Design," *Working paper.*

Marra, Marleen (2021), "Estimating an auction platform game with two-sided entry," *Working paper.*

Robles-Garcia, Claudia (2019), "Competition and incentives in mortgage markets: The role of brokers," *Working paper.*

Sullivan, Michael "Price controls in a multi-sided market," Technical report (2022).

Sun, Liyang and Sarah Abraham (2021), “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199.

Wang, Zhu and Julian Wright (2017), “Ad valorem platform fees, indirect taxes, and efficient price discrimination,” *The RAND Journal of Economics*, 48 (2), 467–484.

Weyl, E Glen and Michal Fabinger (2013), “Pass-through as an economic tool: Principles of incidence under imperfect competition,” *Journal of Political Economy*, 121 (3), 528–583.

Appendices

A Robustness Checks of Staggered DID with Anticipation of 1 Time Period

Table 10: Treatment Effect of Policy 2 with Anticipation of 1 Month

Methodology	SDID
ln(Price)	0.0336 (0.0519)
ln(Post+1)	0.4947* (0.0630)
ln(Q&A+1)	0.1417* (0.0459)

Notes: (i) $^*p < 0.05$. (ii) Control group includes never-treated and not-yet-treated planets. (iii) We constraint the event-study aggregation of ATT's according to planets' length of exposure to Policy 2. We further balances the sample with respect to event time. (iv) We calculate doubly robust standard errors.

Figure 7: Treatment Effects of Policy 2 with Anticipation of 1 Month

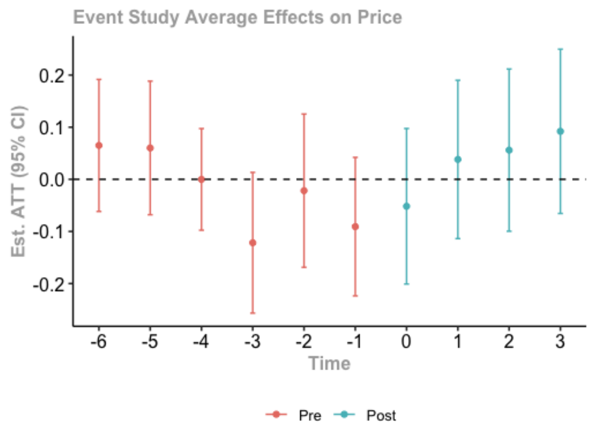

(a) ln(Price)

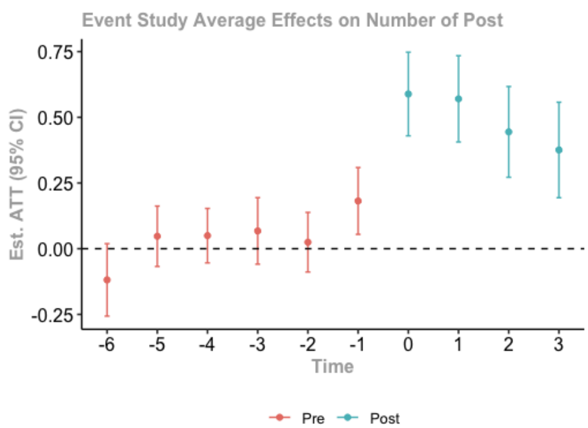

(b) ln(Post+1)

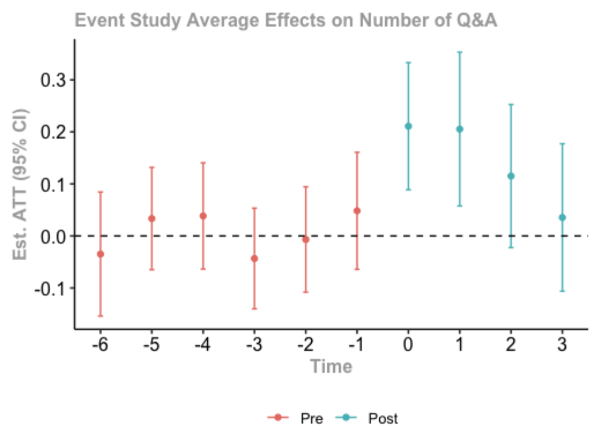

(c) ln(Q&A+1)

B Robustness Checks of Estimation for Studying the Impact of Policy 1: Using Different Time Windows around Policy 1

Table 11: Interrupted Time Series Model Results with Various Time Windows

	3-Month Window around Policy 1			9-Month Window around Policy 1		
	(1) ln(Price)	(2) ln(Post + 1)	(3) ln(Q&A + 1)	(4) ln(Price)	(5) ln(Post + 1)	(6) ln(Q&A + 1)
Time	0.0265***	−0.0826***	−0.0435***	0.0178***	−0.0180***	−0.0124***
	(0.0032)	(0.0091)	(0.0065)	(0.0015)	(0.0043)	(0.0026)
Policy 1	0.0390***	−0.0690***	−0.0272*	0.0261**	−0.2425***	−0.0696***
	(0.0072)	(0.0167)	(0.0109)	(0.0080)	(0.0181)	(0.0108)
Since Policy 1	−0.0249***	−0.0124	0.0238**	−0.0133***	−0.0342***	−0.0150***
	(0.0045)	(0.0113)	(0.0078)	(0.0022)	(0.0056)	(0.0033)
Num. obs.	13409	22805	22805	32594	57525	57525
R^2	0.9492	0.8602	0.8846	0.8904	0.7303	0.7909
Planet FE	YES	YES	YES	YES	YES	YES

Notes: (i) $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. (ii) For all six models, we consider planets which have data before and after August 2019 when the platform increases its commission from 5% to 20% simultaneously to every planet, and we constrain the time window to be either 3 or 9 months before and after August 2019. For planets whose platform commission decrease from 20% to 5% after August 2019, we eliminate their observations since the month of platform commission reduction. (iii) Standard errors are clustered at the planet level.

C Robustness Check of the Effect of Policy 2 given Competition among Planets

Table 12: Robustness Check: Heterogeneous Effects of Policy 2 w.r.t. Planet Competition

	Price Range is 10		Price Range is 20	
	(1)	(2)	(3)	(4)
Policy 2	0.1853***	0.0936†	0.1976***	0.0538
	(0.0497)	(0.0527)	(0.0578)	(0.0605)
ln(Num. Competitor + 1)	−0.1693***		−0.2472***	
	(0.0099)		(0.0124)	
Policy2×ln(Num. Competitor + 1)	−0.0177		−0.0218	
	(0.0154)		(0.0163)	
ln(Subscriber of Competitor + 1)		−0.1316***		−0.1760***
		(0.0053)		(0.0063)
Policy2×ln(Subscriber of Competitor + 1)		0.0060		0.0102
		(0.0088)		(0.0093)
Num. obs.	83782	83782	83782	83782
R^2	0.8499	0.8555	0.8553	0.8597
Planet FE	YES	YES	YES	YES
Time FE	YES	YES	YES	YES

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$. Standard errors are clustered at the planet level.