Blocking the Giants: Theory and Evidence from the Great Firewall

Ruiqi (Rachel) Sun*†

Job Market Paper

[Click here for latest version]

9 December 2024

Abstract

Bans on digital products from foreign tech giants have become a policy staple. Yet the welfare implications of these bans are unknown. Policy-targeted products are typically large and highly differentiated, leading to rich substitution patterns that are essential for welfare analysis. I empirically reveal these substitution patterns through an event study of India's ban on Chinese apps: pairwise elasticities of substitution are correlated with pairwise similarities between product descriptions. Motivated by this finding, I develop a general equilibrium model in which granular products with heterogeneous productivities and hedonic attributes engage in Bertrand oligopolistic competition. This model is purpose-built to feature the pairwise elasticities of substitution, which map to the product similarities constructed from a series of large language model tasks. I use the estimated model to evaluate the impact of China's Great Firewall (GFW) policy. I find that while the GFW increased Chinese real incomes, the loss in leisure utility from using inferior apps overwhelms the benefits, resulting in a net welfare loss of 7.5%. This framework is easily implemented and widely applicable to research questions where pairwise substitutions are essential and rich text is available.

Keywords: digital trade policy; granular trade; competition; LLMs

JEL codes: F12, F13, F14, F52

*I am deeply indebted to my supervisors Daniel Trefler and Xiaodong Zhu, for their continued guidance and support, as well as Heski Bar-Isaac, Kevin Lim, Matthew Mitchell, and Jeffrey Sun for detailed and insightful comments. I also benefit from the discussions with Bernardo Blum, Laurent Cavenile, Mitsuru Igami, Guangbin Hong, Myeongwan Kim, Sijie Lin, Siyuan Liu, Peter Morrow, Regina Seibel, and Irisa Zhou. All errors are my own.

†Rotman School of Management, University of Toronto. Email: ruiqi.sun@rotman.utoronto.ca.

1. Introduction

Trade policies targeting specific products or firms, rather than entire industries or countries, have become increasingly prevalent worldwide. These regulations are often a response to the escalation of geopolitical tensions, where technology, trade, and innovation intersect with national security concerns and global power dynamics (Bradford, 2023). For example, China's Great Firewall (GFW) policy, implemented in 2009, effectively excluded three of the U.S. Big Four – Amazon, Meta, and Google – from its market. Concurrently, we have witnessed the rise of China's own technology giants: Alibaba, Tencent, and Baidu. Each of these Chinese firms mirrors a corresponding American counterpart: Baidu serves as a substitute for Google, Tencent for Meta, and Alibaba for Amazon. Do such trade policies targeted at specific large firms causally foster the growth of local substitutes in China? What are the welfare implications of these policies? The recent implementation of similar bans by countries like the U.S., Canada, India, and Brazil makes finding answers to these questions increasingly important.[1]

However, answering these questions is not trivial. Banned digital products typically have a vast user base and a diverse mix of hedonic attributes, leading to complex substitution patterns with the competitors. This complexity makes detailed knowledge of pairwise elasticities of substitution essential for welfare analysis. I document a novel fact: pairwise elasticities of substitution are correlated with pairwise similarities in product descriptions. Building on this finding, I anchor my analysis around a general equilibrium model where granular, single-product firms with heterogeneous productivities and hedonic attributes engage in Bertrand oligopolistic competition. Millions of elasticity parameters are estimated based on cosine similarity between large language model (LLM) embeddings. This methodology, both theoretically and empirically, provides a practical approach to capture rich demand-side heterogeneity (hedonic attributes) alongside supply-side heterogeneity (Melitz productivity) and highlights its relevance for the welfare analysis of policy interventions targeted at granular firms.

In Section 2, I begin with an event study to empirically document the substitution patterns following a ban. I examine the impact of an unanticipated ban on Chinese apps in India and find that apps not directly affected by the ban expanded as a result of reduced competition. For each app, I construct a measure of similarity to the banned apps. Apps with the highest similarities to banned apps had a substantial increase in market share, while those with the lowest similarities exhibited almost no increase. This finding provides strong evidence of heterogenous responses to a trade shock, which is a feature absent in workhorse models (e.g., Melitz (2003)). Furthermore, the similarity

[1] Appendix Table A1 lists the ten most banned apps and demonstrates the worldwide prevalence of bans as a new trade tool.

measure – constructed using textual product description and LLMs – effectively reflects the elasticity of substitution. The estimates from this event study are used later for structural estimation of the quantitative model.

Motivated by the empirical finding, I develop a general equilibrium model in Section 3. The model captures the rich structure of elasticity of substitution while maintains high tractability. In the model, there is a discrete number of potential entrants heterogenous along two dimensions: hedonic attributes on demand side and productivity on supply side. Inspired by the work of Pellegrino (2023) in industrial organization, I introduce a Translog expenditure function, as proposed by Diewert (1976), where the key parameters emerge endogenously from a micro-foundation of hedonic demand. I integrate the demand system into a granular trade model, building on Eaton, Kortum, and Sotelo (2012) and Gaubert and Itskhoki (2021), and solve for the static equilibrium. In equilibrium, the potential entrants, each born with a mix of hedonic attributes and a productivity, first decide on entry, and upon entry, choose prices under Bertrand oligopolistic competition in the second stage. [2]

The vast number of demand parameters can be estimated with a parsimonious set of restrictions. Standard procedures for estimating elasticities within nested CES frameworks, such as Berry, Levinsohn, and Pakes (1995), require extensive market-level data, which is often unavailable in international trade contexts. In this model, the challenge is even more pronounced as the nesting occurs at the product level. However, the demand system here yields pairwise elasticity of substitution as a function of similarities in hedonic attributes. I propose an estimation strategy that uses LLM embeddings of textual product descriptions as proxies for the hedonic attributes. This approach transforms the complex estimation of pairwise elasticities into a series of product-specific LLM tasks, alongside the identification of one single scale parameter. More importantly, this approach provides a structural interpretation for the output of language models, a tool increasingly used in recent empirical studies.

To validate the model as well as the estimation strategy, I conduct a quantitative analysis using a dataset on mobile apps. I choose this sector and dataset for four reasons. First, the mobile app sector yields complex yet intuitive substitution patterns to the consumers. For example, Google Chrome is highly substitutable with Baidu but is less substitutable with Weibo, which is primarily like X/Twitter but also incorporated search capabilities. Second, I have rich textual data and trade data at product variety level in this

[2] While the baseline model is applicable to any sector, I tailor it to a quantitative model for mobile apps in Section 5.1, where prices correspond to ad rates in this application.

sector, which is essential for my quantitative analysis.[3] Third, natural experiments like the India's app ban and China's GFW policy provide exogenous variation to identify key parameters and validate model predictions. Fourth, the primary players in this sector are superstar firms central to recent trade disputes. Among the top ten largest companies in the world, eight specialize in software services or hardware support and they are all involved in international bans, which underscores the policy relevance of this paper.[4]

Section 4 details the dataset and presents the reduced-form evidence for the impact of the GFW. The dataset covers almost the universal of apps available on the iOS store and consists of three key components. The first component is the estimated hedonic attributes of each app. Each app has a detailed product description displayed in the app store. Using this raw data, I generate synthetic text with focused on product design and convert it into embeddings, which are 3072-dimensional unit vectors that serve as estimates for the attribute vectors in the model. The second component of the dataset comprises trade shares by app and country in 2021. The third component is a manually compiled list of iOS mobile apps on China's GFW blacklist. These three components are linked via unique app ids. Equipped with this data, I isolate the impact of the GFW from other trade frictions and provide reduced-form evidence in support of the model's predictions regarding the bans. I find that the GFW spurred the entry of additional 10 Chinese apps, with these apps capturing a 15 percentage point higher market share compared to those whose foreign competitors were not banned.

To go further and quantify the welfare implication, I adopt a structural estimation in Section 5. I begin by developing a quantitative model that captures both the utility from real consumption and the leisure from app usage. In the full model, app publishers engage in the two-stage games outlined in the baseline model, monetizing user attention by selling advertisements. I classify the model parameters into two groups. The first group includes parameters that correspond directly to observable data moments and can be separately identified. Notably, a scale parameter of EoS in app sector is recovered from the coefficient in the event study presented in Section 2. The second group comprises parameters that cannot be separately identified. I estimate them using a simulated method of moments (SMM) approach, which fully accounts for the general equilibrium effects in the quantitative model. The estimated model is successful at reproducing the rich patterns of trade, time allocation and income observed in the data.

In Section 6, I use the estimated model to quantify the impact of the GFW. I conduct a

[3] A comparable dataset is NielsenIQ for manufactured goods, as used in Jaccard (2023). However, NielsenIQ offers limited product-level textual data. Instead, Argente, Baslandze, Hanley, and Moreira (2023) address this by clustering barcode-level products into 400 groups and supplementing descriptions from Wikipedia. For this paper, however, I explore variation within a sector, making NielsenIQ data less suitable.

[4] Appendix Table A2 lists the top 10 largest companies measured by market capitalization in 2024.

counterfactual analysis by removing the additional trade barriers imposed by the GFW. Quantitatively, the GFW policy has two distinct effects on welfare for Chinese consumers. First, GFW increased Chinese real incomes by shifting profits away from foreign firms, resulting in a 1.05% welfare gain. However, this benefit was more than offset by the leisure loss from using inferior domestic products, e.g., replacing Chrome with Baidu. Although the overall time spent on apps only slightly decreased, leisure utility for Chinese consumers dropped by 29% . Given the calibrated utility weight (75% on real income), the net welfare loss from the GFW policy is 7.48% for Chinese consumers. This is a large effect. These finding suggests that policymakers have an incentive to impose such bans if their objective is to boost real income. However, the welfare quantification does not account for potential losses arising from privacy and national security concerns, which are important areas for future research.

In addition, I examine the impact on the level of product differentiation among active products. The model exhibits a novel and continuum measure of product innovation by the dispersion of hedonic attributes. I find that the GFW reduced the product innovation among Chinese apps by 4.71%. This finding suggests that the GFW not only replaced foreign products with inferior domestic alternatives but also led to the survival of products with less distinctive hedonic attributes. The policy intervention stifles production innovation as a result of reduced competition. It aligns with the finding in Hsieh, Klenow, and Shimizu (2021), and instead of assuming exogenous innovation efficiency, the results here are an endogenous outcome of the trade policy.

Related Literature This paper connects to several strands of literature. First, it contributes to the role of granular firms in international trade. The existing literature highlights the importance of granular firms in explaining trade flows and shaping comparative advantages (Hanson, Lind, and Muendler, 2015, Bernard, Moxnes, and Ulltveit-Moe, 2018, Bernard, Moxnes, and Saito, 2019). In particular, Bernard *et al.* (2018) emphasize that demand-side heterogeneity is as important as supply-side heterogeneity, yet remains under-explored. Relatedly, this work connects to empirical studies on country-level differences in exported products (Hausmann, Hwang, and Rodrik, 2007, Hummels and Klenow, 2005, Baldwin and Harrigan, 2011, Hallak and Schott, 2011, Sutton and Trefler, 2016). This paper – both methodologically and empirically – provides a practical approach to capturing rich demand-side heterogeneity.

In doing so, this paper adds to the extensive body of work on demand system estimation beyond CES in international trade. This paper is built on the Translog functional form proposed by Christensen, Jorgenson, and Lau (1975) and Diewert (1976). The recent papers, such as Feenstra and Weinstein (2017) and Diewert (2023), employ special cases of the (Generalized) Symmetric Translog for identification. In contrast, this

paper microfounds a demand system through hedonic demand and departures from the symmetry assumption. Relatedly, papers such as Adao, Costinot, and Donaldson (2017), Lashkari and Mestieri (2019), Lind and Ramondo (2023) use a cross-nested CES structure to capture flexible substitution patterns. This paper shares a similar flavour by deriving substitution from underlying product attributes (or nests) but differs in two key ways: the demand system (1) is designed to exploit the LLM representations, and (2) yields finite choke prices, accommodating entry and exit decisions.

This paper connects to the literature on trade shocks, competition and welfare. For example, vast of empirical works document that trade liberalization reallocation resources to high-performing firms and products (Pavcnik, 2002, Trefler, 2004, Mayer, Melitz, and Ottaviano, 2014). Under the CES framework, as discussed in Arkolakis, Costinot, and Rodríguez-Clare (2012), this reallocation is efficient and welfare changes depend on a single parameter. Moving beyond CES, Arkolakis, Costinot, Donaldson, and Rodríguez-Clare (2019) find that the pro-competition effect on welfare is quantitatively small. More recently, Baqaee, Farhi, and Sangani (2023) highlights the importance of demand elasticity heterogeneity in fully capturing the gains from trade. This paper also addresses demand-side heterogeneity and extends the literature by accounting for the granularity of trade shocks – where pairwise elasticities play a critical role.

This paper is related to product differentiation in industrial organization (Lancaster, 1966, Rosen, 1974, Petrin, 2002, Pellegrino, 2023). I borrow from the recent progress in this field, in particular from Pellegrino (2023). Pellegrino (2023) characterizes a product network and explains the source of market power for the public corporations in the United States. This paper differs in several key ways. Theoretically, I adopt a Translog functional form, which yields a closed-form solution for EoS. I also introduce two random processes to characterize the endogenous entry. Empirically, while Pellegrino (2023) uses data from Hoberg and Phillips (2016) with similarity measured by keywords frequency, the model in this paper is calibrated to product-variety level with similarities constructed from prompted LLMs. Lastly, regarding the research question, this paper centers on the interaction between trade policy intervention and competition.

This paper is part of the rapidly growing literature on textual analysis in economics. Embedding models form the backbone of recent advances in Natural Language Processing. The semantic similarity between two pieces of text can be quantified using the cosine similarity between their corresponding embeddings, which offers a straightforward and efficient way for economists to analyze textual data (Kelly, Papanikolaou, Seru, and Taddy, 2021, Sun and Trefler, 2023, Bartik, Gupta, and Milo, 2023, Boeing, Brandt, Dai, Lim, and Peters, 2024).[5] In this paper, rather than simply using similarity as a measure,

[5]Embedding models are integral to advanced language models such as ChatGPT, BERT and Cohere, as well as traditional models like Word2Vec and GloVe.

I develop a model that offers an economic interpretation to LLM embeddings within an estimation context.

Lastly, this paper contributes to the literature on policy interventions in digital sector. The most closely related paper is Bernard, Boler, Chor, Gurund, and Lu (2024), which examines the impact of GFW on knowledge spillovers from foreign to Chinese scholars. In contrast, my focus is on product competition rather than technology spillovers. Sun and Trefler (2023) find the impact of AI deployment on app exports is halved in countries with strong cross-border data flow restriction. This paper differs by studying a different policy and the substitution patterns in trade. In addition, recent papers use similar mobile app data to examine the impacts of privacy regulations (Janssen, Kesler, Kummer, and Waldfogel, 2022, Bian, Ma, and Tang, 2021, Aridor, 2022), but none of them is about international trade and competition.

Outline The remainder of the paper is structured as follows. I start with the event study in Section 2, which motivates the model structure and estimation strategy. Section 3 establishes the baseline model and demonstrate the role of pairwise EoS. In Section 4, Section 5 and Section 6, I apply the model to a dataset of mobile apps. Section 4 describes the dataset, LLM-based estimation and presents reduced-form evidences supporting the main model predictions. Section 5 calibrates the quantitative model to data. Section 6 uses the estimated model to quantify the impact of the GFW policy. Section 7 concludes.

2. Motivating Fact: An Event Study

Before delving into the theoretical model, I begin with an event study that provides empirical evidence of substitution patterns resulting from a granular trade shock. The estimates obtained here are later used in the structural estimation of Section 5 to identify a key demand parameter, as well as in the quantitative analysis of Section 6.

Under a CES framework, consider a granular trade shock on some granular products. The predicted response in the market share of product i that is not targeted by the shock, is given by $\sigma\left(\mathrm{d}\log P\right)$, where $\mathrm{d}\log P$ denotes the change in the aggregate market price induced by the shock. There are two implications: (1) the source of the trade cost is irrelevant, and (2) all competing products respond with the same magnitude. If the CES assumption holds broadly, there is no need for a more complex demand system.

To test this hypothesis, I estimate the market share responses of mobile apps to India's 2020 ban on Chinese apps. On June 29 2020, in response to escalating border tensions between India and China, particularly after a violent clash at the Galwan Valley, the Indian government announced a ban on 59 Chinese apps, including TikTok and WeChat. I manually identified 29 of these apps available on iOS App Store and matched them

Figure 1: Estimated effect of India's ban on market share

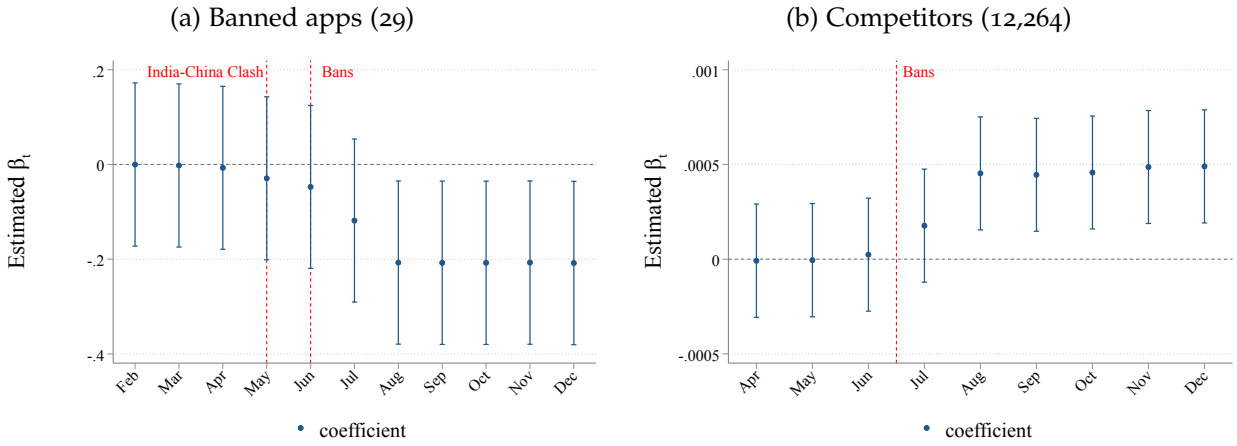

Notes: Estimated event-study coefficient β_τ of equation 1. Panel (a) plots the estimations with the observations on 29 banned apps, and Panel (b) plots the estimations with the observations on 12,264 competing apps not banned. The solid dots are coefficients and the vertical lines reflect the corresponding 95% confidence intervals. The event is defined as the bans implemented on June 2020 and indicated by the red dash line. The plotted coefficients are shown in Table A3 in Appendix B.

with the monthly-active-user data segmented by app, month and market.[6] The event study framework takes the following form:

$$s_{it,IN} = \sum_{\tau} \beta_\tau \, \mathbb{1}_{t=\tau} + \delta_i + \varepsilon_{it} \tag{1}$$

where $s_{it,IN}$ is the market share of app i in month t and in India, given by the fraction of iOS monthly-active-users in percentage points, $\mathbb{1}_{t=\tau}$ are indicators for month, δ_i are fixed effects at app level and ε_{it} is the error term that is clustered at app level. The coefficients measures the difference in market shares between month τ and February 2020 in India. Given that the ban happened in late June 2020, the coefficients on τ before June are expected to be zero and the coefficients on τ after June capture the effects from the ban.

Figure 1 presents the estimation results. First, Figure 1a plots the estimated coefficients on τ for the 29 banned apps. Compared to a typical month before the ban, market shares for these apps slightly decrease on May and June 2020 amid the military conflicts, although this decrease is statistically insignificant. After the ban, the estimated coefficients $\widehat{\beta}_\tau = 0.2$ indicate that the market shares of banned apps fell by 0.2 percentage points – approximately 85% of the market share for a medium banned app. More importantly, Figure 1b plots the estimated coefficients on τ for the remaining 12,264 apps that were not directly affected by the ban, and these coefficients reflect the effects from competition.

[6]The remaining 30 apps were only available on Android devices. All major apps that readers would recognize were successfully matched with the app usage data. Details on the data will be discussed in Section 4.2.1.

Figure 2: The heterogenous estimated effects by deciles in app similarity

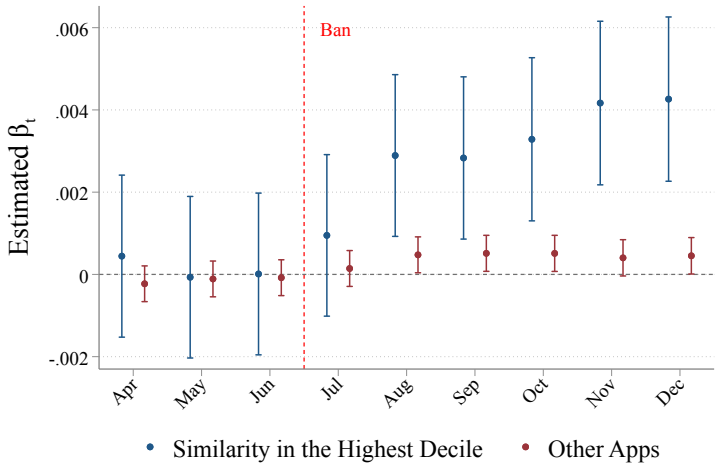

Notes: Estimated event-study coefficient β_τ of equation 1. The solid dots are coefficients and the vertical lines reflect the corresponding 95% confidence intervals. The event is defined as the bans issued on June 2020 and indicated by the red dash line. The plotted coefficients are shown in Table A3 in Appendix B. The robustness check with dependent variable as $s_{it,US}$ and $s_{it,RoW}$ are shown in Table A2 in Appendix B.

Since the ban reduced competition in India, I expect positive coefficients on post June dummies. Indeed, the estimated coefficients $\widehat{\beta}_\tau = 0.0005$ are significantly positive, albeit small in magnitude. On average, competing apps gained an average market share increase of only 0.0005 percentage points.

Next, I investigate the heterogenous responses across different app groups. Leaving measurement details for later discussion, I use detailed app descriptions and advanced large language models to construct a product similarity measure for app pairs: for each app i and a banned app j, I calculate a pairwise product similarity $\omega_{ij} \in [-1, 1]$. I then construct a treatment intensity measure for each of the 12,264 competing apps as $\lambda_i = \sum_j \omega_{ij}$, which reflects the overall substitutability between app i and banned apps. I categorize the apps into groups by deciles in the distribution of λ_i. Figure 2 displays the estimated coefficients for apps in the highest decile and for those in the remaining nine deciles. The estimated coefficients for the two groups are 0.004 and 0.0005, respectively, indicating a statistically significant difference. The results show that apps with higher similarity to the banned apps experience greater gains in market share than those with low similarity. This finding contradicts two key predictions of a CES framework: (1) granularity matters, with product similarity introducing a pairwise effect, and (2) responses to a common trade cost are heterogeneous across apps.

Moreover, this substitution patterns can be identified by LLM-based similarities. I

Table 1: Estimated effect of India's ban on market share

	Banned apps	Competitors: $s_{it,IN}(\times 100)$			
		all	20^{th} percentile	10^{th} percentile	5^{th} percentile
	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}_{\tau \geq Jun\text{-}20}$	-0.70	0.03	0.22	0.31	0.57
	(0.07)	(0.01)	(0.03)	(0.05)	(0.09)
R^2	0.59	0.99	1.00	1.00	1.00
Apps	29	12,264	2,404	1,270	497
Observations	406	163,377	30,241	15,901	6,258
Fixed Effects	i	i	i	i	i

Notes: Estimated coefficient on $\mathbb{1}_{\tau \geq Jun}$. The unit of dependent variable is percentage point for banned apps, and percentage point $\times 100$ for competitors. Standard errors are clustered at app level.

employ a similar specification where the explanatory variable $\mathbb{1}_{\tau \geq Jun}$ indicates month τ is post-June, and I show the estimation results in Table 1. In Table 1, each column presents estimates for a specific group of apps. Column 1 and 2 repeats the estimations shown in Figure 1a and Figure 1b. Columns 3 to 5 focus on competing apps with λ_i in the top 20%, 10% and 5% percentile, respectively. The results indicate that market share gains increase with an app's similarity to the banned apps, with a statistically significant difference. In an extreme case, Moj and Josh – the knock-off apps that replicated TikTok's core functions – rapidly expanded their market shares as replacements after the ban. This empirical finding suggests that the LLM-based similarity measures effectively capture granular substitution patterns without relying on predefined industry or product nests.

Since the event occurred during the COVID-19 pandemic, one identification concern is that the observed pattern might be driven by COVID-related shocks. To address this, I conduct a robustness check by replacing the dependent variable with market shares in the United States and in the rest of the world, excluding India. I find that the response to the shock for competing apps is not present in these other markets. Furthermore, there are no significant differences between apps in different deciles of λ_i. This robustness check helps rule out alternative channels that could contribute to the observed pattern and supports the results shown in Table A2 in Appendix B.

3. Theoretical Model

In this section, I set up a granular trade model and define the equilibrium. The framework relates closely to Eaton, Kortum, and Sotelo (2012) and Gaubert and Itskhoki (2021). The key difference is that I employ a micro-founded Translog utility function, featuring the pairwise EoS shown in Section 2. In the baseline model, there are N countries and one sector. Each firm produces a single final good and can sell to many markets. Let i represent a firm/product, $\mathcal{M}$ represent the set of firms/products, o represent the exporting country, and d the importing country. The baseline model is generally applicable to any sector, and I will present a quantitative model tailored to digital service trade embedded in mobile apps in Section 6.

The competition outcome is achieved through a two-stage game. In the first stage, each exporting country o is born with a discrete number of potential entrants. Each potential entrant perfectly foresees the outcome of the second stage and decides whether to enter each market d. In the second stage, products that chose to enter compete in a Bertrand oligopoly, while products that did not enter remain inactive, charging a reservation price. In equilibrium, goods markets and the labour markets clear in all countries.

Products Let $\mathcal{M}^o$ represent the set of potential entrants originating from country o. The number of products from country o, $M^o = \|\mathcal{M}^o\|$, is a realization of a Poisson random process with the scale parameter $\overline{M}^o$ so that $\mathbb{E}(M^o) = \overline{M}^o$. Each of the potential entrants is endowed with (1) a product blueprint, and (2) a product technology. Both are drawn simultaneously from random processes. First, following the tradition of modelling differentiated products in industrial organization, a product blueprint describes a product variety as a mix of attributes. The space of product attributes is fixed and indexed by $k \in \{1, 2 \ \ldots, K\}$, and a blueprint is a unit vector $\mathbf{a}_i$ defined in the K-dimensional attribute space:

$$\mathbf{a}_i = [a_{i1}, \ a_{i2}, \ \ldots, \ a_{iK}]^T \text{ and } \sum_s a_{is}^2 = 1 \tag{2}$$

where a_{is} is the contribution of attribute s in product i. Each potential entrant takes an *i.i.d.* draw from a von Mises-Fisher (hereforth vMF) distribution, which is a multivariate normal distribution whose support is the surface of a unit sphere.[7] Formally, the probability density function of $\mathbf{a}$ is given by:

$$f_K(\mathbf{a}, \ \boldsymbol{\mu}_o, \ \kappa) = C_K(\kappa) \exp(\kappa \boldsymbol{\mu}_o^T \mathbf{a}) \tag{3}$$

[7] vMF may not be familiar to economists but is used heavily in text mining literature. Its ability to model directional data makes it highly suitable for various NLP tasks, providing a robust framework for handling the unique characteristics of text data.

where $\boldsymbol{\mu}_o$ is a K-dimensional unit vector, $\kappa > 0$ is a scalar and $C_K(\kappa)$ is a normalization constant.[8] The distributions are independent across countries. The country-specific parameter $\boldsymbol{\mu}_o$ denotes the central direction around which product attributes $\mathbf{a}$ are concentrated. The parameter κ common to all countries, controls the dispersion of the distribution: the greater the value of κ, the more concentrated are product attributes $\mathbf{a}$ around $\boldsymbol{\mu}_o$. In the context of differentiated products, κ reflects the costliness of innovation, as unique products located further from the distribution center are rarer. By assuming the distribution, I am able to track high-dimensional product differentiation by low-dimensional directional statistics [9].

Second, each product variety is supplied by a unique blueprint owner with a productivity φ_i. The owner uses local labour to produce with a linear production technology, formally represented by the production function $q_i = \varphi_i L_i$. Consistent with trade literature (e.g., Melitz (2003)), I assume the productivity draws follow a Pareto distribution with shape parameter ν and scale parameter $\overline{\varphi}_o$. The distribution is defined as

$$h(\varphi_i,\ \overline{\varphi}_o,\ \nu) = 1 - (\varphi_i/\bar{\varphi}_o)^{-\nu},\ \text{where } \varphi \geq \overline{\varphi}_o \tag{4}$$

To summarize, products are heterogenous over states $\chi = (\mathbf{a}, \varphi)$, where $\mathbf{a}$ and φ are demand-side characteristic and supply-side characteristic, respectively. Figure 3 provides an illustrative example with three products defined in two-dimensional attribute space. Each product is depicted as a vector, where $\mathbf{a}$ determines the angle and φ determines the length. In this example, product 1 is biased towards attribute 1 and highly productive, while product 2 is biased towards attribute 2 and is less productive than product 1. Now consider a third product, represented by the blue vector, entering the market. Product 3 is highly productive, posing a competitive threat to both products 1 and 2. However, due to its similarity in attribute mix with product 2, it poses a greater competitive threat to product 2. These dynamics become even more complex when considering thousands of products and millions of pairwise substitution patterns.

Preferences Let us consider the demand in a given market and I drop the the market index d for simplicity. A representative consumer can observe the universal product varieties from all countries, and thus her choice set is $\mathcal{M} = \mathcal{M}^1 \cup \mathcal{M}^2 \cup \cdots \cup \mathcal{M}^N$. To capture the distribution of products in the product space, I can define a product network and map it to demand parameters following the concept from Pellegrino (2023). The

[8] $C_K(\kappa) = \frac{\kappa^{K/2-1}}{(2\pi) I_{K/2-1}(\kappa)}$ where $I_s(x)$ denotes the modified Bessel function.
$I_s(x) = \sum_m^{+\infty} \frac{1}{m!\Gamma(m+s+1)} \left(\frac{x}{2}\right)^{2m+s}$.

[9] More details are provided in Appendix A.1.

Figure 3: Illustrative example of three products in 2-dimensional space

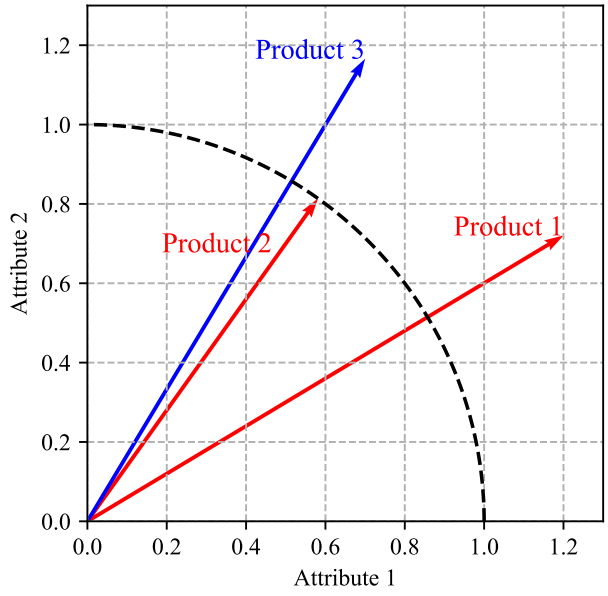

Notes: A demonstration of products when $K = 2$. There are three products depicted by three vectors. For each vector, $\mathbf{a}_i$ determines its angle and φ_i determines its length. The black dash line denotes a unit circle.

product network is summarized by a matrix $\mathbf{A}$:

$$\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_M] = \begin{pmatrix} a_{11} & a_{12} & \ldots & a_{1K} \\ a_{21} & a_{22} & \ldots & a_{2K} \\ \vdots & \ldots & \ldots & \vdots \\ a_{M1} & a_{M2} & \ldots & a_{MK} \end{pmatrix}^T \tag{5}$$

where $M = \|\mathcal{M}\|$ is the number of products. I introduce three concepts from directional statistics (Mardia and Jupp, 2009), which is useful for describing a product network and the properties of demand system.

Definition 1 *The resultant vector* $\overline{\mathbf{a}}$ *for product network* $\mathbf{A}$ *is the sum of individual attribute vectors and given by*

$$\overline{\mathbf{a}} = \left[\sum_i^M a_{i1}, \sum_i^M a_{i2}, \ldots, \sum_i^M a_{iK} \right]. \tag{6}$$

The length of the resultant vector, denoted by $\|\overline{\mathbf{a}}\|$*, captures the intensity of* $\mathbf{A}$*; the normalized resultant vector, denoted by* $\overline{\mathbf{a}}/\|\overline{\mathbf{a}}\|$*, captures the average orientation of* $\mathbf{A}$ *as a unit vector.*

The resultant vector $\overline{\mathbf{a}}$ is a single vector that provides a useful summary of the overall positioning within a network. First, $\|\overline{\mathbf{a}}\|/M \propto \left(\sqrt{var(\mathbf{A})}\right)^{-1}$, ranging from 0 to 1, captures the dispersion of $\mathbf{a}$ in attribute space. A greater value indicates that the individual

vectors are more aligned, suggesting that the products are more homogenous in terms of the mix of attributes. Conversely, a smaller value suggests greater dispersion, implying that the products are more differentiated. When $\|\overline{\mathbf{a}}\|/M$ equals to 1, all products locate in the same place and the underlying $\kappa \to \infty$. In addition, resultant vector $\overline{\mathbf{a}}$ locates at the central direction of $\mathbf{A}$. I further define the centrality of each product accordingly.

Definition 2 *Ω_i is the centrality of product i and given by the inner product of vector $\mathbf{a}_i$ and the resultant vector:*

$$\Omega_i = \mathbf{a}_i^T \cdot \overline{\mathbf{a}} = \sum_{j=1}^{M} \omega_{ij} \tag{7}$$

Ω_i measures the angular distance between product i and resultant vector $\overline{\mathbf{a}}$ and indicates how closely the product aligns with the average orientation of the network. When Ω_i is larger, product i is positioned closer to the center, indicating it shares a similar mix of attributes with the majority of products. In this case, product i is perceived as a homogeneous good by consumers. When Ω_i is small, product i is located farther from the center, indicating that product i offers a distinctive combination of attributes. Therefore, product i is perceived as a differentiated good by consumers.

Definition 3 *The cosine similarity ω_{ij} is the inner product of vector $\mathbf{a}_i$ and vector $\mathbf{a}_j$:*

$$\omega_{ij} = \mathbf{a}_i^T \cdot \mathbf{a}_j = \sum_{s=1}^{K} a_{is} a_{js} \tag{8}$$

which indicates the product similarity between product i and j.

The cosine similarity is a statistic to describe the relative position of two vectors and its value ranges from -1 to 1. In the context of product attribute space, the similarity feeds directly into the patterns of product substitution. A cosine similarity closer to 1 suggests that two products share a highly overlapping mix of attributes, leading to greater substitutability between them. Conversely, lower values (closer to 0) would indicate that products are dissimilar, potentially even complementary in nature.[10]

A representative consumer faces an expenditure cost, which is a quadratic function of attribute and product prices. The functional form is a case of the generalized hedonic demand as defined in equation B.1 of Pellegrino (2023), under certain homotheticity assumptions. The economic intuition for the functional form is as follows: consumers value products as bundles of attributes, fundamentally caring about the price of each attribute. However, the composition of these bundles – how attributes are "packed" together – also

[10]While I allow for the possibility that products are complementary , I restrict the analysis to cases where κ is sufficiently large such that $\omega_{ij} \geq 0$ for most j and all i.

influences consumer preferences. This demand system yields a Translog expenditure function defined on products (Christensen *et al.*, 1975, Feenstra and Weinstein, 2017, Diewert, 2023), with parameters corresponding to the network statistics described above. Formally, the expenditure costs required to obtain a fixed level of utility is given by:

$$\ln(e) = \alpha_0 + \sum_{i}^{M} \alpha_i \ln p_i + \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} \theta_{ij} \ln p_i \ln p_j \tag{9}$$

where p_i is the price of product i and the θ_{ij} are

$$\theta_{ij} = \begin{cases} \frac{\theta}{\|\bar{\mathbf{a}}\|} \omega_{ij} & j \neq i \\ -\frac{\theta}{\|\bar{\mathbf{a}}\|} (\Omega_i - 1) & j = i \end{cases} \tag{10}$$

See Appendix A.2 for detailed derivations.[11] There are three components in θ_{ij}. The parameter θ governs the overall second order effects. The length of resultant vector $\|\bar{\mathbf{a}}\|$ captures a weighted love-of-variety effect – incorporating both the number of products and their dispersion in attribute space. The competition effect is amplified when dispersion increases. For the case $i \neq j$, θ_{ij} is proportional to their product similarity ω_{ij}, indicating that the demand for product i is more sensitive to price changes in product j if they locate closer to each other. For the case $i = j$, θ_{ii} is a function of $\Omega_i - 1 = \sum_{j \neq i} \omega_{ij} > 0$, indicating that the demand for product i is more responsive to its own price if it occupies a central position in $\mathbf{A}$. These three components together determine the substitution patterns. The expenditure share of product i can be solved from the *F.O.C.* with a budget constraint:

$$s_i = \alpha_i + \sum_{j}^{M} \theta_{ij} \ln p_j = \alpha_i - \widetilde{\theta} \Omega_i \left[\ln p_i - \sum_{j} \left(\frac{\omega_{ij}}{\sum_{j'} \omega_{ij'}} \right) \ln p_j \right] \tag{11}$$

where $\widetilde{\theta} = \theta / \|\bar{\mathbf{a}}\|$ for the simplicity of notations.

Discussion: It is worth pointing out how the demand equation given in 11 differs from other recent studies. First, the Translog utility function has been used in other papers, which are the special cases of this paper. Feenstra and Weinstein (2017) examine a Symmetric Translog function where $\theta_{ij} = \theta / M$, which is a special case of my model where all products are identical and $\omega_{ij} = 1$. Diewert (2023) considers a case where

[11] This function satisfies two restrictions on θ_{ij}: (1) *symmetric restrictions* where $\theta_{ij} = \theta_{ji}$; and (2) *homotheticity restrictions* where $\sum_j \theta_{ij} = 0$ for any i and j. Additionally, by imposing a homotheticity restriction on α such that $\sum_i \alpha_i = 1$ and $\alpha_i \in (0, 1)$, the function becomes homogeneous of degree one, as discussed in Diewert (1976) and Christensen *et al.* (1975). This paper focuses on the impacts of competition, and the main conclusions do not rely on the specific parameterization of α_i. I will discuss α_i further in the quantitative exercise of Section 5.

$\omega_{ij} = \omega_i\omega_j$, which assumes firms are differentiated along one dimension such that $\frac{\theta}{\sum_{i'j'} a_{i'}a_{j'}} a_i a_j$. These two cases fall within the class discussed in Matsuyama and Ushchev (2017), resulting in a single and separable price index. Therefore, they cannot generate the pairwise substitution patterns which is the core of this paper.

Second, equation 11 has a similar expression to CES-like demand functions, such as in Baqaee *et al.* (2023), but with two key differences. On one hand, s_i on the left-hand side is expressed in levels instead of in logarithms. This difference means that while the choke price in the CES case is infinite, a sufficiently high price can lead to zero sales for a product. On the other hand, the price aggregator faced by product i is given by $\sum_j \left(\frac{\omega_{ij}}{\sum_{j'}\omega_{ij'}}\right) \ln p_j$, a weighted average of competitor prices where weights correspond to product similarity with product i. This competition effect depends on the joint distribution of $\mathbf{A}$ and $\{p_j\}_{j=1}^M$, a feature absent in other models.

As a consequence, none of the papers mentioned above can capture the pairwise competition that is central to my paper. I formalize the concept in proposition 1 by the expressions of price elasticities and elasticities of substitution:

Proposition 1 *The price elasticities between the product i and j is*

$$\eta_{ij} \equiv \left.\frac{\partial \ln q_i}{\partial \ln p_j}\right|_{U \text{ const.}} = \begin{cases} s_j + \frac{\theta_{ij}}{s_i} & j \neq i \\ -1 + s_i + \frac{\theta_{ii}}{s_i} & j = i \end{cases} \tag{12}$$

The Allen-Uzawa partial elasticities of substitution:

$$\sigma_{ij} \equiv \frac{\eta_{ij}}{s_j} = \begin{cases} 1 + \frac{\theta_{ij}}{s_i s_j} & j \neq i \\ 1 + \frac{\theta_{ii}}{s_i^2} - \frac{1}{s_i} & j = i \end{cases} \tag{13}$$

See Appendix A.5 for derivations. When $\widetilde{\theta} = 0$, the demand function collapses down to Cobb-Douglas functional form and all EoS are equal to unity. When $\widetilde{\theta} > 0$, the pairwise $\theta_{ij} \propto \omega_{ij}$ for product i and product j up to the scale $\widetilde{\theta}$. The expression indicates that EoS between pairs with higher product similarity are greater. Conditional on the product similarity, large products are less substitutable. There are $M(M+1)/2$ pairwise EoS. My model structure reduce the estimating parameters to M number of K-dimensional vectors and one scale parameter $\widetilde{\theta}$.

Similarly, the pairwise competition also affects the extensive margin. Due to the Translog functional form, a product at $\mathbf{a}_i$ can capture a zero sale when its price is sufficiently higher or its productivity is sufficiently low. I define a concept of a choke price in the context of my model in Proposition 2:

Proposition 2 *Given a product network* $\mathbf{A}$*, the choke price at* $\mathbf{a}_i$ *position is the solution for* $s_i = 0$ *in equation 11 and given by:*

$$\ln \bar{p}(\mathbf{a}_i) = \frac{\alpha_i}{\widetilde{\theta}\Omega_i} + \sum_j \left(\frac{\omega_{ij}}{\sum_{j'} \omega_{ij'}} \right) \ln p_j \tag{14}$$

Therefore, the entry probability at $\mathbf{a}_i$ *is given by* $Pr[\varphi \geq \overline{\varphi}(\bar{p}|\mathbf{a}_i)]$*:*

1. *If* $\mathbf{a}_i$ *is closer to a product with high productivity* φ*, the entry probability is lower.*
2. *If* $\mathbf{a}_i$ *is more central in the network, the entry probability is lower.*

See Appendix A.6 for derivations. Figure 4 illustrates proposition 2. Suppose there are two products in the market: product 1 and product 2. Now consider an additional potential entry, represented by a vector pointing to a specific position in the figure and determined probabilistically by the random processes. The blue-shaded region represents areas where the entrant cannot survive if it enters. The edge of this exit region indicates the "choke productivity" corresponding to the choke price at each angle. Notably, the edge is not a circular arc, indicating that choke productivity varies by angles ($\mathbf{a}$). Choke productivity is higher in more central positions and lower in areas closer to the axes. Thus, entry probability increases the farther the entrant is positioned from existing products. Furthermore, choke productivity is even higher in areas closer to product 1 compared to product 2, as product 1 is more productive and thus harder to compete with.

Market Structure Let us consider the trade and competition. The equilibrium is solved backward, so I start with the second stage. I introduce impediments to trade in a standard way: selling one unit of product i to market d from source o requires exporting $\tau_{i,od} \geq 1$, where domestic costs are set to $\tau_{i,oo} = 1$ for all i. The effective marginal cost of product i from country o in market d is

$$c_{i,od} = \tau_{i,od} c_i = \tau_{i,od} \left(\frac{w_o}{T_i} \right) \tag{15}$$

Due to the linearity of the production function, each firm's profit maximization problem is separable across markets and hence can be considered one market at a time. Given firm i enters market d in the first stage, it choose its price $p_{i,od}$, taking as given the wage rate and the prices of its competitors, to maximize its profit from serving the market:

$$\Pi_{i,od} = s_{i,od} E_d - c_{i,od} q_{i,od} \tag{16}$$

where the expression for market share $s_{i,od}$ is given by equation 11. The solution for this Bertrand-Nash competition game is a markup price setting rule:

$$p_{i,od} = \frac{\theta_{ii} - s_{i,od}}{\theta_{ii}} \, c_{i,od} = \left(1 + \frac{s_{i,od}}{\widetilde{\theta}(\Omega_i - 1)} \right) c_{i,od} \tag{17}$$

Figure 4: Illustrative example of choke prices

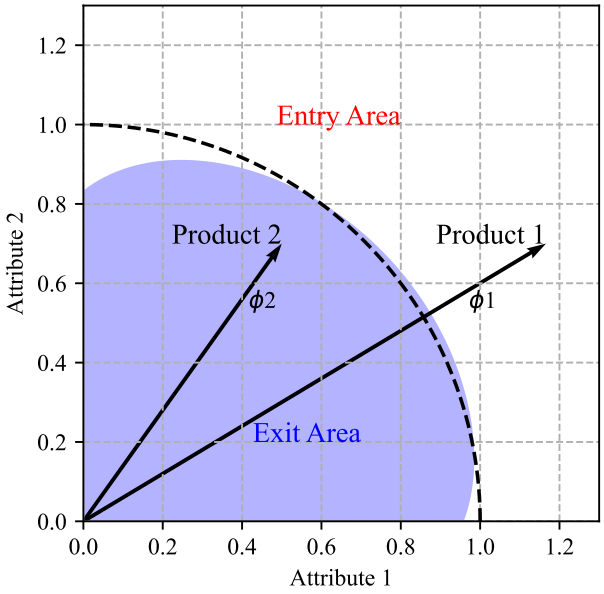

Notes: A demonstration of choke prices when $K = 2$. There are two existing products depicted by product 1 and product 2. I let the price be $1/\varphi$. The edge of blue-shaded areas represent the productivity that capture zero market share at angle $\mathbf{a}$ following equation 14. The black dash line denote a unit circle.

This expression is analogous to a nested logit structure where the nests are infinitesimally distributed over the surface of a sphere:

Proposition 3 *Markup in equation 17 is the generalization of nested logit (e.g., Fajgelbaum, Grossman, and Helpman (2011)) such that the nests are continuum.*

See Appendix A.7 for derivations.

Entry In the first stage, an equilibrium of the entry game is achieved when the equilibrium profits of a subset of firms, as given by equation 16, are non-negative, while for any additional entrant profits upon entry would be negative. With a discrete number of potential entrants, there may exist multiple equilibria in the entry game. We therefore consider a sequential entry game in each market separately. Specifically, consumers in market d can observe the universal set of product varieties from all countries $o \in \{1, 2, \ldots, N\}$, so the initial choice set is $\mathcal{M}$. In the first hypothetical round, all firms are assumed to enter and engage in the Bertrand price-setting game as described above. The profits in this hypothetical equilibrium are derived from the *F.O.C.*

In the the first round, I reassign indexes i such that $\Pi_{1^{(1)}d} \geq \Pi_{2^{(1)}d} \geq \cdots \geq \Pi_{M^{(1)}d}$ where $^{(1)}$ indicate the round number. If the least profitable firm earns a non-negative

profit as $\Pi_{M^{(1)d}} \geq 0$, then this constitutes an equilibrium where all potential entrants remain active in market d. If not, it indicates that at least one firm should exit market d. I let the least profitable firm, indexed by $M^{(1)d}$, exit, and the remaining $M-1$ firms, $i \in \mathcal{M} \setminus \{M^{(1)d}\}$, repeat the Bertrand price-setting game. I consider the equilibrium is defined as the sequential equilibrium following Proposition 4:

Proposition 4 *A sequential equilibrium in market d is reached when:*

1. $\Pi_{M_d+1^{(M-M_d-1)d}} < 0$ *in the* $(M-M_d-1)^{th}$ *round*
2. $\Pi_{M_d{}^{(M-M_d)d}} \geq 0$ *the* $(M-M_d)^{th}$ *round*

$\mathcal{M}_d = \mathcal{M} \setminus \{M^{(1)d}, M-1^{(2)d}, \ldots, M_d+1^{(M-M_d-1)d}\}$ *is the set of active firms. Firms choose their prices as follows:*

$$p_{i,od} = \begin{cases} \left(1 + \frac{Ms_{i,od}}{\theta(\Omega_i - 1)}\right) c_{i,od} & i \in \mathcal{M}_d \\ \\ \bar{p}_{i,d} & i \in \mathcal{M} \setminus \mathcal{M}_d \end{cases} \tag{18}$$

where $\bar{p}_{i,d}$ is a reservation price from equation 14, which captures zero profits.

The sequential entry equilibrium is defined in a way reminiscent of Gaubert and Itskhoki (2021). In their model, firms are heterogeneous in productivity only, allowing for an ex-ante ranking by productivity. Starting from the most to least productive firms ensures that the resulting equilibrium is the most efficient among all possible equilibria. However, in this model, such a ranking is not straightforward due to the complex interaction between product attributes and productivity, which prevents a single ex-ante indicator. Instead, I address this through the dual problem on the exit side. In the appendix, I experiment with alternative sequences and test the robustness of welfare outcomes. Moreover, consistent with Feenstra and Weinstein (2017), the equilibrium expenditure costs fully account for the presence of new and disappearing goods, with finite reservation prices set by inactive firms

General Equilibrium is defined by vectors of wages and prices such that goods and labor markets clear in all countries. Specifically, the goods markets clear at

$$\sum_{i \in \mathcal{M}} p_{i,od} q_{i,od} = E_d = w_d L_d \tag{19}$$

where $q_{i,od} = 0$ if firm i is inactive in market d. Additionally, labour market clearing requires that the aggregate labor supply equals the total labor demand from all domestic firms in each country:

$$\sum_{i \in \mathcal{M}^o} \sum_{d}^{N} \frac{q_{i,od}}{T_i} = L_o \tag{20}$$

For any realizations of the random processes, there exists an unique equilibrium following the sequence described above.

Proposition 5 *Given parameter* $\{L_n, \overline{M}_n, \overline{\varphi}_n, \nu, \boldsymbol{\mu}_n, \kappa, \theta\}_{n=1}^N$ *and trade barriers* $\{\tau_{od}\}_{o,d=1}^N$, *there exists a unique equilibrium defined by* $\{\mathcal{M}_n, w_n, \{p_{i,od}, s_{i,od}\}_{i\in\mathcal{M}_n}\}_{n=1}^N$ *for any realization of* $\{\varphi_i, \mathbf{a}_i\}_{i\in\mathcal{M}^o}$ *following the entry sequence described above.*

See Appendix A.8 for a proof.

4. Model Implications and Empirical Evidence

4.1. Granular bans in the model

Bans targeted at specific products or bans have become increasingly common. For example, TikTok faces restrictions in many countries and is banned in India; X/Twitter has had legal issues with Brazil's Supreme Court; and Huawei and ZTE face restrictions in the United States and Canada. Understanding how these bans shape competition is essential. In this model, I represent bans as non-tariff trade barriers. Specifically, the iceberg costs defined in equation 15 have two components:

$$\ln \tau_{i,od} = \ln \tau_{od} + \ln \overline{\tau}\, \mathbb{1}_{i,\text{policy}_d} \tag{21}$$

where $\ln \tau_{od}$ is the standard dyadic gravity costs and $\ln \overline{\tau}$ is the costs from policy intervention. Specifically, a ban on product i by country d increase i's marginal costs by $\overline{\tau}$ log points in market d. Modelling a ban as an iceberg cost reflects the fact that the policy interventions often result in slowdown in response times and limited local services. Even when an app is fully banned, a non-negligible number of local users may still access it via virtual private networks (VPNs).[12]

I evaluate the impact of bans by a comparative statics between the initial equilibrium without the bans to the new equilibrium with the bans. Following the convention of trade literature, I use hat algebra notation to denote the changes from the initial equilibrium to the new equilibrium – formally, $\hat{x} = x_1 - x_0$.[13] The policy intervention is a ban on a set of products $j \in \{j \mid \mathbb{1}_{j,\text{policy}} = 1\}$. First, the ban induces the entry of new products as the changes in choke prices are given by:

$$\widehat{\ln \bar{p}}_{i,d} = \left(\sum_{j|\mathbb{1}_{j,\text{policy}}=1} \frac{\omega_{ij}}{\Omega_i} \right) \ln \overline{\tau} \tag{22}$$

[12]Prior to 2022, Apple almost never removed banned products from the market. However, amid escalating political tensions, Apple began removing apps like WhatsApp and Threads from the Chinese App Store in response to government orders, starting in April 2024.

[13]Compared with Dekle, Eaton, and Kortum (2007), the changes as defined in levels not in ratios for simplicity purpose. This difference is due to the Translog utility function. I will also use ratios as the robustness check for empirical evidences.

Figure 5: Illustrative example of policy intervention

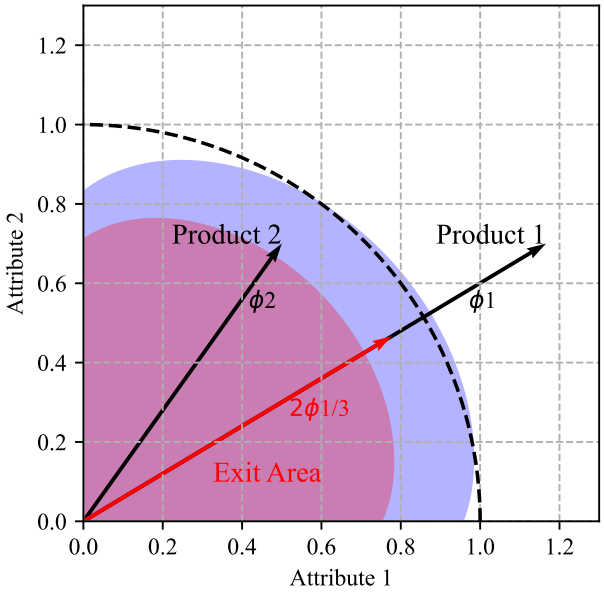

Notes: A demonstration of the impact of a granular ban on product 1. The black vectors denote the products and the red vector denote the effective productivity after product 1 is banned. The blue-shaded region denote original exit area and the red-shaded region denote new exit area after the ban. The black dash line denote a unit circle.

Figure 5 illustrates the changes in choke prices and the impact on firm entry. Consider a ban on product i, which reduce the effective productivity of product i from φ_1 to $\frac{2}{3}\varphi_1$. The policy intervention reduces the competition and shifts the exit area from the blue-shaded region to the red-shaded region. The changes $\widehat{\ln \bar{p}_{i,d}}$ are shown as the difference between two region edges. More additional entry appeared in the places relatively closer to the banned product.

Second, the ban change the market shares of the continuing products. The changes of the banned products are given by:

$$\widehat{s}_{i,od} = -\widetilde{\theta}\left(1 - \sum_{j|\mathbb{1}_{j,\text{policy}}=1} \omega_{ij}\right) \ln \overline{\tau} \tag{23}$$

This expression includes the direct effect on the product as well as the competition from other banned products. The changes of products that are not targeted by the ban are given by:

$$\widehat{s}_{i,od} = \widetilde{\theta}\left(\sum_{j|\mathbb{1}_{j,\text{policy}}=1} \omega_{ij}\right) \ln \overline{\tau} \tag{24}$$

The changes merely come from the reduced competition and the magnitude depends on the similarity with all banned products. Notably, Equation 23 and 24 correspond to the empirical specifications presented in Section 2. To fully account for the long-term responses, I turn to a different policy implemented decade ago and treat the current data as the equilibrium with policy intervention.

4.2. Data and institutional background

4.2.1. Mobile app data

To validate the mechanisms of the model and quantify the welfare implications, I apply the model on a dataset for mobile apps. The primary database is on app usage data purchased from SensorTower. Sensor Tower is the largest and most reliable company providing app-level metadata. [14] For each app, I have detailed metadata scraped from their App Store pages, including app descriptions, release dates, publishers, publisher countries, and app categories. I use the publisher's country as the exporting country of the app and determine the production costs and trade costs faced by each app. More importantly, I leverage rich textual descriptions at the product variety level to estimate the product attributes and pairwise similarities.

The data also track app-level usage by destination country the apps available on the iOS App Store. I have monthly data on the number of monthly active users by destination country for a substantial period, from 2015 to 2021. Using this data, I construct a high-frequency market share measure, calculated as the fraction of users for an app relative to the total users across all iOS apps. This variable is used as the dependent variable in Section 2. In addition, I have a snapshot of the time spent on each app by destination country in 2021, representing the total number of hours from all users on an app in that year. This variable captures equilibrium time allocation and is used as a dependent variable in Section 4.3 as well as the trade flow in Section 5 and Section 6.

4.2.2. Measuring app attributes

One key reason for utilizing this dataset is its rich textual data at the product variety level. App descriptions, provided by publishers, outline the main functionalities, benefits, and use cases, helping users understand what the app does and how it can enhance their experience. A well-crafted description not only explains features but also addresses app regulations, encouraging downloads and long-term engagement. Developers can use descriptions to clarify permissions, subscriptions, or in-app purchases, fostering

[14] A potential concern may be the absence of usage data from Android systems. First, as Google Play is unavailable in China, comprehensive Android market information for the Chinese market is limited. Second, within each platform, an app's market share is highly correlated between iOS and Android.

Figure 6: Example of app description for Netflix

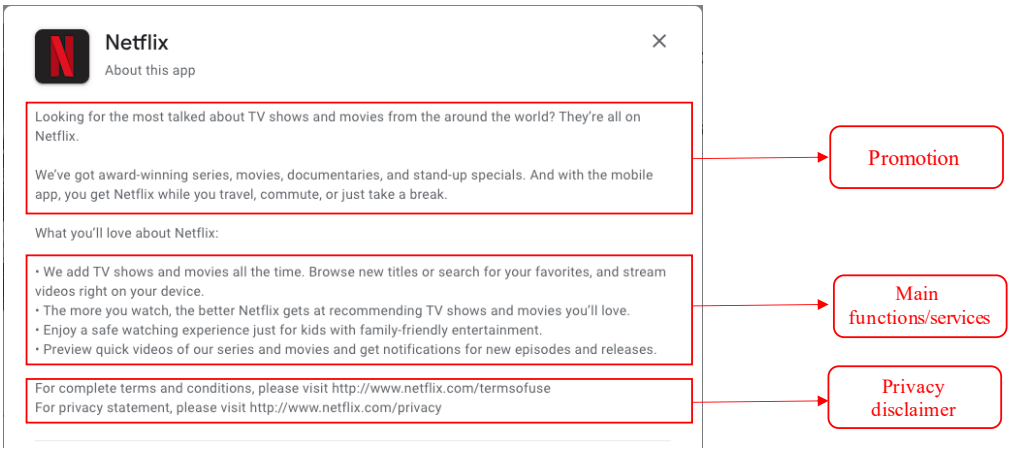

Notes: A screenshot of Netflix's iOS app store page: `https://apps.apple.com/us/app/netflix/id363590051`

transparency with potential users. Thus, the descriptions contain valuable information about product attributes, though they also include unrelated noises. Figure 6 displays an example for Netflix. In this example, there are rich information about the main function and service for Netflix along with the promotion and privacy disclaimer. The goal of the estimation is to get a numerical measure with a focus on product attributes, which is achieved from a two-step procedure with LLMs.

To handle complex tasks and improve accuracy, high-quality labeled data is often required, though it is scarce and costly. As a result, recent advances in natural language processing increasingly rely on synthetic data, particularly data generated by large language models (LLMs) (Bubeck, Chandrasekaran, Eldan, Gehrke, Horvitz, Kamar, Lee, Lee, Li, Lundberg *et al.*, 2023, Liu, Wei, Liu, Si, Zhang, Rao, Zheng, Peng, Yang, Zhou *et al.*, 2024). Josifoski, Sakota, Peyrard, and West (2023) show that when certain tasks cannot be directly solved by LLMs, a more efficient approach is to reverse the task by prompting the LLM to generate input text (which LLMs excel at) from predefined output structures. In this paper, I adopt a similar methodology to that proposed by Josifoski *et al.* (2023) to estimate the product network for universal mobile apps.

First, I utilize the "`GPT-4o-mini`" model, one of the most advanced and efficient LLMs available, to construct a zero-shot prompt for generating synthetic data. The knowledge base for this task is derived from the descriptions of all iOS apps. Specifically, I prompt the LLM to act as an iOS app analyst, identifying key features, target audiences, and unique selling points to evaluate how closely apps compete within the same market segment. The output is structured as a concise paragraph, followed by a bulleted list of core functionalities. The language and tone remain consistent across outputs. To minimize

noise from subjective or inconsistent information, elements such as subscription options, privacy policies, and terms of service are excluded. Potential biases in the data and further details on prompt construction are discussed in Appendix C.1.

Next, I use the synthetic data generated from this process to feed an embedding model, "`text-embedding-3-large`", which estimates the corresponding vector representations. A transformer-based model extracts features from user-submitted text and encodes these features as numeric vectors, known as embeddings. Each embedding is a unit vector in a fixed-dimensional space, conceptually similar to the attribute vectors, as demonstrated in Figure 3, where each vector represents a distinct subject area within the LLM's corpus. The synthetic data generated in the first step effectively bridges the gap between product space and corpus space, ensuring that the vectors are meaningful in both contexts. The outputs are unit embedding vectors with 3072 dimensions.

In summary, the two-step process outlined above allows me to generate unit embedding vectors from synthetic data, based on the app description knowledge base. In the subsequent quantitative analysis, these embedding vectors serve as the measurements of $\mathbf{a}$ in the model. Figure 7 displays a visualization of estimated embeddings on the surface of a three-dimensional ball. Each colour block represents model predicted app categories within which apps are very similar, such as search engines or photo sharing platform.

4.2.3. The Great Firewall policy

This paper uses the data to examine ten impact of the Great Firewall (GFW) policy in China, a set of bans implemented during the early internet era between the two largest economies. The GFW is a combination of legislative actions and technologies enforced by the Chinese government to regulate the domestic internet market. Concerned about the spread of illegal and harmful content – such as pornography, violence, fake news, and politically sensitive information that could threaten social stability and national security – the Chinese government initiated a project to monitor sensitive words and filter content, which became known as the "Great Firewall" project. This initiative began in 1998 and was completed in 2008.

A key function of the GFW is to selectively block access to content using methods like DNS spoofing, URL filtering, and IP range bans. In 2009, China's Ministry of Industry and Information Technology (MIIT) issued the "Notice on Establishing a Blacklist Management System for Illegal Domestic Websites."[15] In the same year, YouTube and Facebook were added to the blacklist. In 2010, Google exited the Chinese market, and all its domains and products were subsequently blacklisted. Today, many mobile apps and

[15] *Notice on Establishing a Blacklist Management System for Illegal Domestic Websites [2009] No.371*

Figure 7: Estimation of **A** for mobile apps and app categories

Notes: This is a visualization of **A** from textual app descriptions. I use UMAP algorithm to reduce 3072-dimensional vectors to 3-dimensional vectors. Each node represents an app in my dataset and all apps are distributed on a unit ball. Each colour represents an app category predicted from aggolomerative clustering algorithm based on the 3-dimensional vectors. Notice that relative locations of the nodes cannot perfect reflect the ω_{ij} in quantitative exercise due to the dimensionality reduction.

websites are blocked, particularly those of companies unwilling to comply with China's strict regulations on data collection and privacy, including Google and Meta.

Although the existence of the blacklist is well known, the standards and timeline for banning sites have been opaque, leaving users uncertain about when a foreign app or website might be banned. Chinese internet users could continue using foreign services while being aware of the risk of sudden, unannounced bans. Even after official bans, some Chinese users bypass the firewall using VPNs. Certain universities and foreign-owned companies provide VPN access for their students and employees under special permissions. However, this method carries increased costs and potential political risks. For instance, in October 2023, a software engineer had 1 million RMB of illegal income confiscated. According to the disclosed verdict, he was penalized for illegally "using a VPN to access Google Chrome, log into GitHub, and set up meetings on Zoom."[16]

[16]See news sources from https://www.sohu.com/a/732098677_121124373

I manually collect the information of blacklist apps from public information and crosscheck with the results from China Firewall Test (chinafirewalltest.co), which is a real-time check on banning status on domains.[17] I found 268 banned apps that can be matched with my dataset. 131 are main banned apps (e.g., Instagram), and the rest are supplementary apps (e.g., Layout from Instagram). The majority of bans happened around 2010. Therefore, I treat the snapshot in 2021 as the equilibrium with the bans. Although I do not observe the pre-treated equilibrium, I will construct a valid control group for each of the banned apps and address the identification with the cross-sectional data variation.

4.2.4. Final dataset

The final dataset includes app attribute vectors, app usage by country, and a list of apps banned in China, merged using unique app IDs from the App Store. Figure 8 illustrates app usage patterns across countries, with monthly active users in December 2020 and time spent in January 2021. On average, a smartphone user actively uses 15 apps on her phone and spends 6 hours per month on each app,[18] amounting to approximately 850 total hours spent on apps. This is a non-trivial utility component for consumers. Table 2 lists the top 40 apps by global time-spent on apps. First, the distribution is highly skewed: total hours spent on the most popular app, YouTube, reaches 390 billion hours, while the 40th most popular app, Reddit, are 3 billion hours. Second, the market is predominantly dominated by apps from the United States and China. Third, among the top 40 apps, China has banned nearly all foreign apps, with the exceptions of Microsoft and VK. These features create a natural experimental setting to examine the impact of bans under the GFW policy.

4.3. Reduced-form evidences

Lacking in information in the initial equilibrium before the policy intervention, I cannot directly test the predictions in equation 22, 23 and 24. Instead, I construct control groups to isolate the impact of the GWF. First, consistent with the trade literature, the trade barriers can be inferred from the difference between domestic market share and foreign market share – formally:

$$\begin{aligned} \Delta s_{i,od} &= s_{i,oo} - s_{i,od} \\ &= \frac{\theta \Omega_i}{M} (\ln \tau_{od} + \ln \overline{\tau}\, \mathbb{1}_{i,\text{policy}} - \Delta \ln P_{i,od}) \end{aligned} \tag{25}$$

[17] See details in appendix Appendix D

[18] The average time is calculated using market share as weights. The unweighted average time spent per app per user is 1.2 hours per month, reflecting intensive use on popular apps.

Table 2: Top 40 apps by worldwide time spent in 2021

	Name	Exporter	Worldwide time spent (bn hours)		Name	Exporter	Worldwide time spent (bn hours)
1	YouTube*	US	390	21	VK	RU	8
2	TikTok^{+}	CN	188	22	Google Maps*	US	8
3	Facebook*	US	142	23	Spotify*	SE	6
4	WhatsApp*	US	114	24	Discord*	US	6
5	Instagram*	US	109	25	Pinterest*	US	6
6	WeChat^{+}	CN	60	26	Waze*	US	6
7	Messenger*	US	29	27	Wattpad*	CA	6
8	Snapchat*	US	29	28	Redbook	CN	5
9	Twitter*	US	26	29	Zalo	CN	4
10	QQ^{+}	CN	21	30	Tencent Video	CN	4
11	Chrome*	US	19	31	Gmail*	US	4
12	iQiyi	CN	19	32	Bilibili	CN	4
13	Netflix*	US	14	33	KakaoTalk^{+}	KR	4
14	LINE*	JP	14	34	Microsoft Teams	US	4
15	Baidu	CN	11	35	Kugou Music	CN	3
16	Kwai	CN	11	36	ZOOM^{+}	US	3
17	Youku	CN	11	37	Yahoo! JAPAN*	JP	3
18	Twitch*	US	9	38	YouTube Kids*	US	3
19	Taobao	CN	9	39	Alipay	CN	3
20	Telegram*	AE	9	40	Reddit*	US	3

Notes: This Table lists the top40 apps ranked by their worldwide time spent in 2021. $*$ indicates apps banned in China and $+$ indicates apps not banned but under restrictions with the services. Almost all foreign big apps are banned or partially banned in China, while two exceptions are VK from Russia and Microsoft Teams from the U.S.

Figure 8: App usage by country

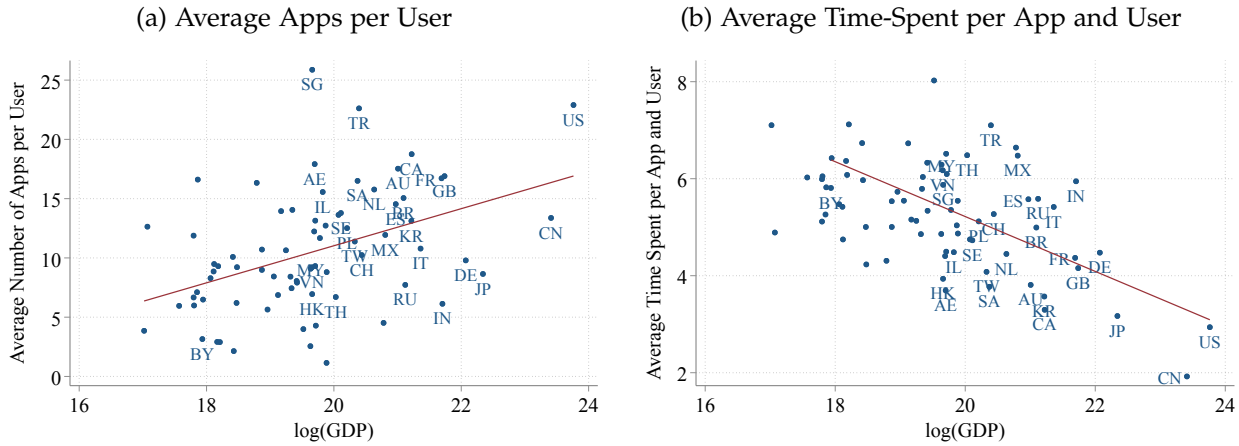

Notes: This figure plot the average apps per user and average time-spent per app and user by country in 2020. Each dot represent a country. The GDP data is from CEPII. There are 30 countries labelled with 2 digital ISO code. They have the largest market share in app market.

where $\Delta \ln P_{i,od}$ is the difference in the multilateral resistant term. Consider two products i and i' developed from the same country o, occupying the same position in attribute space, while product i being banned in market d and product i' not. The net effect of ban is simply given by

$$\ln \overline{\tau} = \frac{M}{\theta \Omega_i} \left(\Delta s_{i,od} - \Delta s_{i',od} \right) \tag{26}$$

The differences eliminate the potential omitted variables including culture barriers, IPRs, productivity and etc. Therefore, for each banned app, I find their control group as the apps from the same exporting country and in the same model predicted category (same colour block in Figure 7). For example, the control group for Chrome is Bing and Firefox. The empirical specification is

$$\Delta s_{i,oCN} = Ban_i + \lambda_c + \lambda_o + \varepsilon_{i,od} \tag{27}$$

where Ban_i is an indicator for whether the app is banned in China, λ_c is the app cluster fixed effects, λ_o is the exporting country fixed effects and $\varepsilon_{i,od}$ is the error term. Table 3 the market shares in the domestic country, mainland China, and Hong Kong for both the ban and control groups respectively. The banned apps are significantly larger than those in the control group, suggesting that the GFW policy is not solely about information censorship; major players with similar functionalities are more likely to be banned. Moreover, both groups have smaller market shares in mainland China and Hong Kong compared to their domestic markets, indicating that gravity factors such as language

Table 3: Summary statistics for market shares

	Observations	mean	25^{th} percentile	medium	75^{th} percentile
$s_{i,oo}$					
Ban Group	78	16.386	0.325	3.245	18.816
Control Group	781	1.146	0.005	0.033	0.282
$s_{i,o-CN}$					
Ban Group	78	3.552	0.001	0.132	1.735
Control Group	781	0.272	0	0	0.009
$s_{i,o-HK}$					
Ban Group	78	13.677	0.102	1.605	15.609
Control Group	781	0.328	0	0	0.009

barriers, cultural differences, and localization still play a role. However, the banned apps experience a relatively larger market share loss in mainland China.

Table 4 presents the estimation results. In column 1 to 3, I include different fixed effects, and the results are robust across these specifications. Compared to the control group, banned apps lost an additional 11 percentage points in market share in China. Column 4 replicates the specification from column 3 but includes all supplementary apps. While the estimated effect is smaller, it remains both statistically and economically significant. In column 6, I replace the dependent variable with $\Delta s_{i,o-HK}$ and find that the market share of banned apps is not significantly different from that of apps in the control group in Hong Kong. This suggests that the identified coefficient captures the causal impact of the GFW policy. The magnitude of this effect is substantial.

Table 4: The effect of bans on foreign apps in China (mainland) and Hong Kong

	$\Delta s_{i,o-CN}$				$\Delta s_{i,o-HK}$
	(1)	(2)	(3)	(4)	(6)
Ban	11.72***	11.51***	11.06***	7.46***	1.09
	(0.41)	(2.47)	(2.56)	(1.54)	(1.02)
Observations	859	859	859	2,335	848
R^2	0.138	0.163	0.172	0.116	0.064
Fixed Effects	-	*c*	*c*, *o*	*c*, *o*	*c*, *o*

Next, I further investigate the impacts on Chinese apps competing with the banned foreign apps. I categorize apps into approximately 300 narrowly defined categories based

Table 5: The effect of the GFW on the entry and market share of Chinese apps

	Number of Product		Market share	
	(1) Main	(2) All	(3) Main	(4) All
ban	10.79***	5.83***	14.27***	7.02***
	(1.29)	(0.70)	(3.71)	(1.93)
Observations	222	293	222	293
R^2	0.051	0.038	0.032	0.029

on their similarities and classify these categories into two groups: those with banned foreign apps and those without. I expect that more Chinese apps enter and that these apps are larger in size in categories with banned apps. Referring to the illustrative example in Figure 5, the differences between app categories resemble the distinction between regions near banned products and regions near other products, reflecting variations in choke prices. The empirical model is specified as follows:

$$y_{c,CN} = Ban_i + \lambda_{\bar{c}} + \varepsilon_c \tag{28}$$

where $y_{ct,CN}$ denotes the number of Chinese apps and their market share in category c, $\lambda_{\bar{c}}$ denotes the fixed effects of categories defined in the App Store and ε_c is the error term. Table 5 presents the estimation results. Columns 1 and 3 show results for app categories with at least one main banned app, while columns 2 and 4 include categories with at least one supplementary app as well. Categories with banned foreign apps have around 10 additional Chinese apps, and their market shares are 15 percentage points higher compared to categories without bans. The median size of an app category is 102 products with 0.63 million users, indicating that the effect is substantial.

To summarize, the reduced-form evidence supports the main predictions on the impacts of granular bans, as presented in equations 22, 23 and 24. To understand the welfare implications, I will establish a full quantitative model and calibrate the model to data in the next section.

5. Quantification

5.1. Quantitative Model

This section introduces a quantitative trade model incorporating both physical goods and app services, where app service producers engaging in Bertrand competition as

described in the baseline model. I rationalize the monetization of app sector through an Ad Exchange (Greenwood, Ma, and Yorukoglu, 2024, Kopytov, Roussanov, and Taschereau-Dumouchel, 2023, Rachel, 2021). [19] In this framework, app producers offer their services to consumers for free and sell user attention as ad slots to advertisers, with a market-specific ad price. The advertisers are physical goods producers who endogenously determine their ad expenditure. Both physical goods and app services are tradable across countries, while ads are non-tradable.

Demand The utility function of a representative consumer from country d is given by:

$$\ln u_d = \beta \ln Q_d + (1-\beta) \ln l_d \tag{29}$$

where Q_d represents consumption of physical goods, and l_d represents leisure consumption, specifically from mobile app services. Each consumer has one unit of time that can be allocated either to work or to using mobile apps. The budge constraint for consumers is given by

$$P_d^p Q_d = w_d (1 - h_d) + \pi_d \tag{30}$$

where P_d^p is the price index for physical goods, h_d is the time spent on app usage, w_d is the wage rate and π_d is profit rebate to local workers.

App Services Potential app entrants are genrated from exogenous random processes based on exporting country o and category c. Specifically, the set of potential entrants $\mathcal{M}^{o,c}$ are generated as follows: (1) the number of potential entrants $M^{o,c} = \|\mathcal{M}^{o,c}\|$ is drawn from Poisson process with the scale parameter $\overline{M}^{o,c}$; (2) product attributes $\mathbf{a}_i$ is drawn from vMF distribution $f_K(\mathbf{a}, \boldsymbol{\mu}^{o,c}, \kappa)$; and (3) productivity φ_i is drawn from Pareto distribution $h(\varphi, \overline{\varphi}^{o,c}, \nu)$. The set of potential entrants from country o is given by $\mathcal{M}^o = \mathcal{M}^{o,1} \cup \mathcal{M}^{o,2} \cup \cdots \cup \mathcal{M}^{o,C}$. I allow for endogenous entry across exporting country o and category c blocks.

App producers employ software engineers to produce services that are free in monetary terms but require time from consumers. The "price" of an app is defined as the amount of time required to gain one unit of leisure utility. For example, if a consumer's goal is to find a piece of news, the "price" of using Chrome is measured by the time spent searching. The "markup" reflects the fact that app producers intentionally slow down user interactions to maximize the time spent on the app, thereby increasing the "eyeballs" available to advertisers.

[19] This is a common modelling approach for leisure enhancing technologies and services, including television, web browsing, and mobile apps.

Taking wage rate w_o and total screen-time h_d in each country as given, app producers engage in the two-stage game described in the baseline model and equilibrium is solved by vectors of $\{\mathcal{M}_d, \{p_{i,od}, q_{i,od}\}_{i\in\mathcal{M}}\}_{d=1}^{N}$.The leisure utility derived from using apps in market d is given by $l_d = h_d/P_d^a$, where P_d^a represents the effective cost of the app usage and follows the Translog expenditure function specified in equation 9. I explicitly consider the trade costs from the GWF such that $\mathbb{1}_{i,\text{GFW}} = 1$ if app i is banned in China.

Physical Goods I introduce the physical goods in the simplest way possible, namely, using a Krugman (1980) model. There are one unit of continuum goods from each country o. The consumption on physical goods in market d be expressed as:

$$Q_d = \left(\sum_o^N \int b_{od}^{\psi}(i) q_{od}^P(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}} \tag{31}$$

where $q_{od}^P(i)$ is the quantity of product i from country o, σ denotes the elasticity of substitution in physical goods sector, $\psi \in (0, 1)$ is a parameter governs the efficiency of advertisement, and $b_{od}(i)$ is the brand equity built from advertisement investment associated with product i (Cavenaile and Roldan-Blanco, 2021, Rachel, 2021). [20] All producers from country o use the same production technology and choose prices and expenditure on advertisement to maximize their profit in each market:

$$\Pi_{od}^P = \max_{b_{od}, p_{od}} \left\{x_{od}^P - c_{od}^P q_{od}^P - b_{od}\right\} \tag{32}$$

where x_{od}^P is trade flow from o to d, $q_{od}^P = T_o^P L_{od}^P$ where T_o^P is the productivity and $c_{od}^P = w_o \tau_{od}^P$ where τ_{od}^P is the iceberg costs in physical goods sector. The optimal monopolistic constant markup is solved as $\sigma/(\sigma-1)$, and the optimal advertisement expenditure is solved as a fraction of trade flows from o in d:

$$b_{od} = \frac{\psi}{\sigma-1} x_{od}^P \quad \text{and} \quad x_{od}^P = \frac{b_{od}^{\psi}(w_o^P \tau_{od}^P / T_o^P)^{1-\sigma}}{\sum_{o'} b_{o'd}^{\psi}(w_{o'}^P \tau_{o'd}^P / T_{o'}^P)^{1-\sigma}} E_d \tag{33}$$

Ad Exchange There is a single Ad Exchange platform setting a uniform price of ads in each market d (Decarolis and Rovigatti, 2021).[21] For simplicity, I assume that the platform charges no profits. Ads can only be sold locally. App producers sell the attention from the consumers in market d to the physical goods producers operate in the market. The

[20] This demand is a special case in Rachel (2021) where $\chi = 1$. An alternative approach is to model advertisement of searching cost to reach more consumers as in Arkolakis (2010). We adopt this approach in our prior work Sun and Trefler (2023). The two approaches do not change our main conclusions but this way is computationally easier

[21] In this case, the Ad Exchange is Apple Transparent Tracking for iOS apps

monetary profit of app i acquired from market d is given by

$$\Pi_{i,od} = p_d^b \left(s_{i,od} h_d\right) - c_{i,od} q_{i,od} \tag{34}$$

where p_d^b is the price for ads in market d. The ads price is pinned down when the total expenditure on ads equals to the total revenue of app sector in each market:

$$p_d^b h_d = \sum_o^N b_{od} = \frac{\psi}{\sigma - 1} E_d \tag{35}$$

General Equilibrium is vectors of wages and labour supplies that clear labour markets, app markets, physical goods markets and ads markets in all countries. The consumer achieve the optimal time allocation following the *F.O.C.* such that

$$\frac{w_d h_d}{P_d^p Q_d} = \frac{1 - \beta}{\beta} \tag{36}$$

which pins down the screen-time of consumers in market d. Suppose manufacturing workers and software engineers are perfectly substitutable, and the total number of labour forces in each country is fixed and denoted by $\overline{L}_d$. The labour markets clear when

$$\sum_d^N \frac{q_{od}^P}{T_o^P} + \sum_{i \in \mathcal{M}^{o,c}} \sum_d^N \frac{q_{i,od}}{\varphi_i} = H_d^P + H_d^E = (1 - h_d)\overline{L}_d \tag{37}$$

5.2. Estimation Procesure

In this section, I discuss the estimation procedure of parameters in the quantitative model and model fit to data.

5.2.1. Additional assumption and data

Table 6 summarizes all parameters in the quantitative model by sectors. I use the vMF distribution and Pareto distribution to characterize the product-specific attributes **a** and productivity φ. I further parameterize α_i as $\alpha_i = \frac{\Omega_i^{-\eta}}{\sum_j \Omega_j^{-\eta}}$ where $\eta \geq 0$. This expression captures the "tech enthusiasm" phenomenon, where consumers value product uniqueness and assign it greater weight in the utility function. When $\eta = 0$, meaning $\alpha_i = 1/M$, the differences in market share only comes from the quadratic terms, capturing the competition effect. These assumptions enable me to translate high-dimensional estimation at the product level into estimation on low-dimensional statistics.

The product generation processes occur at country-category level.Figure 9 shows global time spent on apps by exporting country, revealing a high concentration of exports

Figure 9: The density of pairwise elasticity of substitution

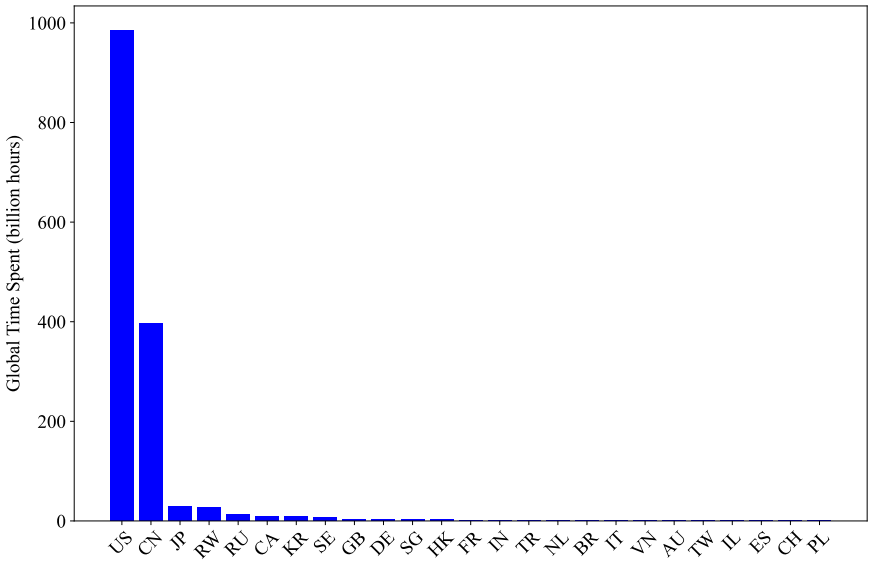

Notes: The global time spent is the summation of time-spent aross apps produced by exporting countries and all markets in 2021

in the US and China. Consequently, I model a world with three economies ($N = 3$): the United States, China and the rest of the world (RoW).using an Agglomerative Clustering method based on embeddings, so the random processes occur independently within each of the $N \times C = 30$ blocks. The model scales with nominal GDP, I normalize the nominal GDP in China to 1. The labour force simply determines the unit of labour, so I also normalize $\overline{L}_{CN} = 1$.

I estimate the parameters using a cross-sectional data for three economies in 2021. The estimation focuses on 24 economies that account for 98.3% of global app service exports, 84% of physical goods exports, 81% of world GDP, and 62% of global employment. The 22 economies except US and China are aggregated into one representing rest of the world. The core of my analysis relies on the dataset described in Section 4.2, which captures bilateral trade flows measured by time spent on apps by product and market. Additionally, data on bilateral trade flows for physical goods, GDP, population, and gravity covariates are sourced from the CEPII BACI Database. Employment data and average annual working hours are taken from the Penn World Table 10.0, while the ICT sector employment share is drawn from the OECD Going Digital Toolkit, supplemented by national statistics bureaus for economies not covered by the OECD.

Table 6: Summary of Parameters

	Mobile App	Physical Goods	Labour
Demand	$\{ \boldsymbol{\mu}^{o,c}, \kappa, \eta, \theta \}$	$\{ \sigma, \psi \}$	$\{ \beta \}$
Supply	$\{ \overline{M}^{o,c}, \overline{\varphi}^{o,c}, \nu \}$	$\{ T_o, \tau_{od}^P \}$	$\{ \overline{L}_o \}$

5.2.2. Estimation strategy

I estimate the model parameters in two steps. First, I calibrate the parameters that can be directly inferred from observed data moments or identified externally. In the second step, I use the simulated method of moments (SMM) to estimate the remaining five parameters. Due to the model's granularity, the simulation relies on realizations of the random processes. I fix the random seeds throughout the estimation and conduct robustness checks using different random seeds. I start with introducing the separately identified parameters and then discuss the SMM procedure.

Estimating $\widetilde{\theta}$ is a key scale parameter in the pairwise EoS. The estimates of $\widetilde{\theta}$ is inferred from the event study presented in Section 2. The empirical specification corresponds to the model moment shown in equation 24, which gives a structural interpretation of the estimated coefficient. Particularly, the model predicted changes in market share are:

$$s_{i\text{-}post,IN} - s_{i\text{-}pre,IN} = \widetilde{\theta} \left(\sum_{j|\mathbb{1}_{j,\text{policy}}=1} \omega_{ij} \right) \ln \overline{\tau} = \lambda_i \, \widetilde{\theta} \times ln\overline{\tau} \tag{38}$$

The corresponding empirical specification is as follows:

$$s_{it,IN} = \beta \times \lambda_i \times \mathbb{1}_{\tau \geq Jun\text{-}20} + \delta_i + \varepsilon_{it} \tag{39}$$

where the estimated coefficient $\widehat{\beta}$ reflect the value of $\widetilde{\theta} \times ln\overline{\tau}$ in the model. Table 7 presents the estimation results. Column 1 shows results with only the post-period dummy as the explanatory variable, while columns 2 to 4 include interactions between the post-period dummy and λ_i. After controlling for the interaction term, the coefficients on the dummy become negative and insignificant. In addition, the explanatory power of the empirical model, measured by within R^2, increases when I restrict the sample to large apps.

The unique nature of digital bans provides a reasonable interpretation of $\ln \tau$. After the ban, local consumers still can get access to the banned services via VPN. However, using a VPN can slow down internet speeds due to factors like data encryption, server distance, and network congestion. On average, VPN usage can reduce speeds by 30% (for a paid VPN at \$10 USD/month) to 50% (for free VPNs), although the exact impact

Table 7: Dependent Variable: $s_{it,IN}$

	(1)	(2)	(3)	(4)
$Post_t \times \omega_{ij}$		0.01	0.04	0.17
		(0.00)	(0.01)	(0.06)
$Post_t$	0.02	-0.17	-0.46	-1.97
	(0.01)	(0.05)	(0.09)	(0.80)
Observations	165,227	165,227	129,965	10,501
R^2	0.99	0.99	0.99	0.99
Within R^2	0.0001	0.0002	0.0004	0.0016
Fixed Effects	app	app	app	app
Sample	all	all	weighted	$s \geq 10^{-5}$

varies based on the VPN service and server location; higher-quality VPNs experience less slowdown compared to free or poorly optimized ones. I set the shock magnitude to $\ln \tau = 0.69$. Consequently, the estimated $\widetilde{\theta} = \hat{\beta} / \ln \tau = 1.7 \times 10^{-5} / 0.69 = 2.46 \times 10^{-5}$. Figure 10 plots the kernel density of the pairwise EoS in India before the ban, with a market share-weighted average EoS of 6.05.

Estimating vMF mean I assume that apps from a country o and category c are drawn from the vMF distribution $f_K(\mathbf{a}, \boldsymbol{\mu}^{o,c}, \kappa)$. Therefore, the resultant vector for apps $i \in \mathcal{M}^{o,c}$ satisfies the convergence condition:

$$E\left[\frac{\overline{\mathbf{a}}^{o,c}}{\|\overline{\mathbf{a}}^{o,c}\|}\right] = \boldsymbol{\mu}^{o,c} \tag{40}$$

with a sufficient large number of realizations. I can directly calibrate $\boldsymbol{\mu}_o^c$ by calculating the resultant vectors in data. The dimension of embedding vectors is fixed with the embedding model, which is 3,027 in this exercise. For computational efficiency, I reduce the dimensionality using the Truncated SVD Algorithm, while retaining most of the information on pairwise cosine similarity.[22] Figure 11 shows the mean square error (MSE) in predicting pairwise cosine similarity with lower-dimensional vectors. The MSE is less than 1×10^{-5} when $K \geq 500$, with only minor changes beyond this point. Thus, the additional dimensions do not contribute significantly to predicting cosine similarity between embeddings, and I use 500-dimensional vectors as proxies for $\overline{\mathbf{a}}^{o,c}$.

Estimating $\overline{M}^{o,c}$, $\overline{\varphi}^{o,c}$, τ_{od} and $\overline{\tau}$ I estimate the number of potential entrants, scale parameter in productivity and iceberg trade costs to match observed trade flows in the

[22] A more familiar approach is Principal Component Analysis (PCA), but the Truncated SVD Algorithm performs better in predicting cosine similarities.

Figure 10: The density of pairwise elasticity of substitution

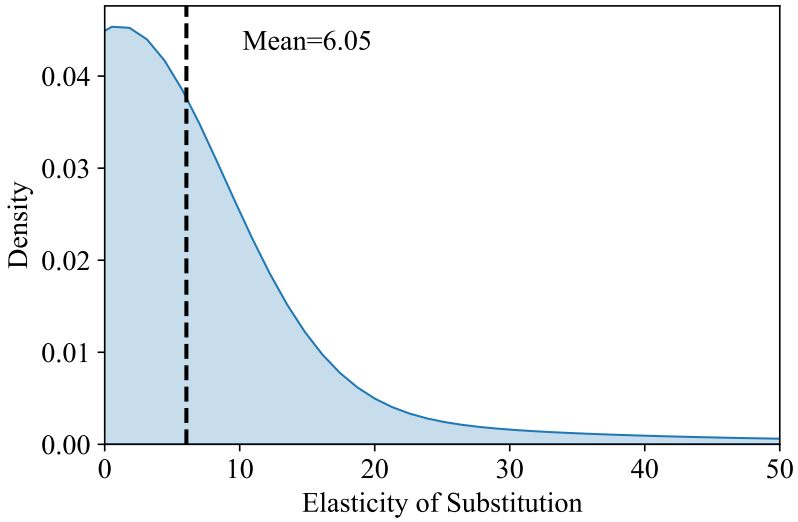

Notes: This figure is plotted by app usage data in India on May 2020. $\widetilde{\theta}$ is set to be 2.46×10^{-5}.

Figure 11: Mean Square Error (MSE) with low dimensionality

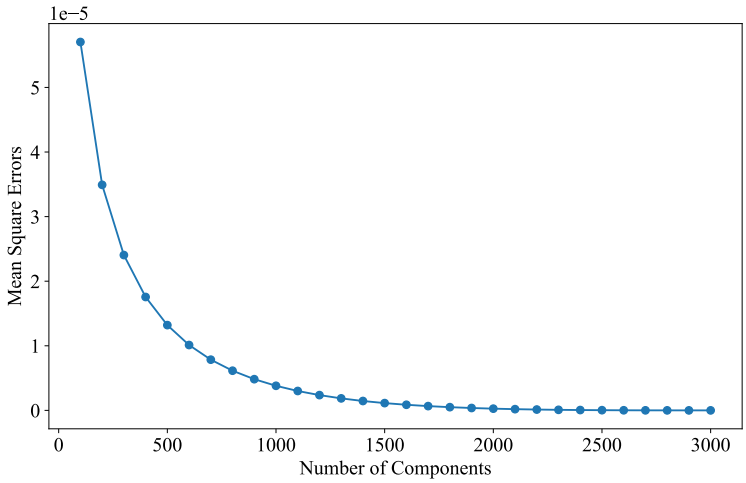

app sector. With the Poisson-Pareto productivity structure, any observed productivity distribution can be realized by various combination of $\overline{M}_o$ and $\overline{\varphi}^{o,c}$, making it impossible to identify these two parameters separately. However, the Translog functional form accounts for the effect of inactive firms on competition. I set all $\overline{M}^{o,c}$ to be the mean of $\overline{M}^{o,c}$ in the data, focusing sectoral differences on productivity without introducing an excess of shadow products into the economy.

As with many trade models, this model faces an identification challenge between productivity and unilateral trade costs. However, since the counterfactual analysis focuses on the extra costs from direct bans on U.S. apps in the Chinese market, the conclusions do not hinge on separating productivity from unilateral trade costs. In my approach, I

Table 8: Trade Costs in App Setor

	Importer		
	CN	US	ROW
CN	1.00	2.33	1.01
US	8.37	1.00	0.59
ROW	5.56	1.55	1.00

use the market share in RoW as a reference point, take the auxiliary value of the relative wage as given, and normalize $\varphi^{CN,1}$ to be 1. Table 8 presents the estimation results. In addition, the extra trade costs due to the GFW are estimated as $\overline{\tau} = 10$, indicating that banned U.S. apps face trade costs 10 times higher than those not banned.

Estimating σ and ψ σ and ψ are identified from matching two data moments. The gravity equation derived from 41 is as follows:

$$\frac{\pi_{od}^P}{\pi_{dd}^P} = \left(\tau_{od}^P \frac{w_o^P / T_o^P}{w_d^P / T_d^P} \right)^{\frac{1-\sigma}{1-\psi}}. \tag{41}$$

where the trade elasticity of physical goods is given by $(1-\sigma)/(1-\psi)$. I borrow the externally identified value from the literature (e.g., Caliendo and Parro (2015)) and let $(1-\sigma)/(1-\psi)$ to be 4. Second, the expenditure share on advertisement is given by $\psi/(\sigma-1)$ from equation 35. I let the parameter match the average expenditure share in data and set $\psi/(\sigma-1)$ to be 0.05. These two data moments fix the value of σ and ψ.

Estimating T_o^P and τ_{od}^P I estimate the T_o^P and τ_{od}^P from the bilateral trade flows in physical goods sectcor. I follow Head and Ries (2001) to estimate bilateral trade costs. Assume that $\tau_{od}^P = \tau_{do}^P$ and given trade elasticity, we can back out symmetric trade barriers as the symmetric Head-Ries index:

$$\tau_{od}^P = \left(\frac{\pi_{od}^P}{\pi_{dd}^P} \frac{\pi_{do}^P}{\pi_{oo}^P} \right)^{\frac{1-\psi}{2(1-\sigma)}} \tag{42}$$

Taking the auxiliary value of the relative wage as given, I can estimate T_o^P as the residual in equation 41 with a normalization $T_{CN}^P = 1$. The relative productivity between two normalizers, $\varphi^{CN,1}/T_{CN}^P$, is calibrated in SMM estimation.

Table 9 summarizes all parameters in the first group. These parameters are separately identified by a direct match with observable data moments. The targeting data moments and value of parameters are displayed in the table.

Table 9: Summary of parameters separately identified

Parameter	Interpretation	Data Moments/Source	Value
App sector			
$\widetilde{\theta}$	scale paramter in EoS	$\hat{\beta}/\ln\overline{\tau}$	2.46×10^{-5}
$\boldsymbol{\mu}^{o,c}$	unit resultant vectors	LLM embeddings	$\overline{\mathbf{a}}^{o,c}/\|\overline{\mathbf{a}}^{o,c}\|$
$\overline{M}^{o,c}$	number of potential entrants	the average number	150
τ_{od}	dyadic trade costs	trade shares	$s_{i,od}$
$\overline{\tau}$	GFW costs	banned apps	10.49
Physical goods sector			
$(\sigma-1)/(1-\psi)$	trade elasticity	Caliendo and Parro (2015)	4
$\psi/(\sigma-1)$	revenue share of apps	expenditure share on ICT	0.05
τ^P_{od}	country-pair iceberg costs	Head-Ries Index	$\left(\frac{\pi^P_{od}}{\pi^P_{dd}}\frac{\pi^P_{do}}{\pi^P_{oo}}\right)$

SMM Estimation We use simulated method of moments (SMM) to estimate the remaining five parameters of the model, $\Theta = \{\kappa, \nu, \eta, \beta, \varphi^{CN,1}/T^P_{CN}\}$. The estimation proceeds as follows: for a given parameter vector Θ, we simulate the model, compute a list of cross-sectoral moments $\mathbb{M}(\Theta)$, and contrast them with the equivalent moments in the data $\widetilde{\mathbf{m}}$. We search for parameter vector $\widehat{\Theta}$ that minimizes the distance between the model and data moments, according to the loss function $\mathcal{L}(\Theta) \equiv (\mathbb{M}(\Theta) - \widetilde{\mathbf{m}})' \mathbf{W} (\mathbb{M}(\Theta) - \widetilde{\mathbf{m}})$, where $\mathbf{W}$ is the weighting matrix. Specifically, we search for the best fitting points from the grid. The details for SMM procedure is described in appendix.

We target five groups of data moments, which corresponds the sectoral outcomes in three economies. Table 10 summarize the parameters and the main targeted moments for each parameter. First, the concentration parameter κ is identified by the dispersion of product attributes. Theoretically, the concentration parameter κ is the solution for

$$\overline{a} = \frac{I_{K/2}(\kappa)}{I_{K/2-1}(\kappa)} \tag{43}$$

Table 10: Summary of parameters in SMM

Parameter	Interpretation	Main targeted moments	Value
κ	spherical concentration	$var(\mathbf{a}_i)$	360
ν	Pareto shape parameter	tail distribution of s_i	1.21
η	love-of-innovativeness	$cov(s_i, \Omega_i)$	1.03
β	utility weight on real consumption	time-allocation	0.75
$\varphi^{CN,1}/T^P_{CN}$	relative producivity	employment shares	1.1

where $\overline{a} = \|\overline{\mathbf{a}}\| / M$ is the resultant vector length and I_v is the modified Bessel function of the first kind and order v. For a large number of observations, κ can be estimated using the following approximation (Sra, 2012) :

$$\hat{\kappa} \approx \frac{\overline{a}(K - \overline{a}^2)}{1 - \overline{a}^2} \tag{44}$$

However, the competition reduce the density in more central locations. Thus, the estimation of $\hat{\kappa}$ cannot be directly mapped to κ. Second, following Gaubert and Itskhoki (2021), I estimate the Pareto shape parameter ν from the tail distribution of market shares. Third, η captures the love-of-innovativeness from the demand side and target the covariance between market shares and product centrality. Fourth, the utility weight on real consumption is informative on the time-spent on apps relative to working hours. Fifth, the relative level of the two normalizations on productivities in China target to the employment shares of app sector relative to the manufacturing sector.

Model Fit Table 11 reports the model-based values of the moments used in the estimation, comparing them to their empirical counterparts. Overall, the model provides a reasonable fit to the data across all five groups of moments. Due to computational constraints, the current SMM estimation is performed on a parameter grid of only 20 points, so further refinements on a finer grid are likely to improve accuracy. First, the estimation closely matches the hedonic attribute dispersion for all active apps in each market. Notably, the model accurately replicates bilateral trade shares in both sectors. Figure 12 shows the bilateral trade shares predicted by the model against those observed in the data for the physical goods and app sectors. The points cluster tightly around the 45-degree line, indicating a strong match. However, the model predicts lower skewness in market shares than the data exhibit, which may stem from its assumption of constant returns to scale and the absence of externalities. Second, the model aligns well with observed time allocation and labor allocation across countries and sectors. Third, it also performs well in predicting non-targeted moments, such as relative wages and relative nominal GDP.

6. Counterfactuals

In this section, I use the quantitative model to evaluate the aggregate and distributional implications of trade interventions targeting specific foreign products. I focus on the policy-induced trade costs, represented by $\ln \overline{\tau}\, \mathbb{1}_{i,\text{policy}_d}$, to replicate the effects of actual bans. These additional costs operate over and above the uniform dyadic trade costs and any productivity differences, ensuring that the quantitative results are not tied to the identification of baseline trade barriers or productivity parameters.

Table 11: Moments in data and estimation

	Moments		Data $\widetilde{\mathbf{m}}$	Model $\mathbb{M}(\Theta)$
Targeted Moments				
1.	Variance of hedonic attributes	$var(\mathbf{a}_{i\in\mathcal{M}})$	0.6074	0.6235
	– active apps in China	$var(\mathbf{a}_{i\in\mathcal{M}_{CN}})$	0.6733	0.6564
	– active apps in the US	$var(\mathbf{a}_{i\in\mathcal{M}_{US}})$	0.5921	0.5432
	– active apps in RoW	$var(\mathbf{a}_{i\in\mathcal{M}_{ROW}})$	0.5514	0.5259
2.	Top-product sales share	$\overline{s}_{1,od}$	0.386	0.164
	Top-3 product sales share	$\sum_{i=1}^{3}\overline{s}_{i,od}$	0.620	0.382
	Top-5 product sales share	$\sum_{i=1}^{5}\overline{s}_{i,od}$	0.687	0.454
	Top-10 product sales share	$\sum_{i=1}^{10}\overline{s}_{i,od}$	0.783	0.511
3.	Covariance of market share and centrality	$cov(s_i, \Omega_i)$	-0.05	-0.08
4.	Average time-spent on apps	h_d		
	– time-spent in China	h_{CN}	0.274	0.295
	– time-spent in the US	h_{US}	0.326	0.309
	– time-spent in RoW	h_{RoW}	0.335	0.313
5.	Employment share of app sector	$L_d^A/\overline{L}_d$		
	– employment in China	$L_{CN}^A/\overline{L}_{CN}$	0.035	0.059
	– employment in the US	$L_{US}^A/\overline{L}_{US}$	0.034	0.013
	– employment in RoW	$L_{RoW}^A/\overline{L}_{RoW}$	0.083	0.051
Non-targeted Moments				
1.	Number of active products			
2.	Relative wage	w_d		
	– US/China	w_{US}/w_{CN}	1.57	1.29
	– RoW/China	w_{RoW}/w_{CN}	6.40	6.21
3.	GDP	E_d		
	– US/China	E_{US}/E_{CN}	2.27	1.60
	– RoW/China	E_{RoW}/E_{CN}	1.42	1.28

Notes: This table displays the data moments using the estimation in Table 9 and Table 10. The current is subject to searching on 20 grids and will be improved with finer searching.

Figure 12: Model fit in bilateral trade flows

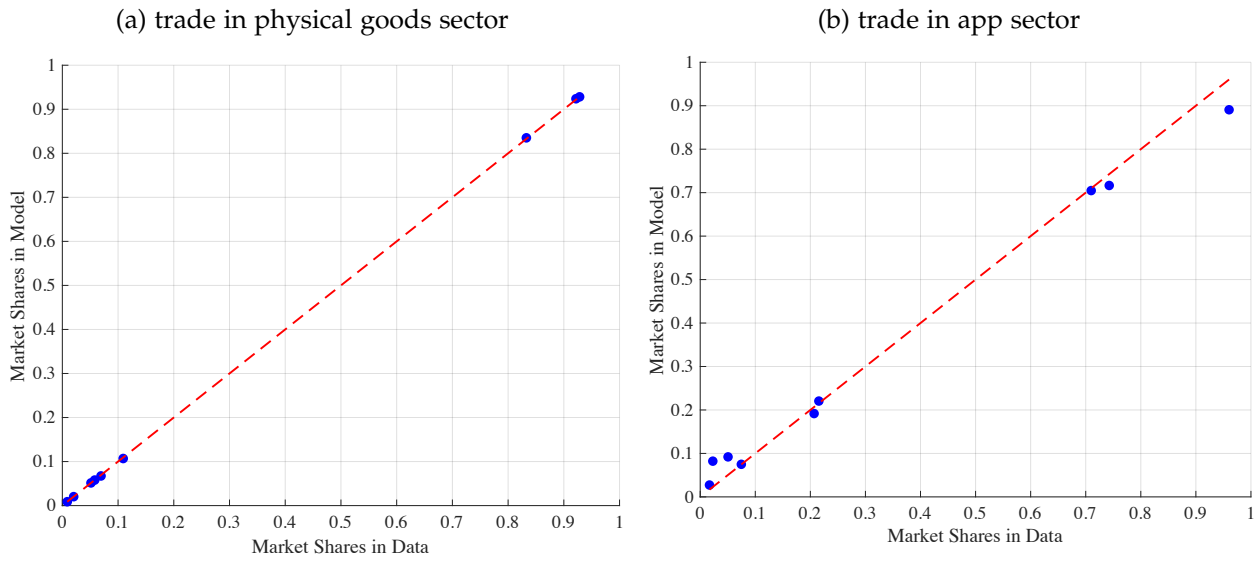

Notes: This figure plot the average apps per user and average time-spent per app and user by country in 2020. Each dot represent a country. The GDP data is from CEPII. There are 30 countries labelled with 2 digital ISO code. They have the largest market share in app market.

I conduct two counterfactual exercises using the estimated parameters. First, I consider the absence of China's GFW policy by removing the policy-induced trade barriers. This exercise allows me to decompose the resulting welfare changes for Chinese consumers. Second, I examine the impact of a ban on TikTok – China's most productive social networking appâĂŤon market outcomes. In this exercise, I uncover the ignored distributional effect among competing apps as a result of heterogenous substitution patterns.

6.1. The impact of the Great Firewall policy

Armed with the estimated model, I now quantify the impact of the GFW policy in China. I consider a counterfactual scenario where $\overline{\tau} = 1$ and take the difference between data and the counterfactual outcome as the impact of the GFW policy. Table 12 displays the changes in welfare components. The changes in x is defined as $x^{\text{Data}}/x^{\text{Counterfactual}} - 1$. Quantitatively, the GFW policy has two distinct effects on welfare. First, GFW caused rent shifting from US to China. Namely, it increased Chinese real incomes by 1.05% while reducing US real incomes by 0.81%. However, this benefit in China was more than offset by a 29% loss in leisure utility within the app sector. Given the calibrated utility weight on real income (75%), the net welfare loss from the GFW policy is 7.48% for Chinese consumers.

Table 13 further decomposes the sources of welfare changes. The income effect primarily arises from rent shifting within the app sector: the number of active Chinese

apps increase only by 0.34% in China, while profits rose by 7.65%, driven by head-to-head oligopolistic competition. For instance, Tencent would not have reached a valuation of $486.24 billion if Meta had remained in China. The welfare loss in leisure stems from using inferior domestic substitutes, such as replacing Chrome with Baidu. While the overall time spent on apps remains constant, the effective cost rises sharply by about 40%. For example, a Chinese consumer receives 40% less valuable information, given equal search time on Baidu compared to a US consumer using Chrome. Importantly, these findings suggest that policymakers may have an incentive to impose such bans on foreign tech companies if their objective is to boost real income. However, this policy overlooks substantial welfare losses on the consumer side.

Table 12: Welfare decomposition of the GFW by country

Country	Welfare	Real Income	Leisure Utility
China	-7.48%	1.05%	-29%
US	0.11%	-0.81%	2.88%
RoW	0.59%	0.00%	2.38%

In addition, I examine the impact on the level of product differentiation among active products. The model exhibits a novel and continuum measure of product innovation by the dispersion of product attributes. I find that the product dispersion among Chinese apps decreased by 4.71%, indicating that the policy not only replaced foreign products with inferior domestic alternatives but also led to the survival of more homogenous goods with fewer distinctive hedonic attributes. This finding suggests that reduced competition stifles product innovation. It aligns with the finding in Hsieh *et al.* (2021), and rather than assuming exogenous innovation efficiency, the results here demonstrate that lower innovation can be an endogenous outcome of reduced competition.

Table 13: The impact of the GFW on Chinese app sector

	Leisure Utility		Real Income		Product
	Leisure hours	Effective costs	Profits	Entry	Dispersion
Changes	-0.63%	39.97%	7.65%	0.34%	-4.71%

6.2. The impact of ban on TikTok

Talk of banning TikTok began years ago as its influence grew in the U.S. In April 2024, President Biden issued a ban requiring ByteDance to sell the platform by January 19,

2025 – just one day before Inauguration Day – making it difficult for President Trump to reverse the decision. Meanwhile, Trump pledged to "never ban TikTok" and argued that "such a ban would only benefit Mark Zuckerberg and Facebook". In this counterfactual exercise, I simulate the policy by imposing a $\overline{\tau}$ barrier on the top Chinese social networking app, representing TikTok.

I examine the distributional effects of the ban on apps in the U.S. market. For each app, I measure its similarity to TikTok using $\omega_{i,\text{TikTok}}$. I then group apps into fifty bins, formed by intersecting deciles of similarity and quintiles of productivity. Figure 13 shows the average changes in revenue for each bin, where revenue changes are defined as $(x^{\text{Data}}/x^{\text{Counterfactual}}) - 1$, and the average change is normalized to 1. A key insight is that apps with high similarity to TikTok and high productivity gain more from the ban. In particular, apps in the highest similarity and productivity category (including Facebook and Instagram) increase their U.S. revenue by 11 times more than the average app. This result is consistent with the political debate surrounding TikTok, but would not emerge from a standard model assuming equal substitution patterns among all product pairs.

Figure 13: Changes in revenue on the US market

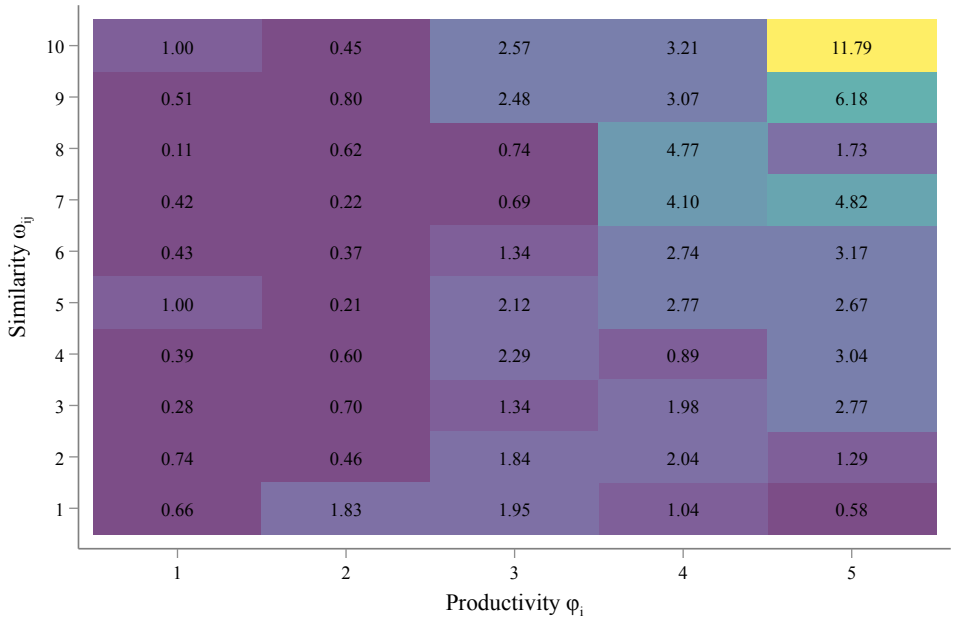

7. Conclusion

This paper examines the welfare implications of trade policies targeted at specific foreign products. The standard tools fall short in addressing this issue, as policy-targeted products typically have a vast user base and a differentiated mix of hedonic attributes,

leading to complex substitutability patterns with competitors. I find causal evidence of these substitution patterns through an event study: apps with high similarity to banned apps captured greater increase in market share after the India's ban on Chinese apps. More importantly, the substitution patterns can be identified using LLM-based similarity measures.

Building on this empirical finding, I develop a model in which large, single-product firms with heterogeneous productivities and hedonic attributes engage in Bertrand oligopolistic competition. I introduce a Translog expenditure function, in which the key parameters endogenously emerge from a micro-foundation of hedonic demand. The model is purpose-built to exploit textual product descriptions and large language models: pairwise elasticities of substitution are proportional to pairwise cosine similarities between product descriptions. I integrate the demand system with a granular trade model and solve for the static equilibrium through a two-stage game.

This model is estimated with a detailed dataset on mobile apps. Equipped with the estimated model, I quantify the welfare implications of China's Great Firewall policy: a censorship system restricting the local availability and functionality of major information-sensitive, foreign (mainly U.S.) apps. I find that the Great Firewall policy causally fostered Chinese tech giants by shifting profits away from their foreign counterparts, resulting in a 1.05% increase in Chinese real incomes. However, this gain was offset by the utility loss from the use of inferior apps, yielding a net welfare change of -7.48%. In addition, I uncover a novel interaction between trade barriers and product innovation (Romerian innovation): reduced competition from policy intervention led to a decline in product differentiation among Chinese domestic products.

This paper contributes to the growing literature on granular firms in international trade. The demand-side heterogeneity has been under-explored, despite its empirical importance alongside supply-side heterogeneity. This paper – both methodologically and empirically – provides a practical approach to capturing rich demand-side heterogeneity and demonstrates its relevance for welfare analysis. Moreover, the model and estimation strategy are easily implemented and widely applicable to research areas where rich textual data is available such as the occupational substitution in labour market or factor substitution in production.

References

Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson. 2017. Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review* 107(3):633–689.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira. 2023. Patents to Products: Product Innovation and Firm Dynamics. Working paper, Federal Reserve Bank of Atlanta.

Aridor, Guy. 2022. Drivers of digital attention: Evidence from a social media experiment. *Available at SSRN 4069567* .

Arkolakis, Costas. 2010. Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy* 118(6):1151–1199.

Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2019. The elusive pro-competitive effects of trade. *The Review of Economic Studies* 86(1):46–80.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. New trade models, same old gains? *American Economic Review* 102(1):94–130.

Baldwin, Richard and James Harrigan. 2011. Zeros, quality, and space: Trade theory and trade evidence. *American Economic Journal: Microeconomics* 3(2):60–88.

Banerjee, Arindam, Inderjit S Dhillon, Joydeep Ghosh, Suvrit Sra, and Greg Ridgeway. 2005. Clustering on the unit hypersphere using von mises-fisher distributions. *Journal of Machine Learning Research* 6(9).

Baqaee, David Rezza, Emmanuel Farhi, and Kunal Sangani. 2023. The darwinian returns to scale. *The Review of Economic Studies* :rdad061.

Bartik, Alexander, Arpit Gupta, and Daniel Milo. 2023. The costs of housing regulation: Evidence from generative regulatory measurement. Working paper.

Bernard, Andrew B., Esther Boler, Davin Chor, Sirig Gurund, and Wei Lu. 2024. The great firewall and knowledge diffusion. Working paper, EITI 2024.

Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito. 2019. Production networks, geography, and firm performance. *Journal of Political Economy* 127(2):639–688.

Bernard, Andrew B, Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2018. Two-sided heterogeneity and trade. *Review of Economics and Statistics* 100(3):424–439.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* 63(4):841–890.

Bian, Bo, Xinchen Ma, and Huan Tang. 2021. The supply and demand for data privacy: Evidence from mobile apps. *Available at SSRN* .

Boeing, Philipp, Loren Brandt, Ruochen Dai, Kevin Lim, and Bettina Peters. 2024. The anatomy of chinese innovation: Insights on patent quality and ownership. *ZEW-Centre for European Economic Research Discussion Paper* (24-016).

Bradford, Anu. 2023. *Digital empires: The global battle to regulate technology*. Oxford University Press.

Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, *et al.* 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* .

Caliendo, Lorenzo and Fernando Parro. 2015. Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies* 82(1):1–44.

Cavenaile, Laurent and Pau Roldan-Blanco. 2021. Advertising, innovation, and economic growth. *American Economic Journal: Macroeconomics* 13(3):251–303.

Christensen, Laurits R, Dale W Jorgenson, and Lawrence J Lau. 1975. Transcendental logarithmic utility functions. *The American Economic Review* 65(3):367–383.

Decarolis, Francesco and Gabriele Rovigatti. 2021. From mad men to maths men: Concentration and buyer power in online advertising. *American Economic Review* 111(10):3299–3327.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2007. Unbalanced trade. *American Economic Review* 97(2):351–355.

Diewert, W Erwin. 1976. Exact and superlative index numbers. *Journal of econometrics* 4(2):115–145.

Diewert, W Erwin. 2023. A generalization of the symmetric translog functional form. *Journal of International Economics* :103821.

Eaton, Jonathan, Samuel S. Kortum, and Sebastian Sotelo. 2012. International trade: Linking micro and macro. working paper, National bureau of economic research.

Fajgelbaum, Pablo, Gene M Grossman, and Elhanan Helpman. 2011. Income distribution, product quality, and international trade. *Journal of political Economy* 119(4):721–765.

Feenstra, Robert C and David E Weinstein. 2017. Globalization, markups, and us welfare. *Journal of Political Economy* 125(4):1040–1074.

Gaubert, Cecile and Oleg Itskhoki. 2021. Granular comparative advantage. *Journal of Political Economy* 129(3):871–939.

Greenwood, Jeremy, Yueyuan Ma, and Mehmet Yorukoglu. 2024. âĂIJyou will:âĂİ a macroeconomic analysis of digital advertising. *Review of Economic Studies* :rdae067.

Hallak, Juan Carlos and Peter K Schott. 2011. Estimating cross-country differences in product quality. *The Quarterly journal of economics* 126(1):417–474.

Hanson, Gordon H, Nelson Lind, and Marc-Andreas Muendler. 2015. The dynamics of comparative advantage. Technical report, National bureau of economic research.

Hausmann, Ricardo, Jason Hwang, and Dani Rodrik. 2007. What you export matters. *Journal of economic growth* 12:1–25.

Head, Keith and John Ries. 2001. Increasing returns versus national product differentiation as an explanation for the pattern of us–canada trade. *American Economic Review* 91(4):858–876.

Hoberg, Gerard and Gordon Phillips. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124(5):1423–1465.

Hsieh, Chang-Tai, Pete Klenow, and Kazuatsu Shimizu. 2021. Romer or ricardo? Working paper, National Bureau of Economic Research.

Hummels, David and Peter J Klenow. 2005. The variety and quality of a nation's exports. *American economic review* 95(3):704–723.

Jaccard, Torsten. 2023. Import demand in differentiated products. Technical report, Tech. rep., University of Toronto.

Janssen, Rebecca, Reinhold Kesler, Michael E Kummer, and Joel Waldfogel. 2022. Gdpr and the lost generation of innovative apps. Technical report, National Bureau of Economic Research.

Josifoski, Martin, Marija Sakota, Maxime Peyrard, and Robert West. 2023. Exploiting asymmetry for synthetic training data generation: Synthie and the case of information extraction. *arXiv preprint arXiv:2303.04132* .

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3):303–20.

Kopytov, Alexandr, Nikolai Roussanov, and Mathieu Taschereau-Dumouchel. 2023. Cheap thrills: the price of leisure and the global decline in work hours. *Journal of Political Economy Macroeconomics* 1(1):80–118.

Krugman, Paul R. 1980. Scale economies, product differentiation, and the pattern of trade. *American Economic Review* 70(5):950–959.

Lancaster, Kelvin J. 1966. A new approach to consumer theory. *Journal of political economy* 74(2):132–157.

Lashkari, Danial and Martí Mestieri. 2019. How substitutable are the products of different countries? working paper.

Lind, Nelson and Natalia Ramondo. 2023. Trade with correlation. *American Economic Review* 113(2):317–353.

Liu, Ruibo, Jerry Wei, Fangyu Liu, Chenglei Si, Yanzhe Zhang, Jinmeng Rao, Steven Zheng, Daiyi Peng, Diyi Yang, Denny Zhou, *et al.* 2024. Best practices and lessons learned on synthetic data for language models. *arXiv preprint arXiv:2404.07503* .

Mardia, Kanti V and Peter E Jupp. 2009. *Directional statistics*. John Wiley & Sons.

Matsuyama, Kiminori and Philip Ushchev. 2017. Beyond ces: three alternative classes of flexible homothetic demand systems. Working Paper 17-109, Global Poverty Research Lab.

Mayer, Thierry, Marc J Melitz, and Gianmarco IP Ottaviano. 2014. Market size, competition, and the product mix of exporters. *American Economic Review* 104(2):495–536.

Melitz, Marc J. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica* 71(6):1695–1725.

Pavcnik, Nina. 2002. Trade liberalization, exit, and productivity improvements: Evidence from chilean plants. *The Review of economic studies* 69(1):245–276.

Pellegrino, Bruno. 2023. Product differentiation and oligopoly: a network approach. Working paper, CESifo.

Petrin, Amil. 2002. Quantifying the benefits of new products: The case of the minivan. *Journal of political Economy* 110(4):705–729.

Rachel, Lukasz. 2021. Leisure-enhancing technological change. *Job market paper, London School of Economics* .

Rosen, Sherwin. 1974. Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy* 82(1):34–55.

Salop, Steven C. 1979. Monopolistic competition with outside goods. *The Bell Journal of Economics* :141–156.

Sra, Suvrit. 2012. A short note on parameter approximation for von mises-fisher distributions: and a fast implementation of $I_s(x)$. *Computational Statistics* 27:177–190.

Sra, Suvrit and Inderjit Dhillon. 2005. Generalized nonnegative matrix approximations with bregman divergences. *Advances in neural information processing systems* 18.

Sun, Ruiqi and Daniel Trefler. 2023. The impact of AI and cross-border data regulation on international trade in digital services: A large language model approach. Working paper, National Bureau of Economic Research.

Sutton, John and Daniel Trefler. 2016. Capabilities, wealth, and trade. *Journal of Political Economy* 124(3):826–878.

Trefler, Daniel. 2004. The long and short of the canada-us free trade agreement. *American Economic Review* 94(4):870–895.

Zhong, Shi and Joydeep Ghosh. 2003. A unified framework for model-based clustering. *The Journal of Machine Learning Research* 4:1001–1037.

Appendix

Appendix A. Derivations and Proofs

Appendix A.1. von Mises-Fisher Distribution

The circular distribution is particularly useful to summarize the high-dimensional unit sphere data by low-dimensional angular statistics and is widely used in machine learning studies. In this model, I use von Mises-Fisher (vMF) distribution to model the arrival of product blueprints. It is normal distribution defined on high dimensional sphere. Starting form a normal distribution with isotropic covariance $\kappa^{-1}\mathbf{I}$ and mean $\boldsymbol{\mu}$ in $K-$dimensional space $\mathbb{R}^K$, we can get the probability density function as follows:

$$f_K(\mathbf{x}, \boldsymbol{\mu}, \kappa) = \left(\sqrt{\frac{\kappa}{2\pi}}\right)^K \exp\left(-\kappa\frac{(\mathbf{x}-\boldsymbol{\mu})^T(\mathbf{x}-\boldsymbol{\mu})}{2}\right) \tag{45}$$

vMF distribution is simply obtained by conditioning on $\|\mathbf{x}\| = 1$ and $\|\boldsymbol{\mu}\| = 1$. By expanding

$$(\mathbf{x}-\boldsymbol{\mu})^T(\mathbf{x}-\boldsymbol{\mu}) = \mathbf{x}^T\mathbf{x} + \boldsymbol{\mu}^T\boldsymbol{\mu} - 2\mathbf{x}^T\boldsymbol{\mu} = 2(1-\mathbf{x}^T\boldsymbol{\mu}), \tag{46}$$

the pdf of Von Mises-Fisher distribution is recovered by adjust the normalization constant as a integral over unit sphere – formally:

$$f_K(\mathbf{x}, \boldsymbol{\mu}, \kappa) = C_K(\kappa)\exp(\kappa\boldsymbol{\mu}^T\mathbf{x}) \tag{47}$$

Consider a case where $K = 2$, the nodes are distributed on a unit circle (e.g., Salop (1979)). The vMF distribution collapses to von Mises distribution, whose *pdf* is given by

$$f_2(x, \mu, \kappa) = \frac{exp(\kappa\cos(x-\mu))}{2\pi I_0(\kappa)} \tag{48}$$

Figure A1 displays 10 random draws (red triangles) from a uniform distribution and 10 random draws (blue rounds) from von Mises distribution. Both of the distribution are centred at $(1, 0)$. In figure A1c, figure A1b and figure A1a, we set $\kappa = 0.1$, 4, and 10. When κ increases, the spherical variance of the nodes from the von Mises distribution decreases.

The vMF distribution's ability to model directional data makes it highly suitable for various tasks in text mining. Its applications range from topic modeling and clustering to improving information retrieval and NLP tasks, providing a robust framework for handling the unique characteristics of text data represented in high-dimensional spaces (Banerjee, Dhillon, Ghosh, Sra, and Ridgeway, 2005, Sra and Dhillon, 2005, Zhong and Ghosh, 2003).

Figure A1: Illustration of circular normal distribution

(a) $\kappa = 10$ (b) $\kappa = 4$ (c) $\kappa = 0.1$

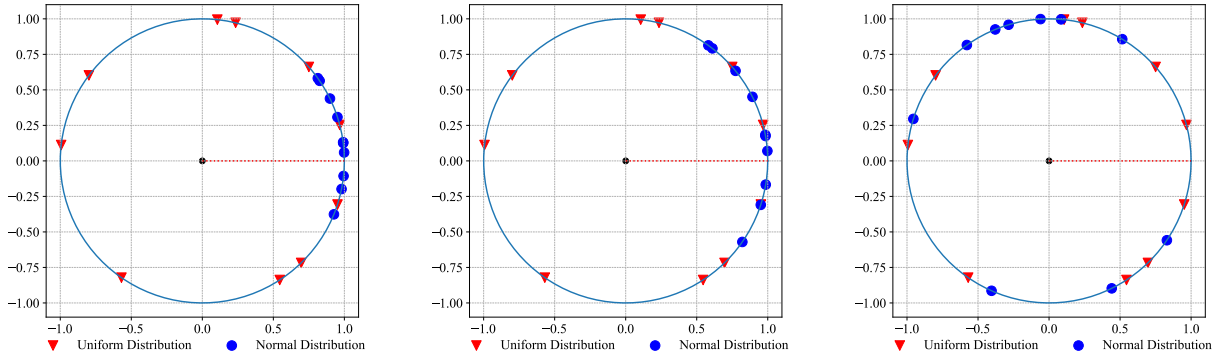

Notes: These figures plot 10 random nodes drawn from two distributions. The random seed is set to 10. The red triangles are drawn from a uniform distribution. The blue rounds are drawn from a circular normal distribution.

Appendix A.2. Generalized Linear Hedonic demand

Suppose there are M products, indexed by $i \in \{1, 2, \ldots, M\}$, in the economy. These products are distributed on the K dimensional attribute space and are described by the matrix $\mathbf{A}$. The price index for a product attribute s is defined as follows

$$\left(\sum_{i=1}^{M} a_{is}\right) \ln p_s = \sum_{i=1}^{M} a_{is} \ln p_i \tag{49}$$

We can rewrite equation 49 as vectors and matrices. The vector of attribute price, indexed by $\tilde{p}$, is a function of product matrix $\mathbf{A}$ and the vector of product price $\mathbf{p}$:

$$\ln \tilde{\mathbf{p}} = diag(\overline{\mathbf{a}})^{-1} \mathbf{A} \ln \mathbf{p} \tag{50}$$

where $\overline{\mathbf{a}}$ is the resultant vector of matrix $\mathbf{A}$, $\ln \tilde{\mathbf{p}} = [\ln \tilde{p}_1, \ \ln \tilde{p}_2, \ldots, \ln \tilde{p}_K]^T$ and $\ln \mathbf{p} = [\ln p_1, \ \ln p_2, \ldots, \ln p_M]^T$. The diagonal matrix $diag(\overline{\mathbf{a}})$ is a matrix where the s^{th} diagonal entry is $\overline{a}_s$ – formally given by:

$$diag(\overline{\mathbf{a}}) = \begin{pmatrix} \sum_i^M a_{i1} & 0 & \ldots & 0 \\ 0 & \sum_i^M a_{i2} & \ldots & 0 \\ \vdots & \ldots & \ldots & \vdots \\ 0 & 0 & \ldots & \sum_i^M a_{iK} \end{pmatrix}^T \tag{51}$$

The representative consumer face an expenditure cost, which is a quadratic function of attribute prices and product prices. The function is as follows:

$$\ln e \ = \ a_0 + \ln \mathbf{p}^T \boldsymbol{\alpha} + \frac{\theta}{2} \ln \tilde{\mathbf{p}}^T \tilde{\mathbf{D}} \ln \tilde{\mathbf{p}} - \frac{\theta}{2} \ln \mathbf{p}^T \mathbf{D} \ln \mathbf{p} \tag{52}$$

where $\widetilde{\mathbf{D}}$ and $\mathbf{D}$ are weight matrices and defined as

$$\widetilde{\mathbf{D}} = \begin{pmatrix} \frac{\overline{a}_1^2}{\|\overline{\mathbf{a}}\|} & 0 & \dots & 0 \\ 0 & \frac{\overline{a}_2^2}{\|\overline{\mathbf{a}}\|} & \dots & 0 \\ \vdots & \dots & \dots & \vdots \\ 0 & 0 & \dots & \frac{\overline{a}_K^2}{\|\overline{\mathbf{a}}\|} \end{pmatrix} \text{ and } \mathbf{D} = \begin{pmatrix} \frac{\Omega_1}{\|\overline{\mathbf{a}}\|} & 0 & \dots & 0 \\ 0 & \frac{\Omega_2}{\|\overline{\mathbf{a}}\|} & \dots & 0 \\ \vdots & \dots & \dots & \vdots \\ 0 & 0 & \dots & \frac{\Omega_M}{\|\overline{\mathbf{a}}\|} \end{pmatrix}$$

where $\|\overline{a}\| = \sum_s (\sum_i a_{is})^2$ is the length of resultant vector. The weights within each matrix are sum to 1 such that $\sum_s^K \widetilde{\mathbf{D}}_{ss} = 1$ and $\sum_i^M \mathbf{D}_{ii} = 1$. The two weights measure the relative importance of each product or attribute. The intuition of the demand system is as following. Under the unbundling assumption, the consumers value product as a bundle of attributes so they fundamentally care about the price of each attributes. However, the bundle packing itself is also valuable. Hence, each product get a discount on price which is negatively correlated with its centrality – a rare bundle packing gets more credit from the consumer. The equation 52 results in the following expenditure cost function:

$$\begin{aligned} \ln(e) &= a_0 + \ln \mathbf{p}^T \boldsymbol{\alpha} + \frac{1}{2} \ln \mathbf{p}^T \Theta \ln \mathbf{p} \\ &= a_0 + \sum_{i=1}^{N} \alpha_i \ln p_i + \frac{1}{2} \sum_{i=1} \sum_{j=1} \theta_{ij} \ln p_i \ln p_j \end{aligned} \tag{53}$$

where θ_{ij} are given by:

$$\theta_{ij} = \begin{cases} \frac{\theta}{\|\overline{\mathbf{a}}\|} \omega_{ij} & j \neq i \\ -\frac{\theta}{\|\overline{\mathbf{a}}\|} (\Omega_i - 1) & j = i \end{cases} \tag{54}$$

This function satisfies the two restrictions on θ_{ij}: (1) *symmetric restrictions* $\theta_{ij} = \theta_{ji}$; (3) *homotheticity restrictions*: $\sum_j \theta_{ij} = 0$ for any i and j. Further Imposing the homotheticity restriction on $\boldsymbol{\alpha}$ such that $\sum_i \alpha_i = 1$ and $\alpha_i \in (0, 1)$, equation 53 is exact translog expenditure cost function proposed by Diewert (1976) and Christensen *et al.* (1975).

Additional Derivations for θ_{ij}:

$$\begin{aligned} \ln \mathbf{p}^T \Theta \ln \mathbf{p} &= \theta \ln \widetilde{\mathbf{p}}^T \widetilde{\mathbf{D}} \ln \widetilde{\mathbf{p}} - \theta \ln \mathbf{p}^T \mathbf{D} \ln \mathbf{p} \\ &= \theta \left(diag(\overline{\mathbf{a}})^{-1} \mathbf{A} \ln \mathbf{p} \right)^T \widetilde{\mathbf{D}} \left(diag(\overline{\mathbf{a}})^{-1} \mathbf{A} \ln \mathbf{p} \right) - \theta \ln \mathbf{p}^T \mathbf{D} \ln \mathbf{p} \\ &= \ln \mathbf{p}^T \left(\frac{\theta}{\|\overline{a}\|} \mathbf{A}^T \mathbf{A} \right) \ln \mathbf{p} - \ln \mathbf{p}^T \left(\frac{\theta}{\|\overline{a}\|} diag(\Omega_i) \right) \ln \mathbf{p} \\ &= \ln \mathbf{p}^T \left[\frac{\theta}{\|\overline{a}\|} \left(\mathbf{A}^T \mathbf{A} - diag(\Omega_i) \right) \right] \ln \mathbf{p} \end{aligned} \tag{55}$$

Appendix A.3. Comparison with other translog functional forms

This section compares the translog functional form defined in equation 53 with other popular functional forms discussed in rencet trade literature – specifically, Feenstra and Weinstein (2017) and Diewert (2023).

First, Feenstra and Weinstein (2017) consider a special case where all products are symmetric and all of the second order parameters are set to be θ/M. These assumptions yield a Symmetric Translog (ST) Function and give neat analytical solutions. The ST function is a special case in our model with zero variance in product network. Therefore, $\omega_{ij} = 1$ for any i and j and $\|\overline{a}\| = M$. We can solve θ_{ij} as

$$\theta_{ij} = \begin{cases} \frac{\theta}{M} & j \neq i \\ \theta(\frac{1-M}{M}) & j = i \end{cases} \tag{56}$$

This is exactly the same with equation (3) in Feenstra and Weinstein (2017).

Second, Diewert (2023) propose a Generalized Symmetric Translog (GST) Function and allow for product specific second order parameters. It corresponds to a special case of my model where consumer only care product differentiation along one dimension. Let consumer gives all weights on attribute 1 so that weight matrix $\widetilde{\mathbf{D}}$ is given by

$$\widetilde{\mathbf{D}} = \begin{pmatrix} \frac{\sum_{ij} a_{i1}a_{j1}}{\sum_{ij} a_{i1}a_{j1}} & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \dots & \dots & \vdots \\ 0 & 0 & \dots & 0 \end{pmatrix} \text{ and } \mathbf{D} = \begin{pmatrix} \frac{a_{11}(\sum_i a_{i1})}{\sum_{ij} a_{i1}a_{j1}} & 0 & \dots & 0 \\ 0 & \frac{a_{21}(\sum_i a_{i1})}{\sum_{ij} a_{i1}a_{j1}} & \dots & 0 \\ \vdots & \dots & \dots & \vdots \\ 0 & 0 & \dots & \frac{a_{M1}(\sum_i a_{i1})}{\sum_{ij} a_{i1}a_{j1}} \end{pmatrix}$$

We can solve θ_{ij} as

$$\theta_{ij} = \begin{cases} \frac{\theta}{\sum_{ij} a_{i1}a_{j1}} \, a_{i1}a_{j1} & j \neq i \\ \frac{\theta}{\sum_{ij} a_{i1}a_{j1}} \, \left(a_{i1}^2 - a_{i1}\sum_{j'} a_{j'1}\right) & j = i \end{cases} \tag{57}$$

where $\widetilde{\theta}$. This is exactly the same with equation (11) in Diewert (2023). Notice that both of functional forms satisfy the discussion of separable translog in Example 3a from Matsuyama and Ushchev (2017) and therefore yield single price aggregator.

Appendix A.4. Derivations for market shares

The Langrangian function is

$$\min_{q_i} \left(e - \lambda(\sum_{i=1}^{N} p_i q_i - E) \right) \tag{58}$$

F.O.C. is

$$\begin{aligned} \frac{e}{p_i}\frac{\partial \ln e}{\partial \ln p_i} &= \lambda q_i \\ \frac{e}{p_i}(\alpha_i + \sum_j \theta_{ij} \ln p_j) &= \lambda q_i \end{aligned} \tag{59}$$

Summing equation 59 over i, we can derive:

$$\lambda = \frac{e}{E} \tag{60}$$

Therefore, the real utility $U = 1/\lambda$. Taking equation 60 into 59, we have

$$s_i = \frac{p_i q_i}{E} = \alpha_i + \sum_j \theta_{ij} \ln p_j \tag{61}$$

Appendix A.5. Derivations for the elasticities

$$\eta_{ij} \equiv \left.\frac{\partial \ln q_i}{\partial \ln p_j}\right|_{U\ const.} \tag{62}$$

From equation 59, we can derive q_i is

$$q_i = \frac{e s_i}{p_i} \cdot U \tag{63}$$

We can get

$$\begin{aligned} \eta_{ij} &= \frac{\partial \ln\left(\frac{e s_i}{p_i} \cdot U\right)}{\partial \ln p_j} \\ &= \frac{\partial \ln e}{\partial p_j} + \frac{\partial s_i}{\partial p_j}\frac{1}{s_i} - \mathbb{1}_{i=j} \\ &= s_j + \frac{\theta_{ij}}{s_i} - \mathbb{1}_{i=j} \end{aligned} \tag{64}$$

Also, we can derive the Allen-Uzawa partial elasticities of substitution:

$$\begin{aligned} \sigma_{ij} &= \frac{1}{s_j}\frac{\partial \ln q_i}{\partial \ln p_j} \\ &= 1 + \frac{\theta_{ij}}{s_i s_j} - \mathbb{1}_{i=j}\frac{1}{s_i} \end{aligned} \tag{65}$$

The unconditional price elasticities and elasticities of substitution can be derived in similar way:

$$\eta_{ij} \equiv \frac{\partial \ln q_i}{\partial \ln p_j} = \begin{cases} \frac{\theta_{ij}}{s_i} & j \neq i \\ -1 + \frac{\theta_{ii}}{s_i} & j = i \end{cases} \tag{66}$$

and

$$\begin{aligned}\sigma_{ij} &\equiv \frac{\eta_{ij}}{s_j} \\ &= \begin{cases} \frac{\theta_{ij}}{s_i s_j} & j \neq i \\ -\frac{1}{s_i} + \frac{\theta_{ii}}{s_i^2} & j = i \end{cases}\end{aligned} \tag{67}$$

Appendix A.6. Derivations for the choke prices

TBD

Appendix A.7. Derivations for markup with nested-CES

TBD

Appendix A.8. Uniqueness and Existence

TBD

Appendix B. Additional Empirical Results

Table A1: Top ten most banned apps worldwide

	App	Category	Users (Brillion)	Banned in
1	Facebook	social media	3.0	CN, RU, IR, KP, MM, TM, UG
2	WhatsApp	messaging	3.0	CN, KP, SY, QA, AE
3	Youtube	video	2.70	CN, KP, TH, MA, DE(partially)
4	Intagram	social media	2.0	CN, VN, KP, TK
5	Google	search engine	1.5	CN
6	WeChat	social media	1.3	IN, UK, US/CA/NZ (partially)
7	TikTok	social media	1.0	IN, US/CA/NZ/TW/ID (partially)
8	X/Twitter	social media	0.4	CN, RU, BR, IR, KP, MM, TM
9	Baidu	search engine	0.7	IN
10	Spotify	music	0.6	CN, IR, ET, LY, IQ, CU, MM

Notes: This table uses data from Linkee.ai and public information.

Table A2: Top 10 public corporation by market capitalization in 2024

Rank	Company	Market Cap (trillion USD)	Main product
1	Microsoft	3.32	Software
2	Apple	3.23	Electronics, software
3	Nvidia	3.18	Semiconductors
4	Alphabet	2.27	Software
5	Amazon	2.01	Software
6	Meta	1.28	Software
7	TSMC	0.90	Semiconductors
8	Berkshire Hathaway	0.88	Conglomerate
9	Eli Lilly	0.82	Pharmaceutical
10	Broadcom	0.75	Semiconductors, software

Notes: The rank and market cap are up to date as of 30 June 2024. See data resources in https://en.wikipedia.org/wiki/List_of_public_corporations_by_market_capitalization.

Table A3: Event Study Coefficients

	Banned apps	Competing apps ($\times 10^{-2}$)			
		all	20^{th} percentile	10^{th} percentile	5^{th} percentile
	(1)	(2)	(3)	(4)	(5)
$\tau =$ Dec-19	0.032	-0.001	-0.021	-0.040	-0.058
	(0.19)	(0.02)	(0.08)	(0.12)	(0.23)
$\tau =$ Jan-20	0.078	-0.006	-0.033	-0.064	-0.062
	(0.19)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Feb-20	0.100	-0.008	-0.015	-0.025	0.023
	(0.19)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Mar-20	0.101	-0.008	0.068	0.101	0.243
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Apr-20	0.063	-0.008	0.119	0.148	0.229
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ May-20	-0.093	-0.008	0.112	0.098	0.151
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Jun-20	-0.208	-0.004	0.109	0.108	0.247
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Jul-20	-0.543***	0.011	0.150*	0.203*	0.430*
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Aug-20	-0.855***	0.039**	0.265***	0.402***	0.726***
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Sep-20	-0.853***	0.038*	0.286***	0.399***	0.697***
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Oct-20	-0.848***	0.039**	0.318***	0.446***	0.793***
	(0.18)	(0.02)	(0.08)	(0.12)	(0.22)
$\tau =$ Nov-20	-0.848***	0.043**	0.358***	0.536***	1.014***
	(0.18)	(0.02)	(0.08)	(0.12)	(0.23)
$\tau =$ Dec-20	-0.849***	0.043**	0.359***	0.547***	1.137***
	(0.18)	(0.02)	(0.08)	(0.12)	(0.23)
Observations	406	163,377	30,241	15,901	6,258
R^2	0.64	0.99	1.00	1.00	1.00
Fixed Effects	app	app	app	app	app
Number of apps	29	12,264	2,404	1,270	497

Figure A2: Robustness Check on US and RoW

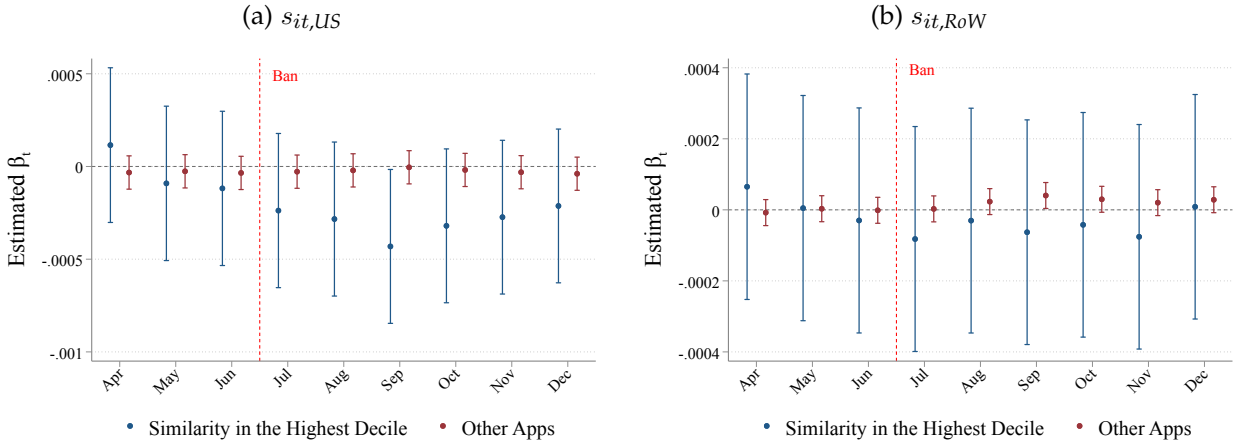

Notes: This figure corresponds to robustness checks for Figure 2.

Figure A3: Heterogeneity by similarity quintiles

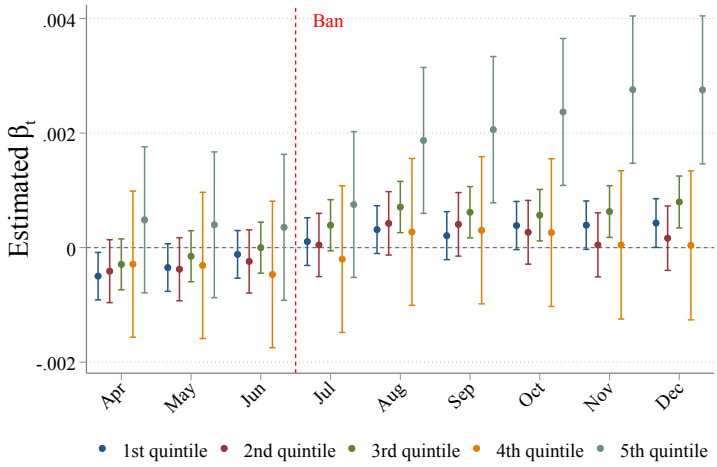

Notes: This figure corresponds to robustness checks for Figure 2.

Appendix C. Estimation Procedure with LLMs

In this paper, I use the large language model for two tasks.

Appendix C.1. Generating synthetic data

The raw data are the app description for consumers in iOS App Store. I use a prompt chat model to generate the synthetic data from raw product descriptions. The pre-trained model is the most efficient GPT chat model up to date. I use one-shot prompt to focus on the functionality and services. Here is the prompt structure:

system: You act as an iOS app analyst, specializing in analyzing market competition among apps. When provided with app descriptions, you identify core features, target audiences, and unique selling points to assess how closely apps compete within the same market segment. You maintain a professional, clear, and concise tone and always give the response in English

user: Here is a description of an iOS mobile app delimited by triple backticks. Please summarize the app in the following format: -A Overall summary of the app,: 'Overall summary: xxxx' -A list of main functions/services in bullet points: '1. xxxx 2.xxx...' Ensure your response is in English, focusing on the most important functions and omitting subscription options, privacy policies, and terms of service. App description: "'{raw_description}"'

model: `gpt-4o-mini`

temperature: 0

seed:123

Following the given prompt, the synthetic description for Netflix is given by

Netflix is a mobile app that provides access to a vast library of award-winning TV shows, movies, documentaries, and stand-up specials, allowing users to stream content anytime and anywhere.

Main functions/services:

1. Browse and stream a wide variety of TV shows and movies.
2. Personalized recommendations based on viewing habits.
3. Create up to five individual profiles for different users.
4. Family-friendly content with a dedicated kids' section.
5. Quick video previews and notifications for new episodes and releases.
6. Download titles for offline viewing.

Appendix C.2. Generating embeddings

The embedding model is `text-embedding-3-large`

Embedding Model = text-embedding-3-large

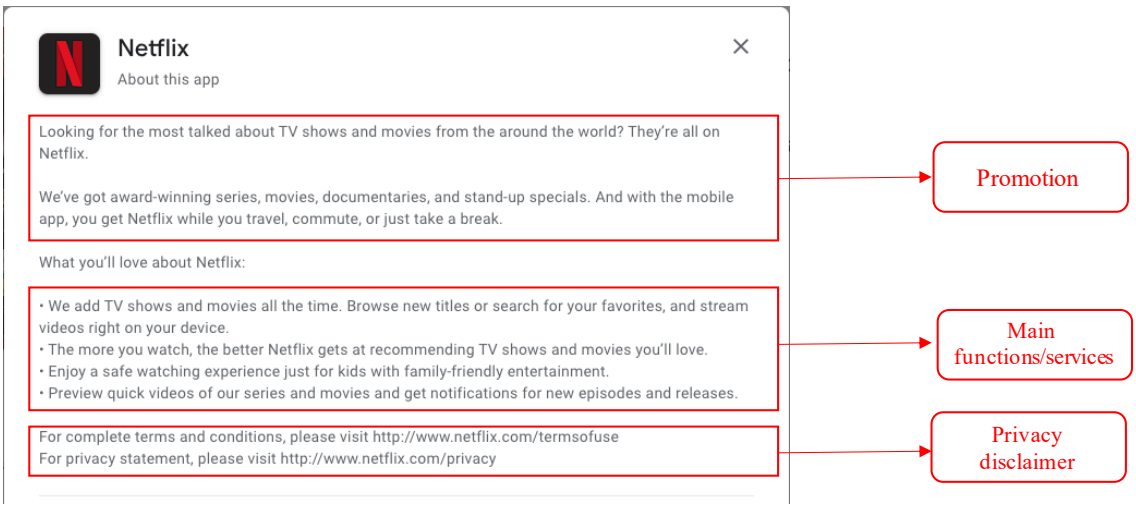

Figure A4: Example of App Description for Netflix

Appendix C.3. Cosine similarities and model validations

We can actually make finer connections between apps, which we will anyways need to construct our instrument. We select the largest App Store category ("Communications') and select the 50 apps in Communications with the most users over our sample. We read each app description and manually classify it into one of 9 sub-categories, displayed in table A4. For each pair of apps a and a' we compute the cosine similarity $\rho_{aa'}$. Because cosine similarities are symmetric and $\rho_{aa} = 1$, we only report elements above the diagonal. There are $50 \cdot (50-1)/2 = 1{,}225$ pairs. We place the $\rho_{aa'}$ into deciles and report the deciles in table A4.

Shaded cells are app pairs that are in the same category and we expect these to have $\rho_{aa'}$ in the 10th (top) decile. False negatives are shaded cells below the 10th decile. Only 22% of pairs are false negatives and all of these are marginally false in that they are either in the 8th or 9th deciles. Unshaded cells are app pairs where a and a' are in different categories and we expect their cosine similarities to be below the 10th decile. False positives are unshaded cells in the 10th decile. For clarity, we only report the false positives. Only 5% of app pairs are false positives. Overall, only 58 of 1,225 pairs are false negatives or positives. This suggests that BERT is able to finely distinguish the similarity of apps. What most surprises us is that in all but 3 of 58 cases, a rereading of

Table A4: Validation of $\rho_{aa'}$ for Communications Apps

	Call-Video-Message													Social network					Email				Browser				Drive			File share			Anti-virus			Emoji			VPN	
	WhatsApp	Messenger	Blackberry	QQ	WeChat	LINE	Skype	imo	Viber	Telegram	KakaoTalk	Zalo	Google Duo	Facebook	Weibo	Baidu Tieba	OK.ru: Mess.	VK social	Gmail	MS Outlook	Yahoo Mail	Mail.ru	Google	Chrome	UC Browser	Opera	Goog. Drive	Googl.docs	Goog. Sheets	SHAREit	Xender	4shared	AVG	App locker	Avast	Hike sticker	Bitmoji	Sticker.ly	Master VPN	Speedtest
WhatsApp Messenger	10	10	10	10	10	10	10	10	10	10	10	9	10	10	·	·	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Messenger		10	10	10	10	10	10	10	10	10	10	10	10	10	·	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	10	·	·	·
Blackberry Message			10	10	10	10	9	10	10	10	9	9		10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
QQ				10	10	10	10	10	10	10	10	9		10	·	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
WeChat					10	10	10	10	10	10	10	9		10	10	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
LINE						10	10	10	10	10	10	9		10	·	·	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	10	·	·
Skype							10	10	9	10	9	10		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
imo video calls								10	10	10	10	10		·	·	·	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Viber Messenger									10	10	9	10		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Telegram										10	9	8		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
KakaoTalk											10	9		·	·	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Zalo												8		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Google Duo														·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Facebook														10	10	10	10		·	·	·	·	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Weibo															10	9	9		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Baidu Tieba																8	8		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
OK.ru: Messenger																	10		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
VK social network																			·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Gmail																			10	10	10		10	10	·	·	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·
Microsoft Outlook																				10	9		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Yahoo Mail																					10		·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Mail.ru																							·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Google																							10	10	9	·	·	10	10	·	·	·	·	·	·	·	·	·	·	·	·
Google Chrome																								10	10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
UC Browser																									10	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Opera Browser																										·	·	·	·	·	·	·	·	·	·	·	·	·	·	·
Google Drive																											10	9	·	·	·	·	·	·	·	·	·	·	·	·
Google docs																												10	·	·	·	·	·	·	·	·	·	·	·	·
Google Sheets																													·	·	·	·	·	·	·	·	·	·	·	·
SHAREit-Connect																														10	10	·	·	·	·	·	·	·	·	·
Xender File Share																															8	·	·	·	·	·	·	·	·	·
4shared Mobile																																·	·	·	·	·	·	·	·	·
AVG																																	10	10	·	·	·	·	·	·
App locker																																		9	·	·	·	·	·	·
Avast																																			·	·	·	·	·	·
Hike sticker																																				9	8	·	·	·
Bitmoji																																					10	·	·	·
Sticker.ly																																						·	·	·
Master VPN																																							9	
Speedtest by Ookla																																								

Notes: Table **??** in the main text did not include the now-defunct Google Duo app. The app reflects an older technology that did internet phone calls without video or messaging. It therefore has a somewhat lower correlation (8s and 9s) with other 'Call-Video-Message' apps.

app descriptions reveals that BERT is more reliable than we were.[23] [24]

We have 50 apps, but table A4 has only 40. Each of the remaining apps has unique functions and so belongs in its own sub-category. We therefore expect none to be in the 10th decile. Of the 445 cells involving at least one of these apps, only 7 are in the 10th decile and, after rereading app descriptions, only 2 are BERT errors.[25] Summarizing, BERT does a great job of describing the relationship between all pairs of apps and does so with considerable nuance.

Appendix D. Firewall Test on Apps

TBD

[23] *False Positives:* We missed two patterns in the false positives. (*i*) Most of the false positives are in the off-diagonal block with one app from Call-Video-Message and the other app from Social Media. These are cases where the two apps share many common functions. For example, Facebook allows for video and text messaging so that the many 10s in the Facebook column are sensible. Likewise, in the app description of OK.ru, the first and second sentences deal with social networking and messenger services, respectively, leading to many sensible 10s in the OK.ru column. (*ii*) Google apps often work in tandem and their descriptions often reference this. As a result, most of the remaining 10s involve pairs of Google apps. This leaves only 2 of 21 false postivies that are an error by BERT.

[24] *False Negatives:* (*i*) Of the 26 false negatives, 18 involve a messaging, social networking or file sharing app produced in an autocracy and/or having end-to-end encryption. These app descriptions have dog whistles that we missed but that BERT correctly identified and classified. (*ii*) 6 of the 26 false negatives involve Google Duo, a now-defunct app that did internet phone calls without video or messaging. BERT correctly recognized that Duo was thus not sufficiently similar to most Call-Video-Message apps. (*iii*) 2 of 26 false negatives involve Hike Sticker, which we incorrectly categorized based on its name (a sticker is an Emoji), and BERT caught our error. This leaves only 1 of 26 false negatives that is an error by BERT.

[25] 5 of 7 are Google products. See point (*ii*) of footnote 23.