

NBER WORKING PAPER SERIES

AI, TRADE AND CREATIVE DESTRUCTION:
A FIRST LOOK

Daniel Trefler
Ruiqi Sun

Working Paper 29980
<http://www.nber.org/papers/w29980>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

This paper is part of the larger book project "Robots and AI: A New Economic Era" edited by Lili Yan Ing and Gene Grossman (Routledge). Gene Grossman and an anonymous referee provided excellent feedback. We thank SensorTower for their help and encouragement in facilitating this project. Keith Head graciously helped interpret the PPML estimates and Mert Demirer answered tough AI questions. We benefited from early conversations with Avi Goldfarb. Shurui Liu provided research assistance. This project was generously supported by the Social Sciences and Humanities Research Council of Canada (SSHRC Grant #435-2021-0149) and the Economic Research Institute for ASEAN and East Asia (ERIA). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Daniel Trefler and Ruiqi Sun. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

AI, Trade and Creative Destruction: A First Look
Daniel Trefler and Ruiqi Sun
NBER Working Paper No. 29980
April 2022
JEL No. F1,F12,F14

ABSTRACT

Artificial Intelligence is a powerful new technology that will likely have large impacts on the size, direction and composition of international trade flows. Yet almost nothing is known empirically about this. One AI-enabled set of services that can be tracked resides in the palm of our hands: the Mobile Apps used by half the world's population. To analyze the impact of AI on international trade in mobile App services we merge 2014-2020 data on international downloads of mobile Apps with data on the AI patents held by each App's parent company. From this we build a measure of AI deployment. We instrument AI deployment using cost-shifters from the theory of comparative advantage: Countries with a large stock of AI expertise will have a comparative advantage producing AI-intensive Apps. We show the following IV results. (1) Bilateral Trade: AI deployment increases App downloads by a factor of six. (2) Variety Effects: AI deployment doubles the number of exported App varieties. (3) Creative Destruction: AI deployment increases creative destruction (entry and exit of Apps) and in 2020 the net effect was an increase in welfare of between 2.5% and 10.6%.

Daniel Trefler
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
dtrefler@rotman.utoronto.ca

Ruiqi Sun
PBC School of Finance
Tsinghua University
43 Chengfu Rd, Haidian District
Beijing 100083
sunrq.17@pbcfsf.tsinghua.edu.cn

The transformative potential of artificial intelligence (AI) is apparent from our daily use of smartphones. We log in using AI-enabled facial recognition, issue commands with AI-enabled speech recognition, conduct AI-enabled internet searches, buy from stores pushing AI-enabled recommendations, and receive goods shipped with AI-enabled logistics systems. Not only has AI enabled the creation of new services, it has improved on existing services and disrupted older services in a familiar process of creative destruction (Schumpeter, 1942). All of these changes can be seen in the palm of our hand and are meticulously tracked by corporations. This paper uses big data on international App downloads and AI patents to track how AI is changing the pattern of trade in services, the variety of services available in each country, and the process of creative destruction.

The early hype about AI has given way to more sober analysis showing that to date AI has had limited effects on tasks (Brynjolfsson, Mitchell, and Rock, 2018), employment and wages (Acemoglu, Autor, Hazell, and Restrepo, 2020). Less is known about AI's impact on international trade either theoretically or empirically. In this paper we explore that impact on (a) bilateral trade flows, (b) the variety of goods imported, and (c) the creation and destruction of varieties. The impact of AI on trade flows is of great interest, but ultimately we care about welfare. We thus also calculate the welfare effects due to AI-induced changes in the availability of varieties to consumers.

There is good reason to expect all three of the above impacts. (a) For bilateral trade flows, McKinsey Global Institute (2019) predicts that AI will reduce outsourced business process and IT services. It will also reduce goods trade by facilitating additive manufacturing that moves production to the point of consumption. McKinsey predicts that together these developments will reduce trade by a trillion dollars. Of course, this reduction in trade tells us nothing about AI's impact on welfare. Indeed, in McKinsey's scenario trade volumes and welfare likely move in opposite directions. (b) For product variety, AI leads both to new services (horizontal differentiation) and to improvements on existing services (vertical differentiation). These are known to affect the pattern of trade and the welfare gains from trade, usually in positive ways. See Krugman (1979), Helpman (1981), Feenstra (1994, 2010), Melitz (2003), Broda and Weinstein (2006), and Hsieh, Li, Ossa, and Yang (2020) for analysis of horizontal differentiation. (c) For creative destruction, AI's impact on vertical differentiation disrupts and displaces existing services. On this process of creative destruction through endogenous innovation see Aghion and Howitt (1992) and Akcigit and Kerr (2018) for closed-economy models and Grossman and Helpman (1991a,b) for both closed- and open-economy models.

Despite intense public interest in AI, research on the impacts of AI on trade, product variety and creative destruction is almost nonexistent. Goldfarb and Trefler (2019a) review the theoretical issues for international trade raised by AI. They argue that key features of AI are scale, local knowledge diffusion, and the degree of international knowl-

edge diffusion. Scale and local knowledge diffusion/externalities have implications for trade flows that have long been understood in the economic geography literature. As well, the degree of local versus international diffusion is central to the endogenous growth literature e.g., Rivera-Batiz and Romer (1991), Grossman and Helpman (1991b) and Irwin and Klenow (1994). Goldfarb and Trefler (2019a,b) also argue that AI affects trade costs in complex ways. For example, privacy concerns create additional trade costs not usually considered by international trade economists. Further, interstate competition can create national regulatory responses best characterized as a privacy race to the bottom. Royal Society-National Academy of Sciences (2019) summarizes the proceedings of a Washington D.C. symposium on international harmonization of AI regulations, including a summary of Goldfarb's and Trefler's views.

The only empirical paper directly on AI and trade is by Brynjolfsson, Hui, and Liu (2019). They show that eBay's introduction of a machine translation system increased its exports by 17.5%. This is the opposite of McKinsey Global Institute's (2019) speculations. Our work is closely related to Brynjolfsson *et al.*. The advantage of their approach is that it carefully identifies the exact AI (machine translation) and the exact mechanism for eBay. In contrast, we will work with a wide set of AIs, companies, and services. This allows us to employ the standard gravity equation for examining impacts on trade as well as product variety and creative destruction.

There are other more distantly related papers. Beraja, Yang, and Yuchtman (2020) show how Chinese government security contracts for facial recognition software provided confidential security data to Chinese firms, data that improved these firms' products. By implication, the paper shows how government subsidies in the AI sphere can improve competitiveness. More tangential to our interests here, Bailey, Gupta, Hillenbrand, Kuchler, Richmond, and Stroebel (2020) use Facebook data to construct bilateral social connections between countries and show that these are a more powerful determinant of bilateral trade flows in *goods* than are traditional determinants such as distance and borders. Though tangential to our main results, we include their bilateral social connections measure and find that it impacts App-based service trade as well.¹

This review, even if missing some citations from the rapidly growing AI literature, clearly demonstrates that the literature on AI and trade is very small. This is in part because trade in AI-enabled services is hard to document. At the core of this paper is the observation that there is actually a vast amount of data available.

Motivated by the tremendous amount of AI that underlies our smartphone Apps, this

¹AI is part of a larger process of automation and is thus part of a larger literature on the impact of trade and technology on employment, wages and inequality. A recent contribution to this literature with an international dimension is Stapleton and Webb (2020) who consider the impact of robots on Spanish multinationals during 1990–2016.

paper is about international trade in mobile App services as well as its implication for product variety and creative destruction. The core of our analysis is based on two types of data. The first is data from a private data provider (SensorTower) on the number of mobile App downloads by App, by producer country, and by user country for the period 2014–2020. The second data source is Bureau van Dijk’s Orbis Intellectual Property patent database. We adopt the methodology behind the WIPO PATENTSCOPE Artificial Intelligence Index to determine whether or not a patent in Orbis is an AI patent. We review this complex methodology in section 2.3 below. A difficult part of building our database is merging the App and patent databases. Each App in the SensorTower data is identified with an ultimate owner. For example, Alphabet owns Google Chrome, Nest Home, YouTube, Waze and Fitbit. We then match ultimate owners with those in Orbis. We do the match by hand for the 834 ultimate owners with the most downloads globally. We show below that the Apps and ultimate owners excluded from our analysis are mostly small and obscure.

We use information about each ultimate owner’s Apps, AI patents, and assets to develop a measure of ‘App-deployment’ by year, exporter, and App category. App categories are defined as follows. The Apple App Store places Apps into 19 App groups (e.g., social networking, productivity). We further refine each group by 19 2-digit NACE industries (e.g., mining, finance). We refer to this cross of groups \times industries as ‘App categories.’ There are 292 categories. Aggregating up from Apps and ultimate owners, we compute AI patent counts by category \times exporter \times year bins. This is our novel measure of AI deployment by category \times exporter \times year bins. (We scale this measure by the value of assets held by firms in the bin; however, our results are not sensitive to this scaling.)

We can summarize our database handily by comparing it to COMTRADE, the standard international trade database used for gravity estimation. We have 53 exporters, 84 importers, seven years (2014–2020), and 292 App categories (App categories are like HS2 or HS4 codes in COMTRADE). Further, many studies of creative destruction and changes in the number of traded varieties (e.g., Broda and Weinstein, 2006) define varieties as US HS10 product lines. There are roughly 20,000 HS10 codes/varieties. In contrast, we have 82,850 Apps/varieties.

Our main results flow from regressions of various outcomes on our AI deployment measure. An obvious concern is the endogeneity of AI deployment. We therefore need an instrument that captures exogenous shocks to the cost of deployment. Heckscher-Ohlin theory provides one. A country with deep AI expertise will have cheap and ready access to the inputs used in deploying AI, which in turn provides a cost advantage that is especially pronounced in App categories that use these inputs intensively. We therefore instrument App deployment with the interaction of (1) a country’s *AI expertise*

as measured by its AI research output and (2) an App category's *AI intensity*. This serves as an exogenous shifter of the costs of AI deployment.²

We have three main IV findings. All of them exploit within-App-category variation.

1. *Bilateral Trade*: We estimate a gravity model of App downloads whose dimensions are importer-exporter dyads, App categories, and years. Using IV, we find that AI deployment causes a sixfold increase in App downloads.

Beyond the Brynjolfsson *et al.* study of eBay's use of machine translation, this is the first and most systematic evidence of the impact of AI on trade.

2. *Varieties*: AI deployment doubles the number of bilaterally traded Apps/varieties.
3. *Creative Destruction and Welfare*:
 - (a) *Entry and Exit*: AI deployment causes high levels of entry into and exit out of the Apps/varieties available in the importer country. That is, it causes creative destruction.
 - (b) *Welfare*: We calculate the welfare implications of entry and exit using Feenstra's (1994, 2010) technique. We find that in 2020, welfare from Apps was between 2.5% and 10.6% higher than it would have been under the counterfactual of no AI deployment. Both are large numbers and the range depends on whether the elasticity of substitution between Apps is high (5) or low (2). An important caveat is that in the Feenstra formula we use download shares rather than expenditure shares.

These three results demonstrate that AI deployment in the mobile App space has already had tangible effects on trade, product variety, creative destruction and welfare.

One might wonder whether our conclusions are the result of a spurious correlation between AI patenting and other unobservables. To examine this, we consider non-AI patents and find that their effects are modest and their inclusion in the analysis does not affect our results.

Section 1 provides background on mobile Apps and AI. Section 2 describes the database. Section 3 uses bilateral gravity equations to estimate the impact of AI deployment on trade. Section 4 estimates the impact of AI deployment on the extensive margin, that is, on the number of Apps/varieties. Section 5 examines the impact of AI deployment on entry, exit, creative destruction and welfare.

²More specifically, AI intensity is measured as total global AI patents associated with the App category. AI expertise is measured as the number of papers presented at AI conferences by researchers affiliated with the country's universities and other research institutions. Data on AI expertise are from Zhang, Mishra, Brynjolfsson, Etchemendy, Ganguli, Grosz, Lyons, Manyika, Niebles, Sellitto *et al.* (2021).

1. A Brief Overview of Apps and AI

As wireless internet technology and personal portable devices have come down in cost and risen in accessibility, mobile applications have become a fixture of daily life. In 2020, the number of mobile internet users hit 4.3 billion globally or 92% of all internet users.³ Each and every day we use mobile applications to read our mail, browse the internet, post to our social network, shop, bank, take photos, play games, watch videos and more. The mobile application industry has been fast-growing and will continue to expand at a significant pace. It currently generates upwards of \$700 billion in revenues and is growing rapidly.⁴

The biggest two application marketplaces, App Store (for iOS) and Google Play (for Android), launched in 2008 alongside the release of the first smartphones (iPhone 3G and T-mobile G1). At the time, these two application marketplaces had about 500 Apps. Today the App Store has 1.82 billion Apps and Google Play has 2.8 billion Apps.

Turning from Apps to AI, Agrawal, Gans, and Goldfarb (2018) define AI as a collection of complementary technologies involving algorithms, data, and computing power that allow predictive programs to automatically improve their performance through experience. The authors date the commercial introduction of AI to 2012. Since then, some of the companies that have pushed the frontiers of AI have grown to be among the biggest in the world. Table 1 lists the eight largest companies in the world by market capitalization. Column 2 is 2020 market capitalization in millions USD. Every one of these companies uses AI to improve its services and expand its service offerings. One, albeit limited, indication of this is the number of AI patents held by these companies. These companies have a large number of such patents. (We do not have data for Tesla, which is not in our dataset.)

Two things stand out in the table. For one, with the exception of Apple and Microsoft, these companies had relatively little presence in the 2011 list of the largest companies in the world. Indeed, Facebook, Tesla and Alibaba were not even in the top-500. This illustrates just how dynamic these companies are and, by implication, how dynamic are the effects of AI likely to be. For another, all of these companies are based either in the United States or China. This has led Kai-Fu Lee (2018), former CEO of Google China, to argue that in the future these two countries will produce all AI-enabled services and the rest of the world will be stuck paying hefty royalties. This potentially has dramatic implications for the pattern of international service trade flows.

³<https://www.statista.com/statistics/617136/digital-population-worldwide>

⁴<https://www.statista.com/statistics/269025/worldwide-mobile-app-revenue-forecast>.

Table 1: The World’s Largest Companies: AI, Growth, Location and Internationalization

Company	Market Cap			Total Downloads		
	(\$B)	AI Patents	2011 Rank	Nationality	Worldwide (millions)	Foreign Share
(1)	(2)	(3)	(4)	(5)	(6)	(7)
1. Apple	\$2,254	1,071	3	USA	151	76%
2. Microsoft	\$1,682	7,088	10	USA	4,023	81%
3. Amazon	\$1,634	509	77	USA	3,015	69%
4. Alphabet	\$1,185	5,675	28	USA	16,155	81%
5. Facebook	\$777	1,243	< 500	USA	21,913	91%
6. Tencent	\$683	2,930	178	China	7,160	22%
7. Tesla	\$668	-	< 500	USA	-	-
8. Alibaba	\$629	1,767	< 500	China	5,065	52%

Notes: Data for 2011 and 2020 are as of December 31. See https://en.wikipedia.org/wiki/List_of_public_corporations_by_market_capitalization#2020. Market capitalization is in millions USD. 2011 ranks are from the *Financial Times* FT500 as of March 31, 2011 (<http://media.ft.com/cms/33558890-98d4-11e0-bd66-00144feab49a.pdf>). AI patents are computed by the authors as described below. ‘< 500’ means the company is not on the list. Data on Tesla’s AI patents are not part of our database, but the company is at the frontier of AI algorithms for autonomous vehicles. Google Play is not available in China and so Android App downloads in China are imputed in this table. We estimate China’s total downloads as China’s iOS downloads divided by the market share of Apple devices in China (21.8% in 2020).

Table 1 makes two other points about these companies. Column 6 shows that these firms all have heavily downloaded Apps, an average of 8 billion per firm. Column 7 shows that these Apps are heavily downloaded internationally. On average, 61% of these firms’ downloads are done outside of the firms’ home countries. This fact is not unique to our top-tier companies: The median value of foreign download shares is 64% in our sample of 834 firms. This is quite remarkable compared to the goods economy where all but a few of the largest multinationals earn most of their revenue in their home markets. Thus, App services are much more internationalized than say manufacturing. Interestingly, the Chinese companies in 1 table are much less internationalized than their US counterparts.

2. The Data

2.1. Mobile Application Data

Our primary database is the App download data purchased from SensorTower. SensorTower is the largest and most reliable company providing App-level metadata. The data track App-level downloads by user country from 2014 to 2020 for the Apps available in the Apple App Store and Google Play, which are the biggest application marketplaces

for the iOS and Android operating systems.⁵ Each App in the Apple App Store and Google Play has a unique, time-invariant product ID and is accompanied by the name of the developer, the name of the selling publisher (App terminology for ‘firm’), and the selling publisher’s website. SensorTower consolidates App IDs to deal with the fact that an App may have different IDs in different countries e.g., TikTok in the US and Douyin in China. We use consolidated IDs to avoid overstating the number of Apps. SensorTower also creates a ‘unified’ firm name that keeps track of the fact that publishers often have different names in different countries and sometimes have different names across wholly owned subsidiaries. We use the unified firm name to link with patent and financial data.

The App Store and Google Play place Apps into groups.⁶ These are displayed in table 2 along with the top-3 Apps in the group. For each App, the table also shows the company and its headquarters country. Most of the top Apps are owned by large digital platforms located in the US and China.

One obvious issue with groups is that they are not fine enough to be useful. For example, the ‘Utilities’ group includes Google’s Chrome and Toyota’s DV, an application for real-time video display. To deal with this we interact the 19 App groups with the 19 2-digit NACE industries to define 292 *App categories* at the level of App group \times NACE industry.⁷

There are billions of Apps in the App Store and Google Play, many of which have no downloads or just one or two. It is not computationally feasible to deal with terabytes of such data. We therefore initially restrict the sample by selecting the 1,000 most downloaded unified firms. Over our 2014–2020 sample period these unified firms had 223 billion downloads of 82,850 Apps.

We are using download data whereas revenue data would be better. To show that the two are correlated, we divide our table 3 subsample into two bins, one for Apps with revenues and one for Apps without revenues. In figure 1 we plot the kernels of log downloads separately for the two bins. The kernel for revenue-generating App downloads is substantially right-shifted relative to the kernel for free App downloads. This illustrates that more-downloaded Apps tend to be revenue-generating Apps.

⁵This database does not cover downloads from the remaining application marketplaces. The largest of these are Huawei App Gallery, Xiaomi App Store, Amazon App Store and Samsung Galaxy Store. Nor do we track downloads done directly from web pages.

⁶The two marketplaces define groups slightly differently, but it is easy to convert the Google Play groups into Apple Store groups.

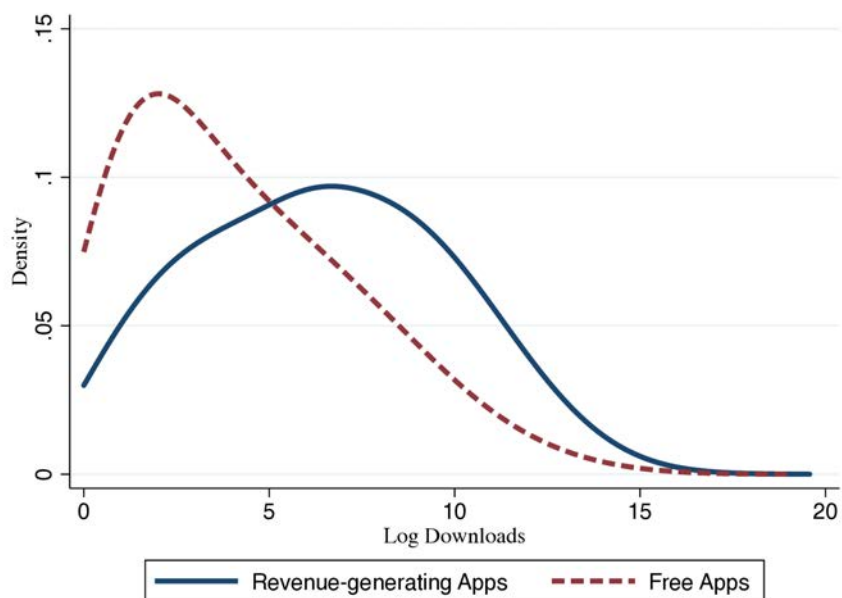
⁷Some of the 19×19 potential App categories have no Apps, leaving us with 292 App categories.

Table 2: App Categories

Group	AI Patents	Single-category AI Patents	Top 3 Apps by Download in 2020
Games	5,560	1,270	Garena Free Fire (Garena, Singapore), PUBG MOBILE (PUBG, Singapore), Subway Surfers (Sybo Game, Denmark)
Photo and Video	4,357	295	Instagram (Facebook, US), Snapchat (Snap, US), Likee (Bigo, US)
Utilities	14,193	632	UC-brosver (Alibaba, China), Truecaller (True Software, Sweden), Chrome (Google, US)
Social Networking	1,448	16	Whatsapp (Facebook, US), Facebook (Facebook, US), Messenger (Facebook, US)
Entertainment	6,417	307	TikTok (ByteDance, China), Netflix (Netflix, US), Youtube (Google, US)
Shopping	3,442	412	Amazon (Amazon, US), Wish (ContextLogic, US), Shopee (Shopee, Singapore)
Music	1,601	528	Spotify (Spotify, Sweden), Youtube Music (Google, US), Shazam (Apple, US)
Finance	11,791	2,608	Google Pay (Google, US), Paypal (Paypal, US), Caixa Tem (Caixa Econômica Federal, Brazil)
Education	13	12	Google Classroom (Google, US), YouTube Kids (Google, US), Duolingo (Duolingo, US)
Productivity	7,621	59	Shareit (SHAREit, China), Gmail (Google, US), Microsoft Word (Microsoft), Word (Microsoft, US)
Business	3,225	2,861	Zoom (Zoom, US), Google Meet (Google, US), Microsoft Team (Microsoft, US)
Lifestyle	20,133	2,871	Pinterest (Pinterest, US), Tinder (IAC, US), Airtel Thanks (Bharti Airtel, Indian)
Sports, Health and Fitness	1,608	1,569	Aarogya Setu (NIC, India), Home Workout (ABISHKING, Singapore), Mi Fit (Xiaomi, China)
Books, News and References	183	175	Wattpad (Wattpad, Canada), Amazon Kindle (Amazon, US), Audible (Audible, US)
Travel	1,892	160	Uber (Uber, US), Google Earth (Google, US), Booking.com (Booking.com, Netherlands)
Food and Drink	52	22	McDonald's (McDonald's, US), Uber Eat (Uber, US), Domino's Pizza (Domino's Pizza, US)
Navigation	6,991	4,041	Google Map (Google, US), Waze (Google, US), Gaode Map (Alibaba, China)
Weather	12,280	10,747	Whether&Radio (WetterOnline, German), The Weather Channel (IBM, US), Whether Forecast (Smart-Pro, Indian)
Medical	303	251	NHS COVID-19(GOV, UK), COCOA(GOV, Japan), Pregnancy + (Health&Parenting, UK)
Total	103,110	28,836	

Notes: Groups are ordered by downloads with the most downloaded group on top. Some firms have patents in multiple categories. Since we cannot assign all patents to groups (see below) and since we want to avoid double counting patents in the “AI Patents” column, we assign all of a firm’s patents to its largest (most downloaded) group. We do this for this table only.

Figure 1: Distributions of log Downloads: Free vs. For-Pay Apps



Notes: The figure displays kernels of the distribution of log downloads. The data are the ‘subsample’ described in table 3. The red and blue lines are for Apps with and without revenues, respectively.

2.2. Linkage to Patent and Financial Data

Our core analysis is about the impact of AI on a variety of trade and welfare outcomes. We will be measuring AI using patent data. We use SensorTower’s unified firm names to link with the Bureau Van Dijk Orbis Intellectual Property database. This provides us with patent and financial data. We were unable to reliably match firm names across the two data sets using machine learning tools. We therefore select the largest 1,000 unified firms in the world (as measured by global downloads) and then find by hand their global ultimate owners in the Orbis database. We match 834 of the 1,000 firms. Unlike many studies, there is no linkage error here.⁸

To investigate the representativeness of our sample we also looked at the 100,000 unified firms with the most downloads – these are not matched to patent and financial data – and call this the ‘full sample.’ Table 3 displays summary statistics for the full sample and our subsample.⁹ Two things stand out. First, our sample is skewed towards unified firms with large downloads (see the 90th percentile column). Second, both samples have 10th percentile downloads that equal 2 so that our sample differs from

⁸We initially used the Python-based FuzzyWuzzy matching algorithm. However, even after extensive pre-cleaning of firm names, a visual inspection of the matching results showed that it was of insufficient accuracy for our comfort. We therefore verify each match by hand. This verification is what constrains us to working with 834 ultimate owners.

⁹Apps that have zero downloads are excluded from this table and from all of our analysis. We re-introduce zeros whenever we do PPML.

Table 3: Summary Statistics of Downloads: Full Sample and Subsample

Sample	N	Mean	Std. Dev.	Percentiles		
				10th	50th	90th
Full Sample	78,042,301	7,819	231,443	2	35	3,514
Subsample	4,733,652	47,204	841,535	2	107	22,496

Notes: This table reports statistics on the number of Apps and their downloads. An observation refers to a unique App×importer×year triplet where importer is the country downloading the App. N is the number of observations. The columns report moments of the distribution of downloads across App-importer-year observations.

the full sample primarily in dropping Apps with extremely small download numbers.

This vividly illustrates that our sample selection criteria do not drop any major apps or firms. One would be hard pressed to recognize any of the Apps excluded from our analysis. The highest-ranked App excluded from our data is slither.io, an obscure action game from Kooapps. The highest-ranked firm not in our data is SayGames, an obscure game startup from Belarus whose most popular App is Twist Hit!. In short, we do not think that our subsample excludes any important Apps or that conclusions drawn from it are biased for our set of questions.

2.3. AI Patent Data

To estimate the impact of AI on trade and welfare we need to be precise about what we mean by AI and how we measure it. We use AI-related patents as the basis of our measure. From the Orbis data we know the 10,144,089 patents assigned to our 834 firms. We categorize each of these patents as AI or non-AI patents following the WIPO (2018, 2019) methodology. For each patent we check if it meets one of three criteria.

1. The main and/or minor CPC codes are on a list of CPC codes that WIPO uses to identify specific AI technologies. For example, CPC subclass G10L-015 is speech recognition.
2. The title and/or abstract contains a phrase that is on a keyword list that WIPO uses to identify specific AI technologies. The list includes phrases such as ‘machine learning’ and ‘neural network’ along with extensions of these phrases such as ‘neural networks’ and ‘neural-network’.
3. Some patents are about AI, but not about a specific AI technology. Here WIPO combines a CPC code with a keyword to identify an AI patent. For example, GTL-

Table 4: Examples for AI Patents

Current Owner	Patent Number	CPC	Specific AI Technology	Keywords	Portion of Abstract
		Classification Criterion			
Method 1: Patent Class					
Facebook	US20190012697A1	G06Q30/0242	G06Q - Data processing systems; 30 - Commerce; 0242 - Determination of advertisement effectiveness	N/A	A client relationship management (CRM) application can generate a ranked list of client engagement tools by computing a rank score for available client engagement tools and determining an order among the available client engagement tools based on the rank scores. The CRM application can use one or more trained prediction models and business rules to compute a prediction for success for client engagement tools.
Method 2: Keywords					
Microsoft	EP3424044A1	N/A	N/A	deep learning	The technology described herein uses a modular model to process speech. A deep learning based acoustic model comprises a stack of different types of neural network layers.
Method 3: Patent Class plus Keywords					
Microsoft	KR1020130110565A	G06Q10/109	G06N - Computer systems based on specific computational models; 10 - Administration; Management; 109 - Time management	predictive models	The present invention relates to a system and methodology to facilitate collaboration and communications between entities such as between automated applications, parties to a communication and/or combinations thereof. The systems and methods of the present invention include a service that supports collaboration and communication by learning predictive models that provide forecasts of one or more aspects of a users' presence and availability.

013 is speech synthesis (text to speech), which may or may not involve AI. However, if a patent in CPC subclass GTL-013 has a title or abstract with keywords such as 'backpropagation' or 'self learning' then WIPO identifies it as an AI patent.

Table 4 gives examples of AI patents identified through each of the above three methods. We have duplicated the WIPO methodology with one exception. Their keyword search is over the English title, English abstract, English claims and English object of invention. Our keyword search is over the English title and English abstract.

Our procedure identifies 103,110 patents as AI patents and the remaining 10,038,168 patents as non-AI patents. Column 2 of table 2 above displays the number of AI patents by App group. Among the 834 firms, 309 firms own at least one AI patent. Finally, our AI patents grow rapidly from 1990 to 2020.

When we speak of a firm's AI patents in year t we will mean its cumulative AI patent applications from 1990 to year t . That is, our AI patents are a *stock* of patent applications. Using applications rather than grants avoids the worst of the right-truncation problem associated with delays in granting patents.

2.4. AI Deployment in Apps

We require a measure of the AI deployed in each App category. If each firm's Apps were in a single App category this would be an easy matter of counting the AI patents of all firms producing Apps in the category. Unfortunately, a large number of firms (ultimate owners) have Apps in multiple categories. For example, the ultimate owner Alphabet controls Google, Nest, YouTube, Waze and Fitbit, all of which operate in different App categories. We therefore purify our measure of the AI deployed in App categories by adapting to our setting the approach of De Loecker, Goldberg, Khandelwal, and Pavcnik (2016). In estimating production-function parameters they only include single-product firms. We define a single-category firm as a firm whose primary category accounts for over 85% of its total downloads.¹⁰ Column 3 of table 2 above shows the AI patents owned by single-category firms. There are 28,836 such patents and they account for a substantial 28% of all AI patents in our sample. There are 549 single-category firms (including some with zero patents) among the 834 firms in our sample. Together they account for 42% of all downloads in our sample.

We construct our measure of the AI deployment of an App category only from the patents of single-category firms. Let $Patent_{cxt}$ and K_{cxt} be AI patent application stocks and total assets summed across all single-category firms in country x in year t that produce Apps only in category c . $AI_{cxt} = Patent_{cxt}/K_{cxt}$ will be our key measure of AI deployment.¹¹

It is of independent interest to know about the AI patents of multi-category firms. In every specification reported below we have also examined the same specification but with the addition of a variable that captures the AI deployment of multi-category firms. To this end, we define $multipleAI_{cxt}$ as the total AI patent application stocks over total assets for multi-category firms from country x with Apps in category c in year t .¹² To control for the general effects of patenting, we also construct $nonAI_{cxt}$ as the total non-AI patent application stocks over total assets for all firms in category c , country x and year t . In our regressions, adding these two variables *never* affects the magnitude or statistical significance of the coefficients on our AI deployment variable AI_{cxt} .

¹⁰Using 95%, 90% or 80% as the threshold does not affect our main results.

¹¹We are grateful to an anonymous referee for suggesting that we scale patents to control for the size of firms in the cxt bin. We choose assets because, relative to other variables in Orbis that we could use for scaling such as employment, assets have few missing values. An earlier version of this paper which did not scale reported similar results.

¹²In constructing $multipleAI_{cxt}$, a multi-category firm's AI patents are given to each of its categories e.g., if a firm has 10 patents and operates in two categories we do not know which patent applies to which category (indeed, some AI patents may apply to both) so we assume that the firm has 10 patents in *each* category.

Table 5: Summary Statistics

	N	mean	sd
App Downloads:			
Total downloads: $\ln(y_{cmxt})$	469,879	6.45	3.95
Number of Apps: $\ln(N_{cmxt})$	469,879	1.05	1.19
Average downloads: $\ln(\bar{y}_{cmxt})$	469,879	5.40	
Patent Variables:			
AI Patents: $\ln(1 + AI_{cxt})$	469,879	0.13	0.60
AI Multiple: $\ln(1 + multipleAI_{cxt})$	469,879	0.60	0.90
Non-AI patents: $\ln(1 + nonAI_{cxt})$	469,879	2.96	2.78

Notes: The table presents summary statistics of variables used in the gravity equations of Section 3. Each of the 469,879 observations is uniquely identified by an App category c , a downloading user country or importer m , a producing country or exporter x , and a year $t = 2014, \dots, 2020$. To be included, the $cmxt$ observation must have strictly positive downloads.

2.5. Summary Statistics

Table 5 reports summary statistics of our data. Each observation is uniquely identified by an App category (292), an exporter (53), an importer (84) and a year (2014–2020). We have 469,879 observations with positive levels of downloads. There are several points to note about the sample size. First, we do not work at the firm level and this requires an explanation. We do not know whether any given App uses AI so we cannot work at the level of a firm’s Apps. What we do know is the extent to which AI is deployed in an App category in an exporter country. So we must aggregate up from firms to App-categories and exporters. Second, we exclude zero downloads, but return to this below using PPML. Third, Google Play is banned in China so that we only have Apple App Store data for Chinese downloads. We thus exclude observations for which China is the importer. Note however that we keep China as an exporter and that including China as an importer makes no difference to our results.

Table 5 reports the dimensions of each variable. These are App category c , importer m , exporter x , and year t . We winsorize the top 1% of observations for the download and patent variables. From the first line of the table, the mean downloads of a $cmxt$ observation is 633 ($= e^{6.45}$), the mean number of Apps is 2.85 ($= e^{1.05}$), and the mean downloads per App is 221 ($= e^{5.40}$). The latter illustrates that when we report results *within* App category, the analysis is at a very fine level.

3. AI and Trade: Bilateral Gravity

We estimate the following gravity equation:

$$\ln(y_{cmxt}) = \beta \ln(1 + AI_{cxt}) + \theta X_{cmxt} + \alpha_{mxt} + \alpha_{cm} + \varepsilon_{cmxt} . \quad (1)$$

In this regression y_{cmxt} is downloads by consumers in country m of Apps in category c produced by firms headquartered in country x . We are interested in international trade in this section so we exclude domestic observations i.e., observations for which the importer is the exporter. Including these observations does not affect our conclusions. Since we only include non-zero trade flows, $y_{cmxt} \geq 1$. Our key independent variable is $\ln(1 + AI_{cxt})$ and our hypothesis is that AI deployment increases trade ($\beta > 0$). X_{cmxt} is a set of gravity variables. α_{mxt} and α_{cm} are the fixed effects. The only other fixed effect that we can add while still identifying β is α_{cx} . Adding these weakens our results because $\ln(1 + AI_{cxt})$ has relatively limited variation across time. Aside from this, our results are not at all sensitive to the choice of fixed effects.

Table 6 reports the OLS results. In column 1, we examine whether standard gravity covariates from CEPII behave the same way for App trade as they do for goods trade.¹³ To this end we consider the full sample, that is, before restricting it by linking to Orbis data. (See section 2.2.) We include an importer fixed effect, an exporter fixed effect, a year fixed effect and a category fixed effect. Log distance between m and x matters, but is much smaller than the median estimate of -0.85 reported in Head and Mayer's (2014) meta-analysis of gravity studies. That distance plays less of a role in digital trade will come as no surprise. The coefficient on contiguity is a little smaller than in Head and Mayer and the coefficient on common language is just a little larger. We also include dummies for whether m and x were ever in a colonial relationship and whether they are in the same regional trade agreement. These covariates are less significant and much smaller than in Head and Mayer. The importer GDP and exporter GDP coefficients are very small, but this is not surprising given that they do not vary much over our period 2014-2020 and so are largely soaked up by the fixed effects. We do not include the populations of either m or x because these are also largely soaked up by fixed effects.

In column 2, we use our subsample matched with patent and financial data to estimate equation 1. The coefficients on the gravity covariates do not change, which provides evidence that our sample is representative in dimensions familiar to trade economists. Crucially, the estimate of β is positive and significant. In OLS, AI deployment is correlated with downloading.

¹³See Head and Mayer (2014) and Head, Mayer, and Ries (2010). Since the CEPII data end in 2019 we linearly extrapolate the time-varying variables by one year to 2020.

Table 6: Gravity Equation: Independent Variable is $\ln(y_{cmt})$

	OLS						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln(1 + AI_{cxt})$		1.21*** (0.194)	1.33*** (0.225)	1.20*** (0.203)	1.13*** (0.210)	1.04*** (0.200)	1.04*** (0.197)
$\ln(1 + multipleAI_{cxt})$					1.02*** (0.239)	0.80*** (0.252)	0.80*** (0.246)
$\ln(1 + nonAI_{cxt})$						0.14*** (0.042)	0.14*** (0.041)
$\ln(\text{Distance}_{mx})$	-0.39*** (0.036)	-0.37*** (0.058)	-0.23*** (0.038)				
Contiguous _{mx}	0.38*** (0.073)	0.35*** (0.105)	0.31*** (0.106)				
Common Language _{mx}	0.76*** (0.106)	0.53*** (0.097)	0.39*** (0.105)				
Colonial Dependence _{mx}	0.08 (0.138)	0.39*** (0.122)	0.31*** (0.115)				
Regional Trade Agreement _{mxt}	0.11** (0.054)	0.03 (0.067)	-0.10 (0.066)				
$\ln(\text{GDP}_{xt})$	0.04 (0.336)	0.18 (0.335)	-0.01 (0.302)				
$\ln(\text{GDP}_{mt})$	-0.11 (0.630)	-0.01 (0.321)	-0.06 (0.340)				
Social Connectedness Index _{mx}			0.18*** (0.055)				
Constant	12.43 (13.964)	5.53 (9.253)	8.29 (9.021)	6.29*** (0.022)	5.69*** (0.144)	5.42*** (0.168)	5.43*** (0.168)
Observations	774,414	468,679	399,873	465,955	465,955	465,955	464,567
Fixed effects	<i>t, m, x, c</i>	<i>t, m, x, c</i>	<i>t, m, x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, m-c</i>
R^2	0.559	0.386	0.396	0.429	0.448	0.451	0.486
Within R^2	0.033	0.040	0.044	0.027	0.059	0.063	0.065

Notes: Each observation is an App category (c), a downloading country or importer (m), an App producing country or exporter (x) and a year ($t = 2014, \dots, 2020$). The dependent variable is the log of the number of downloads, $\ln(y_{cmt})$. In column 1 we use the full sample covering all Apps (we do not restrict the sample to firms that can be linked to Orbis). In columns 2–7, we use our subsample of 834 firms to construct a panel of 292 App categories, 53 exporters, 84 importers, and 7 years. In the fixed effect rows, $t-m-x$ and $m-c$ refer to year-importer-exporter and importer-category fixed effects, respectively. The number of observations is degrees-of-freedom corrected as calculated by Stata's `reghdfe` command and so declines as more fixed effects are added. Standard errors are based on two-way clustering by importer and by exporter. ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels, respectively.

In column 3 we include the Bailey *et al.* (2020) index of pairwise social connectedness. Their index is based on an anonymized snapshot of all friendship links on Facebook. It is the log of the relative probability of a friendship link between a Facebook user in m and a Facebook user in x . The coefficient on social connectedness of 0.18 is significant at the 1% level though smaller than in Bailey *et al.*. However, when we use the full sample of column 1 the coefficient rises to 0.25, which is close to what they report. More importantly, the introduction of social connectedness does not affect the estimated coefficient on our key AI_{cxt} variable and indicates that what we are finding is very different from the channel identified by Bailey *et al.*¹⁴

In columns 4, we introduce year-importer-exporter and category fixed effects. It makes little difference to our estimates of the coefficient on $\ln(1 + AI_{cxt})$.

In columns 5–7 of table 6 we add two additional covariates, $\ln(1 + multipleAI_{cxt})$ and $\ln(1 + nonAI_{cxt})$. In column 5, AI deployment for multiple category firms is significant. More importantly, its inclusion does not affect the coefficient on our key AI variable $\ln(1 + AI_{cxt})$. In column 6, non-AI patents $\ln(1 + nonAI_{cxt})$ is significant, but as we shall see its economic magnitude is half that of $\ln(1 + AI_{cxt})$. If these patents are correlated with AI patents then it is possible that our AI results are just proxying for the effects of patenting in general; however, inclusion of $\ln(1 + nonAI_{cxt})$ has little effect on the coefficient on our key AI deployment variable.

In column 7, we introduce importer-category fixed effects. The coefficient on $\ln(1 + AI_{cxt})$ does not change.

3.1. IV

Our OLS results potentially suffer from the endogeneity of AI deployment. There are two obvious sources of bias. The first is reverse causality and/or omitted variables: firms with high levels of downloads may have other characteristics such as size that justify investing in AI. See Lileeva and Trefler (2010) for a discussion. In this case we expect IV to be smaller than OLS. The second is heterogeneous impacts of the type addressed by Imbens and Angrist's (1994) LATE estimator. We expect that the returns to AI are higher for firms that invest than for firms that do not. If so, IV will overestimate the mean impact of AI deployment (Card, 2001, eq. 11) and, by implication, IV may be larger than OLS.

An ideal instrument is an exogenous cost shock to the deployment of AI i.e., a shock that exogenously drives AI deployment. The comparative advantage logic of

¹⁴The coefficient is 1.21 in column 2 and 1.33 in column 3. However, this difference is entirely due to the difference in samples rather than to the inclusion of social connectedness. If we redo column 2 with the smaller sample of column 3, the coefficient in column 2 rises to exactly 1.33.

Heckscher-Ohlin (HO) provides such a shock. The cost of AI deployment by producers of product c in country x is low if (1) the *country* is abundant in AI and (2) the *product* is AI intensive. We measure a country’s AI abundance using the number of AI conference papers presented by scholars from exporter country x in year t . Denote this by $ConfPaper_{xt}$. This is a commonly used measure of a country’s AI capacity. See for example Goldfarb and Trefler (2019a). Data are from Zhang *et al.* (2021). We measure the AI intensity of a product or App category as the sum of all single-category-firm AI patents for firms in category c divided by the sum of all single-category-firm assets for firms in category c . This is calculated at the global level, meaning we sum across firms in *all* countries. Further, it is calculated separately for each year. Denote this by AI_{ct} and note that as in the HO literature, it is a global variable rather than an exporter-level variable. Our instrument for $\ln(1 + AI_{cxt})$ is then $\ln(1 + AI_{ct}) \cdot (ConfPaper_{xt})$. Note the interaction of country (x) and product (c) characteristics, which is the fundamental core of all comparative advantage theories. More specifically, our first stage will look a lot like the test of HO comparative advantage in Romalis (2004).

Table 7 reports our IV estimates. Panel B reports the first-stage, that is, a regression of our endogenous variable $\ln(1 + AI_{cxt})$ on our instrument $\ln(1 + AI_{ct}) \cdot \ln(ConfPaper_{xt})$. Only the coefficients on the instrument are reported. These coefficients are all positive and statistically significant. Further, the Kleibergen-Paap weak-instruments F -statistic hovers around the Stock-Yogo significance threshold of 20.

The IV estimates of the coefficient on AI deployment appear in Panel A of table 7. Columns 1–6 correspond to columns 2–7 of table 6, respectively. The remaining regressors are included but not reported. The IV results are somewhat bigger than the OLS results, which suggests that heterogenous impacts are more important than reverse causality and/or omitted variables. While the IV results are larger than OLS, the difference is small relative to the IV standard error.

3.2. Economic Magnitudes when Patents are Right Skewed

In this section we explore an alternative specification that makes it easier to interpret the size of the impact of AI deployment on exports of Apps. The specification also addresses a major concern that arises in the patent literature. A small number of firms hold a large fraction of all patents and of all patent citations, leading to a concern that the impacts of AI deployment are significant only for a small number of large firms and insignificant for all other firms. See Aghion, Bergeaud, Lequien, and Melitz (2018) and Lim, Trefler, and Yu (2019) for a discussion. In this section we investigate an alternative specification

Table 7: Gravity Equation: Instrumental Variables

Panel A. IV						
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(1 + AI_{cxt})$	1.75*** (0.415)	1.47*** (0.332)	1.70*** (0.427)	1.46*** (0.433)	1.54*** (0.460)	1.56*** (0.474)
$\ln(1 + multipleAI_{cxt})$				1.01*** (0.244)	0.82*** (0.235)	0.82*** (0.230)
$\ln(1 + nonAI_{cxt})$					0.11*** (0.040)	0.11*** (0.040)
Observations	468,679	399,873	465,955	465,955	465,955	464,567
Gravity covariates	yes	yes	no	no	no	no
Social Connectedness	no	yes	no	no	no	no
Fixed effects	<i>t, m, x, c</i>	<i>t, m, x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, m-c</i>
Panel B. First Stage						
	(1)	(2)	(3)	(4)	(5)	(6)
Instrument	0.613*** (0.138)	0.720*** (0.0835)	0.602*** (0.143)	0.598*** (0.142)	0.622*** (0.135)	0.612*** (0.133)
K-P <i>F</i> -value	19.78	74.45	17.65	17.71	21.16	21.33
Fixed effects	<i>t, m, x, c</i>	<i>t, m, x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, m-c</i>

Notes: This table reports the IV counterparts to the OLS results of table 6. Columns 1–6 correspond respectively to columns 2–7 of table 6. We suppress the estimates of the gravity and social connectedness coefficients. Panel A displays the IV estimates and panel B displays the first stage. In the first stage the dependent variable is $\ln(1 + AI_{cxt})$ and the independent variable is the Heckscher-Ohlin instrument $\ln(1 + AI_{ct}) \cdot ConfPaper_{xt}$. All other first-stage coefficients are suppressed. Standard errors are based on two-way clustering by importer and by exporter. ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels, respectively. See the notes to table 6 for details.

Table 8: Gravity Equation: Nonparametrics and Magnitudes

Panel A. OLS					
	(1)	(2)	(3)	(4)	(5)
$\ln(1 + AI_{cxt})$					
First positive tercile	2.12*** (0.289)	2.14*** (0.299)	2.14*** (0.306)	1.98*** (0.271)	1.97*** (0.275)
Second positive tercile	1.81*** (0.237)	1.81*** (0.237)	1.94*** (0.255)	1.69*** (0.259)	1.67*** (0.253)
Third positive tercile	1.63*** (0.335)	1.61*** (0.346)	1.70*** (0.301)	1.37*** (0.304)	1.37*** (0.303)
$\ln(1 + multipleAI_{cxt})$					
First positive tercile			0.37 (0.331)	0.19 (0.345)	0.20 (0.334)
Second positive tercile			0.97*** (0.296)	0.62** (0.301)	0.63** (0.288)
Third positive tercile			1.99*** (0.436)	1.39*** (0.457)	1.40*** (0.441)
$\ln(1 + nonAI_{cxt})$					
First positive tercile				-0.14 (0.309)	-0.14 (0.310)
Second positive tercile				0.16 (0.212)	0.15 (0.218)
Third positive tercile				0.84*** (0.288)	0.83*** (0.283)
Observations	468,679	465,955	465,955	465,955	464,567
Fixed effects	<i>t, m, x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, m-c</i>
R^2	0.384	0.428	0.442	0.446	0.482
Within R^2	0.037	0.025	0.048	0.054	0.057
Panel B. IV					
	(1)	(2)	(3)	(4)	(5)
$\ln(1 + AI_{cxt})$					
First positive tercile	2.13*** (0.301)	2.09*** (0.312)	1.97*** (0.299)	1.86*** (0.267)	1.87*** (0.274)
Second positive tercile	2.27*** (0.352)	2.21*** (0.350)	2.15*** (0.374)	2.05*** (0.357)	2.08*** (0.362)
Third positive tercile	1.91*** (0.197)	1.84*** (0.198)	1.87*** (0.185)	1.71*** (0.162)	1.75*** (0.171)
K-P F -value	40.91	39.20	39.57	42.49	40.99

Notes: The dependent variable is the log number of downloads ($\ln(y_{cmxt})$). An observation is uniquely identified by the App category (c), the exporter (x), the importer (m), and the year (t). We break AI deployment AI_{cxt} into four dummies. The omitted dummy is for observations with $AI_{cxt} = 0$. The remaining three dummies are for the terciles of the distribution of AI_{cxt} conditional on $AI_{cxt} > 0$. Likewise for $\ln(1 + multipleAI_{cxt})$ and $\ln(1 + nonAI_{cxt})$. Note that the specification in column 1 includes all the same gravity equation regressors as appear in column 1 of table 6, but these are not reported. The first-stage results appear in table A1. See the notes to table 6 for details.

that is more robust to the right skew of the patent distribution and that yields easily interpreted coefficient magnitudes.

We divide AI_{cxt} into four groups. The first is all observations with $AI_{cxt} = 0$. We then take the remaining observations and divide them into terciles of the distribution of strictly positive AI_{cxt} . Table 8 reports the results. Consider column 1 of Panel A which reports our OLS results for each of the three tercile dummies of AI_{cxt} . The omitted category is observations with $AI_{cxt} = 0$. There is no evidence that impacts vary across terciles: The F -statistic for the test of equality of the three tercile coefficients is tiny across all specifications ($F \approx 1.2$, $p \approx 0.30$). This is useful because it shows that our results are not driven by the upper end of the distribution of patents; rather, our estimates are homogeneous across the distribution of patents.

Turning to coefficient magnitudes, consider two exporters of category- c Apps, one exporter having $AI_{cxt} = 0$ and the other having AI_{cxt} in the first tercile. From column 1, the latter has downloads that are 2.12 log points higher or 8.3 times higher ($8.3 = e^{2.12}$).

Adding additional covariates, as in columns 3–5, does not alter this conclusion. In columns 3–4 we add terciles of multiple-category AI and non-AI patents. For non-AI patents we see that the results are driven entirely by the high-patenting observations, as we have come to expect from the patent literature. It is reassuring to see this for non-AI patents where we expect them, but not for our AI deployment measure. Also note that the three non-AI tercile coefficients are jointly insignificant at the 1% level in columns 4–5 ($F \approx 4$, $p \approx 0.012$). In columns 2 and 5 we add finer fixed effects and this has no impact.

IV results appear in panel B of table 8. There are now three endogenous variables (the terciles of $\ln(1 + AI_{cxt})$) so that we must be very cautious in lending too much weight to the results. We create three instruments by interacting our single instrument $\ln(1 + AI_{ct}) \cdot ConfPaper_{xt}$ with tercile dummies. The first-stage results are reported in table A1 and are very strong. This is apparent from the K-P weak instruments F -statistic of approximately 40 reported at the bottom of panel B. It is well above the threshold of 20. What is remarkable about the IV results is that they are almost identical to the OLS results. This raises our confidence in the causal interpretation of our results.

Looking at the IV coefficients on the tercile dummies for $\ln(1 + AI_{cxt})$, the smallest value is 1.71. We use this as a conservative guide to our headline number: *AI deployment leads to a sixfold increase in downloads* ($5.52 = e^{1.71}$). *This is a very large effect.*

4. Product Variety: The Extensive Margin of Trade

We now examine the number of varieties traded in a bilateral relationship. This is called the extensive margin of trade. Let N_{cmxt} be the number of category- c Apps from country x available to consumers in country m in year t . We view each App within an App

category as a variety. For example, Chrome (Google), Baidu (Baidu), Internet Explorer (Microsoft), and Safari (Apple) are varieties of browsers. Following Eaton, Kortum, and Kramarz (2011) we decompose total downloads into average downloads per App times the number of Apps. Mathematically,

$$\ln y_{cmt} = \ln \bar{y}_{cmt} + \ln N_{cmt} \quad \text{where } \bar{y}_{cmt} \equiv y_{cmt} / N_{cmt} . \quad (2)$$

$\ln \bar{y}_{cmt}$ corresponds to the intensive margin and $\ln N_{cmt}$ corresponds to the extensive margin or number of varieties. We estimate

$$\ln(N_{cmt}) = \beta \ln(1 + AI_{cmt}) + \alpha_{mct} + \alpha_{cm} + \varepsilon_{cmt} .$$

That is, we estimate the same equations as before, but with a different dependent variable.

Table 9 reports the results. From panel A, AI deployment is associated with greater numbers of bilaterally traded Apps and this result is robust across specifications. The coefficient is about half the size of the coefficient when the dependent variable is total downloads. Panel B reports IV results. These are larger than the OLS results, but not statistically so.

When looking at the extensive margin, the issue of zero trade flows looms large. To investigate, instead of omitting observations with zero downloads, we change the dependent variable from $\ln N_{cmt}$ to N_{cmt} and keep zero-download observations ($N_{cmt} = 0$). This doubles the number of observations. We then use PPML estimation. The results appear in panel C of table 9 and are very similar to the OLS results, indeed identical in columns 4–5.

To get a clearer sense of magnitudes and to ensure that our specifications are robust to firms with very large numbers of patents, we return to our analysis of terciles. In table 10 we repeat table 8 with just a single change: the dependent variable is now the log of the number of varieties $\ln N_{cmt}$. There is evidence of modest coefficient heterogeneity across terciles, but otherwise the conclusions here about the impact of AI deployment on varieties are very similar to those about impacts on downloads. Since the IV and OLS results are very similar, we do not report the former. Averaging across the tercile coefficients in column 5 we get 0.81, which drives our headline number that *AI deployment doubles the number of imported varieties* ($2.25 = e^{0.81}$).¹⁵

¹⁵For the intensive margin ($\ln \bar{y}_{cmt}$), these results are very similar to the results for $\ln(N_{cmt})$ that we reported in table 9. Restated, the intensive-margin effects are very significant and half the size of the total effects. This is true for OLS, IV and PPML. We do not report these results.

Table 9: Product Variety and the Extensive Margin

Panel A. OLS $\ln(N_{cmxt})$					
	(1)	(2)	(3)	(4)	(5)
$\ln(1 + AI_{cxt})$	0.49*** (0.064)	0.50*** (0.065)	0.47*** (0.067)	0.47*** (0.069)	0.47*** (0.069)
$\ln(1 + multipleAI_{cxt})$			0.37*** (0.090)	0.35*** (0.098)	0.35*** (0.097)
$\ln(1 + nonAI_{cxt})$				0.01 (0.019)	0.01 (0.019)
Observations	468,679	465,955	465,955	465,955	464,567
Fixed effects	t, m, x, c	$t-m-x, c$	$t-m-x, c$	$t-m-x, c$	$t-m-x, m-c$
R^2	0.461	0.478	0.504	0.504	0.525
Within R^2	0.056	0.053	0.100	0.100	0.103
Panel B. IV $\ln(N_{cmxt})$					
	(1)	(3)	(4)	(5)	(6)
$\ln(1 + AI_{cxt})$	0.69*** (0.108)	0.70*** (0.119)	0.62*** (0.111)	0.62*** (0.108)	0.63*** (0.113)
$\ln(1 + multipleAI_{cxt})$			0.36*** (0.089)	0.35*** (0.093)	0.36*** (0.091)
$\ln(1 + nonAI_{cxt})$				0.00 (0.018)	0.00 (0.018)
Observations	468,679	465,955	465,955	465,955	464,567
Fixed effects	t, m, x, c	$t-m-x, c$	$t-m-x, c$	$t-m-x, c$	$t-m-x, m-c$
Panel C. PPML N_{cmxt}					
	(1)	(2)	(3)	(4)	(5)
$\ln(1 + AI_{cxt})$	0.53*** (0.132)	0.55*** (0.137)	0.51*** (0.103)	0.47*** (0.095)	0.47*** (0.095)
$\ln(1 + multipleAI_{cxt})$			0.77*** (0.173)	0.59*** (0.173)	0.59*** (0.173)
$\ln(1 + nonAI_{cxt})$				0.13* (0.070)	0.13* (0.069)
Observations	958,974	944,599	944,599	944,599	914,186
Fixed effects	t, m, x, c	$t-m-x, c$	$t-m-x, c$	$t-m-x, c$	$t-m-x, m-c$

Notes: Each observation is an App category (c), a downloading country or importer (m), an App-producing country or exporter (x) and a year ($t = 2014, \dots, 2020$). We use a panel of 292 App categories, 53 exporters, 84 importers, and 7 years. The dependent variable is the number of imported Apps in category c : $\ln(N_{cmxt})$ in panel A (OLS), $\ln(N_{cmxt})$ in panel B (IV), and N_{cmxt} in panel C (PPML). For PPML we keep observations with $N_{cmxt} = 0$. For IV, the first stage already appears in panel B of table 7 and so is not repeated here. The specification in column 1 includes the same gravity regressors as in column 1 of table 6, but these are not reported. In the fixed effect rows, $t-m-x$ and $m-c$ refer to year-importer-exporter and importer-category fixed effects, respectively. Standard errors are based on two-way clustering by importer and by exporter. ***, ** and * indicates statistical significance at the 1%, 5% and 10% levels, respectively.

Table 10: Extensive Margin $\ln(N_{cm,xt})$: Nonparametric and Magnitudes

	OLS				
	(1)	(2)	(3)	(4)	(5)
$\ln(1 + AI_{cxt})$					
First positive tercile	0.95*** (0.082)	0.96*** (0.081)	0.95*** (0.073)	0.91*** (0.074)	0.91*** (0.074)
Second positive tercile	0.94*** (0.129)	0.95*** (0.131)	1.00*** (0.142)	0.95*** (0.151)	0.95*** (0.149)
Third positive tercile	0.60*** (0.116)	0.61*** (0.122)	0.64*** (0.109)	0.56*** (0.123)	0.58*** (0.121)
$\ln(1 + multipleAI_{cxt})$					
First positive tercile			0.10 (0.115)	0.06 (0.128)	0.07 (0.126)
Second positive tercile			0.31*** (0.099)	0.23** (0.107)	0.24** (0.104)
Third positive tercile			0.81*** (0.151)	0.67*** (0.169)	0.68*** (0.163)
$\ln(1 + nonAI_{cxt})$					
First positive tercile				0.02 (0.118)	0.02 (0.121)
Second positive tercile				0.04 (0.116)	0.03 (0.118)
Third positive tercile				0.22* (0.124)	0.22* (0.124)
Observations	468,679	465,955	465,955	465,955	464,567
Fixed effects	<i>t, m, x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, m-c</i>
R^2	0.465	0.482	0.508	0.510	0.531
Within R^2	0.063	0.061	0.107	0.111	0.115

Notes: This table is identical to Panel A of table 9. The only difference is in the treatment of the patent variables. We break AI_{cxt} into four dummies. The omitted dummy is for observations with $AI_{cxt} = 0$. The remaining three dummies are for the terciles of the distribution of AI_{cxt} conditional on $AI_{cxt} > 0$. Likewise for the multi-category patents and the non-AI patents. The specification in column 1 includes all the same gravity equation regressors as in column 1 of table 6, but these are not reported. See the notes to table 6 for details.

5. Creative Destruction

We start by reviewing the literature on estimating the welfare gains from new products and product churning. This will motivate the empirics. We then turn to a brief review of the literature on creative destruction.

5.1. The Welfare Gains from New Products

Feenstra (1994) considers a consumer having CES preferences with an elasticity of substitution σ . He explains how to construct expenditure functions and how to do welfare analysis in the presence of a changing set of varieties induced by an arbitrary shock such as the introduction of AI or a trade-liberalizing event. Let E_{t-1} and E_t be the expenditure functions pre- and post-shock. His now famous formulation is that the impact of changing sets of varieties on welfare is captured by an extra multiplicative term in the expression for E_t/E_{t-1} . This extra term is constructed as follows. Define

$$\lambda_t = 1 - \frac{\text{year } t \text{ expenditures on new varieties (varieties available in } t \text{ but not } t-1)}{\text{year } t \text{ expenditures on all varieties that are available in } t}. \quad (3)$$

Then the extra multiplicative term is

$$\left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{-1/(\sigma-1)}. \quad (4)$$

Feenstra (2010, ch. 2) offers a nice review of this result.

In this formula one can also interpret λ_t and λ_{t-1} as actual data and counterfactual data, respectively. For example, equation (4) can be related to a familiar result in Arkolakis, Costinot, and Rodriguez-Clare (2012). Reinterpreting new varieties as imported goods, λ_t is the share of expenditures on domestic varieties using actual data. Letting $\lambda_{t-1} = 1$ be expenditure shares on domestic data in the counterfactual of autarky, equation (4) reduces to $(\lambda_t)^{-1/(\sigma-1)}$. This is the familiar Arkolakis *et al.* (2012) formula for the gains from trade when moving from autarky to the existing level of period- t trade restrictions. In similar fashion, we will interpret λ_t as actual data in a world with AI and λ_{t-1} as a counterfactual in a world in which there is no AI.

Broda and Weinstein (2006) is an important empirical application of equation (4) to international trade. They find that the number of new varieties made available to US consumers through imports tripled between 1972 and 2001 and this resulted in welfare gains valued at 2.6% of GDP. Feenstra (2010, table 2.1) finds that if all countries in the world moved from autarky to their 1996 levels of trade, welfare gains would be valued at 12.5% of world GDP. See Melitz and Trefler (2012) for further discussion. The Melitz (2003) model adds firm-level selection to the discussion of why varieties are created and

destroyed by international trade. Trade reduces the number of domestic varieties and increases the number of foreign varieties. The net effect is ambiguous. Hsieh *et al.* (2020) revisits the above Broda and Weinstein (2006) analysis and the Trefler (2004) analysis of the Canada-US Free Trade Agreement and shows that the net effect of changes in varieties was negative.

5.2. *The Welfare Gains from Creative Destruction*

The previous subsection dealt with CES-based models. Because CES goods are complements, more varieties are preferred to fewer varieties. An alternative approach emphasizes vertically differentiated goods, that is, goods differentiated by quality. Vertical differentiation underpins models of growth through creative destruction. By innovating, a firm can generate a profit by displacing an existing lower-quality good with its own higher-quality good. This process has come to be known as *creative destruction*. See Aghion and Howitt (1992) and Akcigit and Kerr (2018) for closed-economy models and Grossman and Helpman (1991a,b) for both closed- and open-economy models. Aghion, Bergeaud, Boppart, Klenow, and Li (2019) explore the role of creative destruction for measuring US growth. While their primarily empirical paper treats innovation as exogenous, they provide formulas related to equation (4). We now turn to estimating the impact of AI on creative destruction and plug the estimates into equation (4) to generate welfare calculations.

5.3. *AI and Creative Destruction in the Global Economy: A First Look*

In this subsection we look at the raw data on AI and creative destruction, that is, on how the download shares of new and exiting Apps are impacted by AI. Consider an App in App category c that is downloaded by country m . The App is ‘new’ in year t if it was downloaded in t , but not $t - 1$. The App is ‘exiting’ in year t if it was downloaded in $t - 1$ and t , but not $t + 1$. The App is ‘continuing’ in year t if it was downloaded in $t - 1$, t and $t + 1$. For each cmt triplet, let Ω_{cmt}^n , Ω_{cmt}^e , and Ω_{cmt}^c be the sets of *new*, *exiting* and *continuing* Apps, respectively.

Let ω index Apps and let $y_{cmt}(\omega)$ be downloads of App ω in category c by country m in year t . For each cmt triplet let $y_{cmt}^k \equiv \sum_{\omega \in \Omega_m^k} y_{cmt}(\omega)$ be type- k downloads where k indexes new Apps ($k = n$), exiting Apps ($k = e$), or continuing Apps ($k = c$). All Apps fall into one and only one of these three types. The share of type- k App downloads is

$$\theta_{cmt}^k = \frac{y_{cmt}^k}{\sum_{k'} y_{cmt}^{k'}} = \frac{\text{country } m\text{'s downloads of type-}k\text{ Apps in category } c \text{ and year } t}{\text{country } m\text{'s downloads of all Apps in category } c \text{ and year } t}. \quad (5)$$

The denominator is total downloads for cmt (including downloads of domestically produced Apps). Since these shares are calculated using data for $t - 1$ and/or $t + 1$ we drop the θ_{cmt}^k for the first and last years and work with $t = 2015, \dots, 2019$.

We are interested in how AI impacts the entry of Apps (θ_{cmt}^n) and the exit of Apps (θ_{cmt}^e). While our interest is in what is consumed at the cmt level, our AI deployment measure is about what is produced at the cxt level. Therefore, for each importer m and App category c we take the average of the AI_{cxt} across exporters x that export to m . We use a weighted average with weights proportional to importer m 's downloads of c . As is common in international trade regressions we will be exploiting how the composition of exporters of c Apps varies across importers m e.g., Vietnam imports social networking from China (WeChat) while Canada imports it from the US (Facebook). Mathematically, let $w_{cmtx} \equiv y_{cmtx} / \sum_{x'} y_{cmtx'}$ be the share of m 's downloads originating from producer country x . (Again, we include domestic downloads $m = x$.) Then our key importer-level independent variable is

$$AI_{cmt} = \sum_x w_{cmtx} \cdot AI_{cxt} . \quad (6)$$

Table 11 reports some basic sample statistics on creative destruction. There are 78,741 category-importer-year (cmt) observations in our data, which includes zero-download observations. The left panel of table 11 reports cross-tabulations for whether there was entry ($\theta_{cmt}^n > 0$) and whether there was AI deployment ($AI_{cmt} > 0$). Among observations with positive AI deployment, 91% have some entry. In contrast, among observations with no AI deployment, only 66% have some entry. Thus, AI is (non-causally) correlated with the entry of new varieties. The right panel of table 11 repeats the exercise using exits. Among observations with positive AI deployment, 83% have some exit. In contrast, among observations with no AI deployment, only 49% have some exit. AI is correlated with the exit of varieties. *Taken together, these two results point to the role of AI deployment for entry and exit i.e., for creative destruction.*

5.4. The Welfare Gains from AI: New Empirics

We saw in equation (4) that the welfare gains from AI deployment can be expressed as

$$\Delta W_t = (\lambda_t / \lambda_t^{\text{no AI}})^{-1/(\sigma-1)}$$

where λ_t equals one minus the share of new Apps in total downloads and $\lambda_t^{\text{no AI}}$ is its counterfactual value in a world with no AI deployment. Before explaining how we estimate ΔW_t we make two observations. The first is the major caveat that we are using download data whereas the welfare calculation should be based on expenditure data.

Table 11: Creative Destruction

Number of Category-Importer-Year Pairs by AI-Deployment Status and					
ENTRY Status			EXIT Status		
	AI Deploy > 0	AI Deploy = 0		AI Deploy > 0	AI Deploy = 0
No entry	1,644 (9%)	20,467 (34%)	No exit	3,072 (17%)	30,887 (51%)
Entry	16,973 (91%)	39,657 (66%)	Exit	15,545 (83%)	29,237 (49%)

Notes: Each observation is an App category (c), a downloading country or importer (m), and a year ($t = 2015, \dots, 2019$). We use a panel of 292 App categories, 84 importers, and 5 years and have 78,741 observations. The numbers are counts of observations. Numbers in parentheses are counts as a percentage of the total observations in the table (78,741).

Second, this is a welfare calculation for the mobile App category. It ignores all other goods.

To estimate ΔW_t , we first need an empirical counterpart to λ_t . From equations (3) and (5), it is natural to equate λ_t with $1 - \theta_{cmt}^n$, that is, with one minus the new downloads share for category- c Apps downloaded by users in country m in year t . Since we do not want to get bogged down in reporting welfare for each App category and importer, we take λ_t to be the download-weighted average of the $1 - \theta_{cmt}^n$.¹⁶ This is the obvious empirical counterpart to λ_t .

Our next challenge is to calculate the counterfactual $\lambda_t^{\text{no AI}}$. To this end, we regress $1 - \theta_{cmt}^n$ on AI deployment $\ln(1 + AI_{cmt})$. We interact this with year dummies so that we can do counterfactuals separately by year. θ_{cmt}^n is not defined for 2014 so we omit the year. Table 12 reports the regressions. The first three columns are OLS. The negative coefficients mean that high AI deployment is associated with low $1 - \theta_{cmt}^n$ and hence with high new-App download shares. This is sensible and expected. Also as expected, the AI deployment coefficient becomes more negative with time: As AI has become more sophisticated, its positive impact on new-App downloads has grown.

In columns 2 and 3 we add our covariates for non-AI patents and multiple-category AI patents.¹⁷ Adding them does not alter the coefficients on our AI-deployment variables.

IV estimates are reported in columns 4–6. We use the same instrument as before, but with one alteration. That instrument was at the cxt level: $\ln(1 + AI_{ct}) \cdot (ConfPaper_{xt})$.

¹⁶ For each t , the cm download weights are the denominator of θ_{cmt}^n , that is, the downloads of all varieties of category- c Apps available to users in country m in year t .

¹⁷ These are the download-weighted average of non-AI patents and the download-weighted average of the AI patents of multi-category firms. See equation (6) for weights.

Table 12: (1 – New Product Share) Regressed on AI Deployment

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2020}$	-2.56*** (0.115)	-2.53*** (0.111)	-2.51*** (0.111)	-2.91*** (0.191)	-2.88*** (0.187)	-2.89*** (0.188)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2019}$	-1.76*** (0.097)	-1.78*** (0.096)	-1.72*** (0.095)	-2.03*** (0.182)	-2.06*** (0.181)	-2.04*** (0.182)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2018}$	-1.60*** (0.105)	-1.54*** (0.106)	-1.48*** (0.105)	-1.88*** (0.177)	-1.83*** (0.177)	-1.81*** (0.177)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2017}$	-0.83*** (0.109)	-0.77*** (0.101)	-0.70*** (0.099)	-0.94*** (0.179)	-0.88*** (0.167)	-0.83*** (0.167)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2016}$	0.17 (0.121)	0.09 (0.120)	0.17 (0.119)	-0.04 (0.201)	-0.11 (0.197)	-0.05 (0.197)
$\ln(1 + AI_{cmt}) \cdot \text{Year 2015}$	0.88*** (0.135)	0.87*** (0.133)	0.95*** (0.133)	-0.95*** (0.309)	-0.93*** (0.302)	-0.90*** (0.302)
$\ln(1 + \text{Multiple } AI_{cmt})$		-1.00* (0.581)	-0.17 (0.448)		-0.97* (0.577)	-0.15 (0.445)
$\ln(1 + \text{non}AI_{cmt})$			-51.43*** (14.812)			-51.29*** (14.783)
Observations	78,614	74,606	74,606	78,614	74,606	74,606
Fixed effects	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>
K-P <i>F</i> -value				215.1	214.9	214.4
R^2	0.335	0.333	0.333			

Notes: The dependent variable is $(1 - \theta_{cmt}^n) \cdot 100$. Each observation is indexed by an App category c , a downloading country or importer m , and a year t . We use a panel of 292 App categories, 84 importers, and 6 years. The independent variable is $\ln(1 + AI_{cmt})$ where AI_{cmt} is defined in equation (6). Columns 1–3 are OLS and columns 4–6 are IV. The first-stage results appear in table A2 in appendix. ‘K-P’ *F*-value is the Kleibergen-Paap weak-instruments test statistic. We include fixed effects for year-importer and for App category. Standard errors are clustered by importer. ***, ** and * indicates statistical significance at the 1%, 5% and 10% levels, respectively.

Table 13: Quantifying Welfare Gains from AI's Creative Destruction

	2020		2015	
	OLS	IV	OLS	IV
1. λ_t	0.878	0.878	0.867	0.867
2. $\Delta\lambda_t$	0.082	0.093	-0.023	0.025
3. $\lambda_t^{\text{no AI}} = \lambda_t + \Delta\lambda_t$	0.960	0.971	0.844	0.892
4. Gains from AI ($\sigma = 5.03$)	1.022	1.025	0.993	1.007
5. Percentage Gains from AI ($\sigma = 5.03$)	2.2%	2.5%	-0.7%	0.7%
6. Gains from AI ($\sigma = 2.00$)	1.093	1.106	0.973	1.029
7. Percentage Gains from AI ($\sigma = 2.00$)	9.3%	10.6%	-2.7%	2.9%

Notes: Row 1 uses observed data. Row 2 is based on estimates from table 12. See the text for an explanation. Row 3 is the sum of rows 1 plus 2. Row 4 is $\Delta W_t = (\lambda_t / \lambda_t^{\text{no AI}})^{-1/(5.03-1)}$. Row 5 is $100 \cdot (\Delta W_t - 1)$. Row 6 is $\Delta W_t = (\lambda_t / \lambda_t^{\text{no AI}})^{-1/(2.00-1)}$. Row 7 is $100 \cdot (\Delta W_t - 1)$.

As in equation (6), we aggregate this to the *cmt* level using importer download weights, that is, $\sum_x w_{cmtx} \cdot \ln(1 + AI_{ct}) \cdot (ConfPaper_{xt})$. We then interact this with year dummies to create six instruments for our six endogenous variables. The first stages appear in appendix table A2 where it is shown that the instruments are highly significant and the first-stage coefficients are sensible. With five instruments we must be especially mindful of the weak-instruments problem; however, our Kleibergen-Paap F -statistics of over 200 are well above the Stock-Yogo threshold of 20. See the second last row of table 12. With instruments in place, columns 4–6 show that the IV estimates are similar to OLS, are statistically significant, are negative in 2020, and decline over time. The exception is 2015.

We can now quantify how AI has influenced the welfare gains from creative destruction. We are interested in $\Delta W_t = (\lambda_t / \lambda_t^{\text{no AI}})^{-1/(\sigma-1)}$. We set σ to 5.03, which is the Head and Mayer (2014) median estimate of σ from their meta-study. However, it seems reasonable given network effects that the elasticity relevant to Apps is closer to unity. If this is the case, we are understating the welfare gains. Table 13 reports calculations of ΔW_t for 2020 and 2015. Consider the first column of numbers. From row 1, one minus the new-App share is 0.878 in 2020. From row 2, we estimate that AI induces λ_t to change by 0.082. This is calculated as follows. From column 1 of table 12, the impact of AI deployment on λ_{2020} is $-2.56/100$. (We divide by 100 because the dependent variable in the table was multiplied by 100.) The weighted average of $\ln(1 + AI_{cmt,2020})$ conditional on $AI_{cmt,2020} > 0$ is 3.19. Hence the estimated change in λ_{2020} from shutting down AI is

$\Delta\lambda_t = (-2.56/100) \cdot (-3.19) = 0.082$. We compute the counterfactual λ_t as $\lambda_t^{\text{no AI}} = \lambda_t + \Delta\lambda_t$ and this appears in row 3. Row 4 reports $\Delta W_t = (\lambda_t / \lambda_t^{\text{no AI}})^{-1/(5.03-1)}$. Row 5 expresses this as the percent $100 \cdot (\Delta W_t - 1)$. Using OLS, we estimate that AI led to welfare gains of 2.2% in 2020. *Our headline number, based on IV, is that AI led to welfare gains of 2.5%*. There is limited evidence on elasticities of substitution between Apps. We therefore also consider a common alternative estimate of σ , namely, $\sigma = 2$. From rows 6–7, this implies that *AI led to welfare gains of 10.6%, a very large number*.

The commercialization of AI is usually dated to 2012 (Agrawal *et al.*, 2018) so that in 2015, the use of AI in mobile Apps was in its early days. We should therefore expect smaller welfare benefits of AI in 2015. This is a bit like a placebo test. In table 13 we repeat the analysis for 2015. The calculations are similar except that now we use the 2015 coefficients from table 12 (0.88 for OLS and -0.95 for IV) and we use the weighted average of $\ln(1 + AI_{cm,2015})$ conditional on $AI_{cm,2015} > 0$, which is 2.64. From table 13 (IV), in 2015 AI led to welfare gains of 0.7%. This is for $\sigma = 5.03$. It is 2.9% for $\sigma = 2.00$. As expected, these are much smaller than the corresponding gains in 2020.

Summarizing, in 2020, AI deployment raised welfare from creative destruction by between 2.5% and 10.6%. Further, in 2015, when AI deployment in mobile Apps was still in its infancy, AI deployment raised welfare from creative destruction by only between 0.7% and 2.9%.

6. Conclusions

Artificial Intelligence is a powerful new technology that will likely have large impacts on the size, direction and composition of international trade flows. Yet almost nothing is known empirically about this process, partly because impacts on goods trade have likely been minimal and partly because researchers have failed to look where the action is most obvious — in the palms of our hands. We observed that mobile Apps provide a large and growing collection of services that billions of people use daily and whose international dimension is captured by mobile App downloads. We developed a new database of App downloads and the AI deployed in those Apps. Using an IV strategy to estimate the impacts of AI deployment we presented three results:

1. *Bilateral Trade*: AI deployment increased App downloads by a factor of six. This is the first systematic evidence of the impact of AI on trade.
2. *Variety Effects*: AI deployment doubled the number of bilaterally traded App varieties.
3. *Entry, Exit, and Creative Destruction*: AI deployment caused high levels of entry into and exit out of Apps/varieties available in the importer country. This has

important welfare implications. Comparing the actual evolution of mobile App downloads to a counterfactual world in which no AI is deployed, AI deployment in 2020 raised welfare from App downloads by between 2.5% (when Apps are highly substitutable) and 10.6% (when Apps are less substitutable).¹⁸

With regards to international App markets, AI deployment has already had tangible impacts on trade, product variety, creative destruction and welfare.

¹⁸An important caveat is that our welfare calculations use download shares rather than expenditure shares.

References

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2020. AI and jobs: Evidence from online vacancies. NBER Working Papers No. 28257, National Bureau of Economic Research,.
- Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2019. Missing growth from creative destruction. *American Economic Review* 109(8):2795–2822.
- Aghion, Philippe, Antonin Bergeaud, Matthieu Lequien, and Marc J. Melitz. 2018. The impact of exports on innovation: Theory and evidence. Working papers 678, Banque de France.
- Aghion, Philippe and Peter Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60(2):323–351.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Cambridge, MA: Harvard Business Review Press.
- Akcigit, Ufuk and William R. Kerr. 2018. Growth through heterogeneous innovations. *Journal of Political Economy* 126(4):1374–1443.
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare. 2012. New Trade Models, Same Old Gains? *American Economic Review* 102(1):94–130.
- Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert J Richmond, and Johannes Stroebel. 2020. International trade and social connectedness. Working paper, National Bureau of Economic Research.
- Beraja, Martin, David Y. Yang, and Noam Yuchtman. 2020. Data-intensive innovation and the state: Evidence from AI firms in China. NBER Working Papers 27723, National Bureau of Economic Research, Inc.
- Broda, Christian and David E. Weinstein. 2006. Globalization and the gains from variety. *Quarterly Journal of Economics* 121(2):541–585.
- Brynjolfsson, Erik, Xiang Hui, and Meng Liu. 2019. Does machine translation affect international trade? Evidence from a large digital platform. *Management Science* 65(12):5449–5460.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108:43–47.
- Card, David. 2001. Estimating the return to schooling: Progress on some persistent econometric problems. *Econometrica* 69(5):1127–1160.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. 2016. Prices, markups, and trade reform. *Econometrica* 84:445–510.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2011. An anatomy of international trade: Evidence from french firms. *Econometrica* 79(5):1453–1498.

- Feenstra, R.C. 2010. *Product Variety and the Gains from International Trade*. MIT Press.
- Feenstra, Robert C. 1994. New product varieties and the measurement of international prices. *American Economic Review* 84(1):157–177.
- Goldfarb, Avi and Daniel Trefler. 2019a. Artificial intelligence and international trade. In Ajay Agrawal, Joshua Gans, and Avi Goldfarb (eds.) *The Economics of Artificial Intelligence: An Agenda*, NBER. University of Chicago Press, 463–492.
- Goldfarb, Avi and Daniel Trefler. 2019b. How Artificial Intelligence Impacts International Trade. In *The World Trade Report 2018: The Future of World Trade: How Digital Technologies are Transforming Global Commerce*. Geneva: WTO, 140.
- Grossman, Gene M. and Elhanan Helpman. 1991a. Quality ladders in the theory of growth. *Review of Economic Studies* 58(1):43–61.
- Grossman, Gene M. and Elhanan Helpman. 1991b. *Innovation and Growth in the Global Economy*. Cambridge, Mass. and London: MIT Press.
- Head, Keith and Thierry Mayer. 2014. Gravity equations: Workhorse, toolkit, and cookbook. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff (eds.) *Handbook of International Trade*. Amsterdam: Elsevier, 131–195.
- Head, Keith, Thierry Mayer, and John Ries. 2010. The erosion of colonial trade linkages after independence. *Journal of International Economics* 81(1):1–14.
- Helpman, Elhanan. 1981. International trade in the presence of product differentiation, economies of scale and monopolistic competition: A chamberlin-heckscher-ohlin approach. *Journal of International Economics* 11(3):305–340.
- Hsieh, Chang-Tai, Nicholas Li, Ralph Ossa, and Mu-Jeung Yang. 2020. Accounting for the new gains from trade liberalization. *Journal of International Economics* 127:1–29.
- Imbens, Guido W. and Joshua D. Angrist. 1994. Identification and estimation of local average treatment effects. *Econometrica* 62(2):467–475.
- Irwin, Douglas A and Peter J Klenow. 1994. Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy* 102(6):1200–1227.
- Krugman, Paul R. 1979. Increasing returns, monopolistic competition, and international trade. *Journal of International Economics* 9(4):469–479.
- Lee, Kai-Fu. 2018. *AI Super-Powers: China, Silicon Valley, and the New World Order*. New York: Houghton Mifflin Harcourt.
- Lileeva, Alla and Daniel Trefler. 2010. Improved access to foreign markets raises plant-level productivity ... for some plants. *Quarterly Journal of Economics* CXXV(3):1051–1100.
- Lim, Kevin, Daniel Trefler, and Miaojie Yu. 2019. Trade and innovation: The role of scale and competition effects. Working paper, University of Toronto.

- McKinsey Global Institute. 2019. Globalization in transition: The future of trade and value chains. <https://www.mckinsey.com/~media/mckinsey/featured%20insights/innovation/globalization%20in%20transition%20the%20future%20of%20trade%20and%20value%20chains/mgi-globalization%20in%20transition-the-future-of-trade-and-value-chains-full-report.pdf>. Accessed: 2020-09-20.
- Melitz, Marc J. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6):1695–1725.
- Melitz, Marc J. and Daniel Trefler. 2012. Gains from trade when firms matter. *Journal of Economic Perspectives* 26(2):91–118.
- Rivera-Batiz, Luis A. and Paul M. Romer. 1991. Economic integration and endogenous growth. *Quarterly Journal of Economics* 106(2):531–555.
- Romalis, John. 2004. Factor proportions and the structure of commodity trade. *American Economic Review* 94(1):67–97.
- Royal Society-National Academy of Sciences. 2019. *Artificial Intelligence: An International Dialogue*. The Royal Society.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism, and Democracy*. New York; London: Harper & Brothers.
- Stapleton, Katherine and Michael Webb. 2020. Automation, trade and multinational activity: Micro evidence from Spain. Working Paper, Oxford and Stanford Universities.
- Trefler, Daniel. 2004. The long and short of the Canada-U.S. free trade agreement. *American Economic Review* 94(4):870–895.
- WIPO. 2018. Data collection method and clustering scheme. Background paper, WIPO.
- WIPO. 2019. *WIPO Technology Trends 2019: Artificial Intelligence*. Geneva: World Intellectual Property Organization.
- Zhang, Daniel, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, et al. 2021. The AI index 2021 annual report. *arXiv preprint arXiv:2103.06312* .

Appendix Tables

Table A1: First Stage for Table 8

	$\ln(1 + AI_{cxt})$ • First positive tercile	$\ln(1 + AI_{cxt})$ • Second positive tercile	$\ln(1 + AI_{cxt})$ • Third positive tercile
Instrument • First positive tercile	0.810*** (0.0506)	-0.036*** (0.00964)	-0.023*** (0.00624)
Instrument • Second positive tercile	-0.053*** (0.0155)	0.911*** (0.101)	-0.030*** (0.00426)
Instrument • Third positive tercile	-0.067*** (0.0249)	-0.049*** (0.0175)	0.898*** (0.0466)
<i>F-value</i>	807	385	596
Fixed effects	<i>t-m-x, c</i>	<i>t-m-x, c</i>	<i>t-m-x, c</i>

Notes: This table displays the first stage for our preferred specification in table 8 (column 2). The dependent variable in the first stage is $\ln(1 + AI_{cxt})$. Its instrument is $\ln(1 + AI_{ct}) \cdot (ConfPaper_{xt})$. Both are interacted with terciles dummies of the distribution of AI_{cxt} conditional on $AI_{cxt} > 0$. There are thus three independent variables and hence three first stages or columns. We include fixed effects for year-importer-exporter and category. Standard errors are clustered by importer. ***, ** and * indicates statistical significance at the 1%, 5% and 10% levels, respectively.

Table A2: First Stage for Table 12

	$\ln(1 + AI_{cmt})$ • Year 2020	$\ln(1 + AI_{cmt})$ • Year 2019	$\ln(1 + AI_{cmt})$ • Year 2018	$\ln(1 + AI_{cmt})$ • Year 2017	$\ln(1 + AI_{cmt})$ • Year 2016	$\ln(1 + AI_{cmt})$ • Year 2015
Instrument • Year 2020	10.25*** (0.289)	-1.040*** (0.0561)	-0.818*** (0.0477)	-1.066*** (0.0344)	-1.162*** (0.0399)	-1.468*** (0.0593)
Instrument • Year 2019	-1.461*** (0.0666)	10.37*** (0.312)	-0.700*** (0.0471)	-0.993*** (0.0360)	-0.990*** (0.0416)	-1.296*** (0.0495)
Instrument • Year 2018	-2.204*** (0.0896)	-1.417*** (0.0776)	14.93*** (0.351)	-1.431*** (0.0510)	-1.424*** (0.0581)	-1.905*** (0.0608)
Instrument • Year 2017	-2.207*** (0.0955)	-1.468*** (0.0846)	-1.082*** (0.0741)	16.39*** (0.377)	-1.670*** (0.0652)	-2.048*** (0.0729)
Instrument • Year 2016	-2.634*** (0.120)	-1.550*** (0.0919)	-1.111*** (0.0869)	-1.737*** (0.0681)	15.80*** (0.543)	-2.451*** (0.106)
Instrument • Year 2015	-3.136*** (0.154)	-1.785*** (0.103)	-1.642*** (0.101)	-2.831*** (0.140)	-2.831*** (0.159)	19.95*** (0.778)
<i>F-value</i>	403	387	783	785	630	535
Fixed effects	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>	<i>t-m, c</i>

Notes: This table displays the first stages for our preferred specification in table 12 (column 4). The dependent variable in the first stage is $\ln(1 + AI_{cmt})$ defined in equation (6). Its instrument is $\sum_x w_{cmt} \cdot \ln(1 + AI_{ct}) \cdot (ConfPaper_{xt})$. Both are interacted with year dummies. With six years there are six dependent variables and hence six first stages or columns. We include fixed effects for year-importer and category. Standard errors are clustered by importer. ***, ** and * indicates statistical significance at the 1%, 5% and 10% levels, respectively.