

NBER WORKING PAPER SERIES

THE IMPACT OF E-COMMERCE ON RELATIVE PRICES AND CONSUMER WELFARE

Yoon J. Jo
Misaki Matsumura
David E. Weinstein

Working Paper 26506
<http://www.nber.org/papers/w26506>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2019

We are especially grateful to the Rakuten Institute of Technology for providing us with data and assistance for this project. We also want to thank Mary Amiti, Kirill Borusyak, Yu Hirate, David Nagy, Amine Ouazad, Roberto Rigobon, Lorenzo Rotunno, Mototsugu Shintani, Takashi Umeda, Takashi Unayama, and Tsutomu Watanabe, Eugene Wong for many useful comments on early drafts. We also thank Yingjie (Angela) Wu for providing excellent research assistance. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Yoon J. Jo, Misaki Matsumura, and David E. Weinstein. 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.

The Impact of E-Commerce on Relative Prices and Consumer Welfare
Yoon J. Jo, Misaki Matsumura, and David E. Weinstein
NBER Working Paper No. 26506
November 2019
JEL No. F14,L86,R32

ABSTRACT

This paper examines the impact of e-commerce on pricing behavior and welfare. Using Japanese data, we find that the entry of e-commerce firms significantly raised the rate of intercity price convergence for goods sold intensively online, but not for other goods. E-commerce also lowered relative inflation rates for goods sold intensively online. We overcome data challenges using long data series and historical catalog sales as an instrument for e-commerce sales intensity. We estimate that reductions in price dispersion raised welfare by 0.3 percent. E-commerce also lowered variety-adjusted prices on average by 0.9 percent, and more in cities with highly educated populations.

Yoon J. Jo
Department of Economics
Texas A&M University
4228 TAMU
College Station, TX 77843
yoonjo@tamu.edu

Misaki Matsumura
Columbia University
1201 N. Fillmore St.
Arlington, VA 22201
mm4387@columbia.edu

David E. Weinstein
Columbia University, Department of Economics
420 W. 118th Street
MC 3308
New York, NY 10027
and NBER
dew35@columbia.edu

1 Introduction

How has e-commerce affected prices and consumer welfare? Answering this question is difficult because researchers often face three data challenges. First, researchers typically do not have access to data on what share of a consumer's purchases of any particular good are made online and offline, making it difficult to assess how important online sales are for most goods. Second, we lack long time series, which hinders our ability to know whether price dynamics associated with e-commerce were present even before e-merchants entered the market. Finally, we lack good instruments for e-commerce sales intensity that can control for endogeneity biases that might stem from a correlation between e-commerce entry and market conditions.

This paper makes use of a unique Japanese data set containing household expenditure shares by retail channel—catalog, e-commerce, and physical store—for the set of goods that make up the Japanese consumer price index (CPI) which enables us to overcome these empirical challenges. The long time series of our data (1991 to 2016) enables us to control for important pre-trends in the pricing dynamics of the types of goods available from online merchants. Finally, the fact that the data contains information on catalog expenditures by product for the period before the advent of e-commerce provides us with a powerful instrument to correct for measurement errors and an endogeneity bias.

Consistent with [Goolsbee and Klenow \(2018\)](#), we find that goods sell intensively online have significantly lower relative rates of price increase. While one might be tempted to attribute this result to the impact of online merchants, we exploit the long time series in our data to show that this pattern was also true before e-commerce firms entered the Japanese market. Thus, the differential pattern in pricing behavior seems to be a characteristic of the types of goods amenable to online sales rather than a feature of e-commerce *per se*. Nevertheless, we document that after the entry of e-commerce merchants, the difference in the rate of price increase between goods not sold intensively online and those sold intensively rose,

suggesting that e-commerce lowered the relative rate of price increase of e-commerce intensive goods.

We also document that e-commerce had important impacts on the convergence rates of intercity price differentials at physical stores. We find that prior to e-commerce, the rate at which the retail price of a good converged across cities was similar regardless of whether that good would ultimately be sold intensively online or not. However, after the advent of the e-commerce, we find that the rate of price convergence for goods sold online rose significantly, but this was not so for other goods. Based on our estimates of how e-commerce differentially affected the ability of merchants to price discriminate across cities, we compute the impact of e-commerce on Japanese welfare using the model developed in [Jensen \(2007\)](#). We estimate the welfare gains due to e-commerce to be 0.3 percent of consumption expenditure in 2014 in our baseline specification.

In addition to the gains arising from improved price arbitrage across cities, we also compute the consumer gains due to new varieties. We model goods purchased online as new varieties either because online purchases enable consumers to buy goods that were not available locally before or because an online shopping experience differs in important ways (convenience, service, etc.) from an offline shopping experience. As [Brynjolfsson et al. \(2003\)](#) and [Dolfen et al. \(2019\)](#) have argued, these variety channels are likely to be quite important, and we follow them by modeling e-commerce as a new method of conducting retail purchases. We use the approach developed in [Feenstra \(1994\)](#) to compute the gains due to varieties. Using this framework, we estimate that e-commerce lowered the price index faced by consumers by 0.9 percentage points. Interestingly, we also find evidence of a digital divide in terms of which consumers benefit the most from new varieties. Since e-commerce expenditure shares are strongly correlated with college education in Japan, these gains accrued far more in cities with populations with a high share of college graduates like Tokyo than in cities with low levels of college education.

1.1 Related Literature

Our results are related to a number of papers examining how information technology has affected pricing and welfare. Following [Cavallo \(2018\)](#), we argue that e-commerce is a technology that promotes uniform pricing across locations. However, our work differs in our focus on convergence rates and the resulting welfare implications. Our work is also related to the large literature that has demonstrated that information technology serves to reduce price dispersion and promote trade. [Freund and Weinhold \(2004\)](#) show that countries with more web hosts export more to each other. [Jensen \(2007\)](#), [Aker \(2010\)](#), and [Allen \(2014\)](#) examine the impact of the introduction of mobile phones on fish or agricultural markets in India, Niger, and the Philippines, and [Steinwender \(2018\)](#) examines the impact of the transatlantic telegraph cables on 19th century textile prices and exports. Our work is complementary to these papers in that we show that e-retail also serves to reduce price dispersion across locations. However, we differ in focus and scope—our study examines the role played by e-commerce in an advanced, modern economy on the prices of hundreds of goods in physical retailers. Our paper also relates to the literature on internet pricing. In particular, [Cavallo \(2017\)](#) shows that online prices and prices in physical stores are quite similar.¹ This fact helps motivate our assumption that local retailers with high prices should face stiff competition from online retailers.

Our work is also related to studies of the impact of e-commerce on welfare. Much of this research has focused on the gains from variety that arise as consumers can purchase products that are not available in local stores. For example, [Brynjolfsson et al. \(2003\)](#) compute the variety gains from internet book sales, [Fan et al. \(2018\)](#) examine the relative variety gains in large and small Chinese cities associated with internet usage; and [Dolfen et al. \(2019\)](#) estimate the gains from e-retail due to shopping convenience and

¹[Cavallo \(2017\)](#) found that the ratio of online to offline prices in Japan tended to be lower than in other countries, but this may reflect the fact that all but one of his offline stores in Japan were electronics stores whereas a much broader sample of stores was used in countries like the US.

new varieties in the U.S. An important difference between these studies and ours is our focus on price convergence across locations and the fact that we make use of representative surveys in order to measure internet sales intensities.

Other papers have examined how the internet affects local markets. [Goldmanis et al. \(2010\)](#) examine regional patterns in online purchase behavior change the market structure in bookstores, travel agencies and car dealers. [Goyal \(2010\)](#) finds that the introduction of internet kiosks raised soy prices in rural India. [Couture et al. \(2018\)](#) conduct a randomized control trial in eight rural Chinese counties and find little effect of the introduction of e-commerce on the local economy. [Brown and Goolsbee \(2002\)](#) show that the creation of online insurance sales systems reduced the variance of insurance pricing. Our work differs from these studies in terms of scope (the large number of different sectors considered), the link to physical retail prices across an entire economy, and identification strategy (the ability to examine differential rates of price convergence before and after the advent of e-commerce).

Finally, our paper is also related to the large literature on purchasing power parity convergence regressions. [Parsley and Wei \(1996\)](#) were the first to document that differences in convergence coefficients across cities were linked to trade costs, an insight that we build upon in this paper. We estimate that intercity convergence rates for Japan are higher than those obtained in [Parsley and Wei \(1996\)](#) and [Cecchetti et al. \(2002\)](#). These studies found no price convergence across U.S. cities once one controlled for city fixed effects. In contrast, we find that the advent of e-commerce is associated with a significant rise in the rate of price convergence for goods with high internet sales intensity after the advent of e-commerce. Our ability to better detect intercity price convergence probably arises from the fact that Japanese CPI data is based on the sampling of identical or extremely similar goods across cities, whereas U.S. price data is based on similar but non-identical sets of goods across cities. Our approach also builds off [Bergin et al. \(2017\)](#), who employ a similar triple difference strategy to show that

rates of price convergence across European countries increased after joining the euro area. It is also related to [Crucini et al. \(2015\)](#) who examine how distance affects the rate of price convergence across cities.

The remainder of the paper is organized as follows. Section 2 introduces the estimation strategy and provides the theory for the welfare calculation. Section 3 presents the data and provides some stylized facts about e-commerce suitability. Section 4.1 presents our results on national prices. We present our main estimates for the impact of e-commerce on price convergence and welfare in Section 4.2. Section 4.3 presents consumer gains in the new trade theory models, and Section 5 concludes.

2 Theory

In Section 2.1, we model the impact that e-commerce has had on interregional price differentials and show how the decline in these differentials raises welfare in Section 2.2.

2.1 Estimating the Impact of the E-Retail on Price Arbitrage

We begin by defining some notation. Let $p_{ict} \equiv \ln P_{ict}$ be the log price of item i in city c in time t . Define the Δ^k operator as $\Delta^k p_{ict} \equiv p_{ict} - p_{ic,t-k}$; thus, if we set $k = 1$, we can examine annual changes, but we can also examine longer differences by setting k equal to a whole number larger than one. Let $x_{ib}^E \in [0, 1]$ be the “e-commerce sales intensity” of a good measured in a survey year b , where zero indicates it is not sold online and one indicates that its share of online sales relative to offline sales is the highest among all goods. Let D_t be an indicator variable that is one if e-commerce is an option in period t and zero otherwise. We assume that the change in the price of any item in a city c can be written as a standard purchasing price parity specification in which the rate of price convergence depends on their availability online, i.e.,

$$\begin{aligned}\Delta^k p_{ict} &= \alpha_{it} + \beta_{ct} + \left(\gamma + \delta_1 x_{ib}^E + \delta_2 D_t + \delta_3 D_t x_{ib}^E \right) p_{ic,t-k} + \epsilon_{ict} \\ &= \alpha_{it} + \beta_{ct} + \gamma_i^E(D_t) p_{ic,t-k} + \epsilon_{ict},\end{aligned}\tag{1}$$

where $\gamma_i^E(D_t) \equiv \left(\gamma + \delta_1 x_{ib}^E + \delta_2 D_t + \delta_3 D_t x_{ib}^E \right)$. In this specification, α_{it} is an item-time fixed effect; β_{ct} is a city-time fixed effect; γ is a parameter that captures the rate of intercity price convergence for goods not available online; δ_1 is a parameter that captures the rate of price convergence for goods available online prior to the entry of e-commerce firms; δ_2 captures whether there was a change in overall convergence rates after the entry of e-commerce firms; δ_3 captures the increase in rate of price convergence for online goods after the entry of e-commerce firms; and ϵ_{ict} is an iid error term. We think of this error as a price shock arising from period t local supply-and-demand conditions for an item in a city that is not common to all items in the city and is uncorrelated with past prices.


If $\gamma_i^E(D_t) = -1$, equation (1) collapses to $p_{ict} = \alpha_{it} + \beta_{ct} + \epsilon_{ict}$, and therefore the price of any item can be decomposed into its national price (α_{it}), a common local market premium (β_{ct}), and an iid error term that is not persistent. In this case, any idiosyncratic price shock to a good in a city (ϵ_{ict}) has no impact on prices in the next period. Hence, price convergence occurs in one period, and prices always equal their conditional mean of $(\alpha_{it} + \beta_{ct})$ plus a random iid shock. At the other extreme, we have the case of where $\gamma_i^E(D_t) = 0$, which implies that prices follow a random walk with a drift term given by $(\alpha_{it} + \beta_{ct})$. In intermediate cases where $\gamma_i^E(D_t) \in (-1, 0)$, prices converge across cities by price differences across cities can persist for more than k years.

2.2 Welfare in Partial Equilibrium

We can map the price changes into welfare gains by using the framework developed in [Jensen \(2007\)](#). Jensen considered a technological change that enabled arbitrage between a high-priced region (H) and a low-priced region

(L). If e-commerce reduces price dispersion, we should expect the price in region H to fall and the price in L to rise as shown in Figure 1. Consumers in H will gain $(A + B)$, and sellers will gain $(C - A)$, yielding a net gain of $(B + C)$. Similarly, in region L , consumers will *lose* $(D + E)$ and sellers will gain $(D - F)$, yielding a net loss of $(E + F)$. Overall, the welfare gain is $(B + C) - (E + F)$, which will necessarily be positive in the case of linear demands with equal slopes as long as the final price in H is at least as large as the price in the region L (i.e., $P(Q_H + \Delta Q) \geq P(Q_L - \Delta Q)$). One can also see this condition holds in the figure because both trapezoids $(B + C)$ and $(E + F)$ have identical bases and differ only in the lengths of their parallel sides.

Figure 1: Welfare Gains from Arbitrage in the Jensen Model



Jensen (2007) considered a case in which the marginal cost of supplying a good is zero, which enabled him to compute the lengths of the parallel sides of the quasi-trapezoids by just using the prices. When thinking about production more generally, however, marginal costs are likely to be positive, so technically we should subtract marginal costs from prices when computing the lengths of the parallel sides of the quasi-trapezoids. However, as

one can see from Figure 1, if we assume constant and equal marginal cost of production for a given good, then $G = G'$, and we can still compute the welfare gain as $(B + C + G) - (E + F + G') = (B + C) - (E + F)$.²

We can use our estimates of the impact of e-commerce on price convergence to calibrate the Jensen model. In order to compute the partial equilibrium welfare gain due to e-commerce, we consider the difference in implied gains in two counterfactual cases. In each case, we assume that the economy has deviated from a steady-state equilibrium but experiences different rates of price convergence. The first case corresponds to one in which consumers do not have access to e-commerce, so price convergence is slow, which results in small movements towards a common price. The second case is one in which consumers do have access, which results in greater price convergence and therefore greater welfare gains. We then set the welfare gain associated with e-commerce to be equal to the difference in the welfare gains arising from the different convergence rates.

Let D denote each counterfactual case, where $D = 0$ corresponds to a counterfactual with no e-commerce, and $D = 1$ corresponds to a counterfactual with e-commerce. We will consider a counterfactual in which observed prices constitute deviations from a free-trade steady state in which goods prices are the same in all cities $p_{ic}^* = p_i^*$. One of the features of thinking about convergence to this free-trade steady state is that we can show that there will always exist a p_i^* such that price convergence towards that level will result in no net increase in demand, as in Jensen (2007). A necessary condition for the existence of steady-state prices (p_{ic}^*) is

$$\Delta p_{ic}^* = 0 = \alpha_{it} + \beta_{ct} + \gamma_i^E(D_t) p_i^*. \quad (2)$$

By definition, the common steady-state prices (p_i^*) do not vary with t . Similarly, $\gamma_i^E(D_t)$ will not vary with t as long as we remain in a counterfactual world in which D_t (which we now refer to as D) always equals zero

²The assumption of equal marginal costs is probably not extreme for Japan given the small physical size of the country (most major cities are within a few hours drive of Tokyo), which means that transport costs are unlikely to produce large price differences across cities.

or one. Therefore, it must be the case that in the free-trade steady state, $\alpha_{it} + \beta_{ct} = \alpha_i^*$. This condition intuitively implies that in the steady state α_{it} and β_{ct} cannot vary with time, so $\alpha_{it} = \alpha_i$ and $\beta_{ct} = \beta_{c't} = \beta^*$. Thus, we can rewrite this equation as

$$0 = \alpha_i^* + \beta^* + \gamma_i^E(D) p_i^*. \quad (3)$$

Suppose that we start in this steady state and perturb steady-state prices (p_i^*) in period $T-1$ to some other values, $p_{ic,T-1}$. In period T , we will observe prices change by $\widehat{\Delta p_{icT}}(D) = \alpha_i^* + \beta^* + \gamma_i^E(D) p_{ic,T-1}$. If we subtract equation (3) from this equation, we obtain

$$\widehat{\Delta p_{icT}}(D) = \gamma_i^E(D) (p_{ic,T-1} - p_i^*). \quad (4)$$

Using the log change as an approximation for the percentage change, the price level in period T , $P_{icT}(D)$, can be written as $P_{icT}(D) = P_{ic,T-1} [1 + \widehat{\Delta p_{icT}}(D)]$. If we assume a constant elasticity of substitution (CES) demand system, the demand for good i in city c in time t in counterfactual D is given by

$$Q_{ict}(D) = \frac{(P_{ict}(D) / \varphi_{ic})^{-\sigma}}{[P_{ct}(D)]^{1-\sigma}} E_c, \text{ where } P_{ct}(D) \equiv \left[\sum_i \left[\frac{P_{ict}(D)}{\varphi_{ic}} \right]^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (5)$$

where σ is the demand elasticity. It will also be useful to denote the log change in the city price index by $\Delta p_{ct} = \ln [P_{ct}(D) / P_{c,t-1}]$.

Following Jensen, we consider a set of price changes that are consistent with equation (4) and imply no aggregate quantity changes, i.e., $\Delta Q_{iT}(D) = \sum_c \Delta Q_{icT}(D) = 0$. In other words, just as Jensen made the partial equilibrium assumption that eliminating deviations in fish prices did not raise aggregate consumption of fish, we are assuming that eliminating regional price variation in e-commerce intensive products (e.g., books) did not increase aggregate consumption of books. Equation (5) implies that

$$\Delta q_{icT}(D) = (\sigma - 1) \Delta p_{cT}(D) - \sigma \widehat{\Delta p_{icT}}(D) + \Delta \ln E_{cT},$$

where $\ln Q_{ict} = q_{ict}$. If we make the partial equilibrium assumption that

aggregate prices and urban expenditures are unchanged, this equation reduces to

$$\Delta q_{icT}(D) = -\sigma \widehat{\Delta p_{icT}}(D). \quad (6)$$

Using the log change as an approximation for the percentage change, the counterfactual change in urban consumption in any city c in time T can be written as $\Delta Q_{icT}(D) = Q_{ic,T-1} \Delta q_{icT}(D)$. Substituting equation (6) into this equation and summing produces

$$\sum_c Q_{ic,T-1} \widehat{\Delta p_{icT}}(D) = 0. \quad (7)$$

If we then substitute equation (4) into equation (7), we obtain

$$\sum_c Q_{ic,T-1} \gamma_i^E(D) [p_{ic,T-1} - p_i^*] = 0,$$

which means that the steady-state price for each good is given by

$$p_i^* = \frac{\sum_c Q_{ic,T-1} \gamma_i^E(D) p_{ic,T-1}}{\sum_c Q_{ic,T-1} \gamma_{ii}^E(D)}. \quad (8)$$

The welfare gain arising from prices moving from their values in $T-1$ towards their steady state values (P_i^*) can be written as

$$\Delta W_{icT} = \frac{1}{2} (2P_{icT-1} + \Delta P_{icT}(D)) \Delta Q_{icT}(D) - m_i \Delta Q_{icT}(D),$$

where m_i is the marginal cost of producing the good. This welfare gain will be positive whenever prices are higher than their steady-state levels and negative otherwise.

The national welfare change associated with a price convergence toward steady state price levels equals

$$\Delta W_{iT}(D) = \frac{1}{2} \sum_c (2P_{icT-1} + \Delta P_{icT}(D)) \Delta Q_{icT}(D) - \underbrace{m_i \sum_c \Delta Q_{icT}(D)}_{=0}.$$

The welfare gain in the economy expressed as a share of national expenditures, E_T , is

$$\Delta W_T(D) = \sum_i \frac{\Delta W_{iT}(D)}{E_{T-1}(D)}.$$

The gain due to e-commerce is simply the difference between the welfare gain obtained from price arbitrage in the e-commerce regime less that obtained in the regime without e-commerce or

$$\Delta W_T = \Delta W_T(D = 1) - \Delta W_T(D = 0). \quad (9)$$

3 Data

Japanese data is useful for testing the theory because e-commerce expanded rapidly in Japan, which allows us to break the time series into pre- and post-e-commerce periods. Moreover, the share of e-commerce transactions in Japan is similar to those in the U.S.³ By April of 2000, when Japan's largest e-commerce firm, Rakuten, announced its initial public offering, it had grown to be a platform in which consumers had access to goods available from 2,300 merchants, and the Rakuten website was getting 95 million hits per month—almost one hit for every man, woman, and child in Japan.⁴ The following year sales on the Rakuten platform exceeded ¥52 billion (about \$430 million). Thus, within five years, Japanese consumers in any city went from only being able to buy locally or from catalogs to being able to purchase goods from thousands of merchants located across Japan. Rakuten's growth was part of a broader e-commerce boom in Japan. Amazon entered the Japanese market in 2001, with Rakuten still accounting for

³Japan's Ministry of Economy Trade and Industry reports that 5.8% of all retail transactions were done online in 2017, roughly equalling the U.S. level in 2014. See http://www.meti.go.jp/english/press/2018/0425_002.html and https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf. While we do not know when each city obtained broadband internet, we do know that every prefecture in Japan had digital subscriber lines (DSL) and cable TV connections available by 2001, and about a quarter of all households were subscribing to DSL by the end of 2004. In addition, many consumers also used cell phones to make purchases. Source: Ministry of Public Management, Home Affairs, Posts and Telecommunications (<https://www.stat.go.jp/english/data/handbook/pdf/2018all.pdf#page=23>). We thank Takashi Unayama for providing us with these numbers.

⁴Phred Dvorak, "Japan's Highly Popular Rakuten Plans IPO Despite Shaky Market," *Wall Street Journal*, April 18, 2000.

30% of Japanese e-commerce transactions in 2010.⁵ By 2017, e-commerce firms accounted for 5.8 percent of Japanese retail sales or about ¥16.5 trillion (about \$149 billion).

A major advantage of Japanese data is that one can obtain measures of consumer expenditures by product and retail outlet. The National Survey of Family Income and Expenditures (NSFIE) is a representative survey of households with two or more members that reports consumer expenditures by product and retail outlet type.⁶ Starting in 2004, the NSFIE also began a quinquennial recording the expenditure share of each product from online merchants.

One of the problems with the NSFIE data is that it tends to under-report aggregate internet sales because mixed-method retailing (e.g., seeing a product in a store and buying it online tends to be placed in the “other” category). Fortunately, the Ministry of Economy Trade and Industry (METI) reports very reliable aggregate estimates of sales by e-commerce and other retailers by surveying actual sales to consumers by retail merchants. We therefore scale the NSFIE data by the ratio of aggregate sales in the METI data relative to the NSFIE data in order to obtain the same value for aggregate e-commerce sales in the two datasets. In order to make sure that sampling problems are not driving our results, we will use catalog sales as an instrument to deal with classical measurement error and also conduct a robustness check for all of our main results using data on Rakuten sales by product category.

We construct a measure of the e-commerce sales intensity for each of the CPI expenditure categories by computing the share of online expenditures in total expenditures. Let e_{ib} denote the average household total expenditures in category i in survey year b ($b \in \{1999, 2004, 2009, 2014\}$). We denote expenditures in category i from retail channel r in survey year b by s_{ib}^r , where r can take on three values (E , C , and R): the first two are based on

⁵Rakuten, Inc. (2010) *Annual Report*.

⁶The retail outlet types are small retail, supermarket, convenience, department, club, discount, catalog, internet, and “other”.

expenditure data from the NSFIE survey of e-commerce and catalog expenditures and the last is based on Rakuten's reported e-commerce sales. We then define the retail channel "intensity" x_{ib}^r of category i by dividing the expenditures through retail channel r by total expenditures (e_{ib}), normalized by the maximum value of this ratio, i.e.,

$$x_{ib}^r = \frac{s_{ib}^r}{e_{ib}} / \max_j \left(\frac{s_{jb}^r}{e_{jb}} \right). \quad (10)$$

Thus, our measure of the expenditure intensity of retail channel r equals zero if there are no expenditures on goods via a retail channel r in expenditure category i and a value of 1 if the expenditure through retail channel r relative to that in the economy is the highest among all categories of goods sold through that retail channel. Expressing retail channel intensity this way makes our intensity variable (x_{ib}^r) invariant to the size of sector i .

Table 1 makes clear some basic stylized facts of our data. In order to see how e-commerce intensity varies across products, we aggregate the FIES product categories into broader ones so that we can display the data in a compact form. The rows are ordered by a category's share of Japanese expenditures on goods. The first column of Table 1 reports the percentage of expenditures in category ℓ among goods in 2009 as reported in the FIES ($E_{\ell b} \equiv \sum_{i \in \Omega^\ell} e_{ib} / \sum_j e_{jb} \times 100$), where Ω^ℓ is the set of items in some more aggregated category ℓ . In the second column, we report the percentage of online expenditure in 2009 that corresponds to that category ($S_{\ell b}^E \equiv \sum_{i \in \Omega^\ell} s_{ib}^E / \sum_j s_{jb}^E \times 100$), where s_{ib}^E is online expenditure from NSFIE). Columns 3-5 report the e-commerce intensity (i.e., $x_{\ell b}^E \equiv S_{\ell b}^E / E_{\ell b} / \left[\max_k \left\{ S_{kb}^E / E_{kb} \right\} \right]$).

There are three important takeaways from this table. First, there is enormous variation in the e-commerce intensity. Some of this reflects the fact that highly perishable, non-standardized items (e.g. fresh foods), and restricted/time-sensitive items (e.g., medicine and physical newspapers) are not sold much online. At the other end of the spectrum, we see that more standardized goods—e.g., electronics, books, clothing, footwear, and

Table 1: E-Commerce intensity of consumer expenditure on goods

Category	Share Expenditure 2009		E-Commerce Intensity				Catalog Intensity	
	Total	E-Commerce	2004	2009	2014	Rakuten 2010	1999	1999
Fruits and vegetables	10.24	1.76	0.01	0.03	0.05	0.03		0.06
Household consumables	10.19	18.00	0.15	0.36	0.28	0.58		0.34
Clothing	9.61	13.45	0.11	0.28	0.22	0.42		0.41
Store-bought cooked food	7.62	1.10	0.03	0.03	0.04	0.03		0.04
Cereal	6.21	1.54	0.02	0.05	0.05	0.07		0.06
Fish and shellfish	6.13	1.40	0.02	0.05	0.05	0.04		0.05
Cakes and candies	5.72	1.62	0.03	0.06	0.04	0.08		0.05
Meat	5.55	0.73	0.01	0.03	0.04	0.03		0.02
Recreational goods	4.65	12.71	0.30	0.55	0.47	0.93		0.22
Household appliances	4.05	6.32	0.21	0.31	0.36	0.35		0.17
Electronics	3.88	19.32	1.00	1.00	1.00	0.53		0.41
Alcoholic beverages	3.36	1.32	0.05	0.08	0.10	0.26		0.06
Medicine and nutritional supplements	3.35	4.85	0.23	0.29	0.31	0.23		1.00
Non-alcoholic beverages	3.17	2.20	0.09	0.14	0.15	0.16		0.27
Oils, fats and seasonings	3.11	0.73	0.02	0.05	0.07	0.05		0.09
Newspapers and magazines	2.96	0.00	0.00	0.00	0.00	0.00		0.00
Dairy products and eggs	2.81	0.29	0.01	0.02	0.04	0.01		0.02
Transportation equipment	2.14	3.01	0.23	0.28	0.18	0.58		0.40
Domestic utensils	2.06	4.04	0.14	0.39	0.49	0.53		0.41
Furniture and furnishings	1.78	3.45	0.33	0.39	0.51	1.00		0.56
Footwear	1.40	2.13	0.14	0.30	0.28	0.92		0.33
Total/Mean	100.00	100.00	0.15	0.22	0.23	0.32		0.24

Note: Shares are expressed as percentages. This table shows the share of consumption expenditure, e-commerce expenditure, and e-commerce sales intensity, and catalog intensity for goods. E-Commerce intensity and catalog intensity are defined in equation (10).

furniture and furnishings—are sold very intensively online. Interestingly, we see that domestic utensils, household consumables (which include non-durable household supplies like paper products and cleaning agents), and recreational goods (which include items like sports equipment and gardening supplies) are sold very intensively online as well. Second, if we compare the values for 2004, 2009, and 2014, we see a lot of persistence in what is sold online. Newspapers, meat, dairy products, and fruit and vegetables were not sold much online in 2004 and have low e-commerce sales intensity in all subsequent years. Similarly, electronics has the highest internet sales intensity in all of the survey years. Third, as one can see in the second-to-last column, Rakuten sales intensity is highly correlated with e-commerce sales in the NSFIE data ($\rho = 0.57$), which suggests that these datasets are in broad agreement as to what goods are sold intensively online.

One of the other striking features of the table is that if we compare catalog sales intensity (column 7) with e-commerce intensity (columns 3-6), we see that there is a lot of similarity between goods that are sold intensively online and goods that were sold intensively by catalogs in 1999. In that year, e-commerce firms in Japan were still in their infancy: Amazon had not yet entered the Japanese market and Rakuten only had 5.5 million dollars worth of sales on its platform (Olson (2012)). Thus, we can be fairly confident that Japanese catalog sales were probably not much influenced by e-commerce sales. Nevertheless, it is clear that goods sold intensively online tend to have characteristics that are similar to those goods historically available in catalogs, which will motivate our instrument.

In addition to the retail sales data that we have been discussing, we also make use of the fact that the Japan Statistical Bureau (JSB), which produces the Japanese CPI, provides detailed information on representative prices. These prices are sampled in 168 cities on average, which gives us the ability to not only track product prices across time but also across space. This information identifies the brand of an item or provides a detailed description (e.g., “Big-eyed tuna, sliced (for sashimi), lean, 100g”). Since the objective of the JSB sampling is to make meaningful intercity price comparisons, the

JSB tries to select products available in all cities. Finally, we use the official quality-adjusted price quotes for Tokyo computed by the JSB in order to adjust the prices in other cities when goods are substituted into and out of the sample.⁷

One potential concern with these data is that the JSB may not be sampling the same goods in different cities, so unlike barcode data in which goods are precisely defined, some of the price variation in our sample may be capturing unmeasured quality variation instead. Since [Hottman, Redding, and Weinstein \(2016\)](#) show that the correlation between price and quality in barcode data is 0.9, we should expect sampling problems to produce greater levels of price dispersion in JSB data because it would be capturing both price and quality variation. In order to check for this, we compute the log relative price of each good in each city ($\tilde{p}_{ict} = p_{ict} - \frac{1}{C} \sum_c p_{ict}$) and take the standard deviation of \tilde{p}_{ict} . When we do this, we find that the standard deviation of intercity price differences for the same good in Japan is 17 percent. By contrast, [Broda and Weinstein \(2008\)](#) find the standard deviation in intercity prices of bar-coded goods is 22 percent in the US and 19 percent for Canadian provinces. The fact that intercity price dispersion of goods in the Japanese CPI is lower than that for bar-coded goods in the US and Canada suggests that the JSB item definitions probably do not include goods that differ substantially in quality in different cities and therefore that quality variation across cities for the same product is unlikely to be a major problem in our data.⁸

Table 2 reports the sample statistics for our data. The first three lines of the table report the measure of e-commerce intensity (x_{ib}^E) computed for each of the survey years $b \in \{2004, 2009, 2014\}$. These summary statistics highlight the skewness in the distribution of e-retail sales intensity that we saw in Table 1. Some goods are sold very intensively online, but most goods are purchased predominantly in physical stores. Looking at the value of e-

⁷Source: <http://www.e-stat.go.jp/SG1/estat/List.do?bid=000001033703&cycode=0>

⁸In order to further reduce the impact of measurement error, we also trimmed 3 smallest and 3 largest price quotes within an item-year observation..

commerce intensity in the middle of the period (x_{i09}^E), we see that goods in the upper 90th percentile of the distribution have an e-commerce sales intensity of 0.19 over the full sample period, which is more than eight times higher than a good with the median intensity. We also can see that there are substantial increases in e-commerce intensity at the 50th and 90th percentiles, which reflects the growing importance of e-commerce across the set of goods in the sample. The fourth line of the table recomputes e-commerce intensity using the data from Rakuten (x_{i10}^R). The values of x_{i09}^E and x_{i10}^R are quite similar, indicating that both datasets indicate similar distributions of the e-commerce intensity. The fifth line of Table 2 presents catalog sales intensity in 1999, which has a similar distribution as e-commerce intensity, especially in the later years. Finally, the last row shows the annual growth rate of prices from 1991 to 2016. Eighty percent of annual product price changes in cities were less than 11 percent in absolute value, with the typical price change being close to zero.

Table 2: Summary Statistics

	N	Mean	St. Dev.	Min	p10	p50	p90	Max
x_{i04}^E	321	0.043	0.095	0.000	0.000	0.011	0.091	1.000
x_{i09}^E	325	0.070	0.111	0.000	0.000	0.022	0.193	1.000
x_{i14}^E	324	0.103	0.143	0.000	0.016	0.037	0.280	1.000
x_{i10}^R	327	0.064	0.114	0.000	0.001	0.017	0.175	1.000
x_{i99}^C	327	0.074	0.111	0.000	0.001	0.029	0.217	1.000
Δp_{ict}	791,747	-0.001	0.112	-1.798	-0.114	0.000	0.110	1.679

Note: This table shows summary statistics of e-commerce intensity, catalog intensity, and annual price growth rates from 1991 to 2016; x_{i04}^E , x_{i09}^E , and x_{i14}^E denote e-commerce intensity in 2004, 2009, and 2014, respectively; x_{i10}^R is e-commerce intensity using Rakuten sales data in 2010; x_{i99}^C denotes catalog intensity in 1999, where these variables are defined in equation (10). Prices are in natural logs. Δp_{ict} is the one-year log difference in prices. N indicates the number of products in the sample for the x_{it}^r variables and the total number of observations for Δp_{ict} .

4 Results

We present our results in three sections. Section 4.1 shows how e-commerce has affected national prices. In Section 4.2, we present plots to show that price convergence is a central tendency in the data and that the internet appears to have changed the rate of convergence for goods available online but not for other goods. This provides some *prima facie* evidence that our focus on relative intercity price movements of goods sold by e-retailers is in line with the data. We next estimate the impact of e-retail on the rate of price convergence and we present our estimates of the welfare gain from e-retail. Finally, in Section 4.3, we show welfare gains in new trade models.

4.1 E-commerce and National Prices

Goolsbee and Klenow (2018) found that goods traded online have rates of price change that were about 1.3 percentage points lower than the rates for goods in the same product categories in the CPI. Here, we present complementary evidence that there are different rates of price increase for goods sold intensively by e-merchants relative to other goods. We also find that differential rates of price increase were present long before the entry of e-commerce firms; these became more pronounced after the entry of e-commerce merchants; and they arose in part because the rate of price increase of goods not available online rose. In order to examine this in the data, we regress annual log price changes of goods (Δp_{ict}) from 1991 to 2016 on good fixed effects along with an indicator variable, D_t , that is one starting in 1997 (the year the first e-commerce firm, Rakuten, opened) and zero before as well as the e-commerce intensity of the good interacted with this dummy ($x_{i09}^E D_t$):

$$\Delta p_{ict} = \alpha_i + \phi D_t + \theta x_{i09}^E D_t + \epsilon_{ict}, \quad (11)$$

where α_i is a parameter to capture any pre-trends in the data that might arise if goods available online exhibit different price increase trends than goods not available online. If any differential trend in prices emerged after 1996, we should pick it up in our estimate of ϕ . The most interesting

parameter for our purposes is the coefficient on the e-commerce intensity interaction term (θ), which indicates the differential rate of price change for goods traded online after the entry of e-commerce firms.

One of the advantages of our specification is that we can eliminate any good-specific pre-trends (α_i) that might confound specifications that compare price growth rates of goods available online with those not sold online. In order to understand whether controlling for these pre-trends is likely to be important, we split the sample into two groups by e-commerce sales intensity. The first sample of goods (X_B) consists of products that have an e-commerce sales intensity (x_{i09}^E) in the bottom quartile, and the second sample is composed of goods with an e-commerce sales intensity in the top quartile (X_T). We then computed the average rate of price increase for the two sets of goods by running the following regression separately for each sample:


$$\Delta p_{ict} = \theta_t + \epsilon_{ict}, \quad (12)$$

where the estimates of the time fixed effect θ_t in each sample tell us the average rate of price increase for the goods in each sample.

We plot these expenditure-weighted estimates and the 95-percent confidence bands in Figure 2. There is clearly some volatility in each of the series. Most notably, there are very negative rates of price increase for low e-commerce intensity goods in 1992, which (as we show in the online appendix) is entirely due to idiosyncratic movements in fruit and vegetable prices that year. The relative magnitude and approximate timing of the price swings of e-commerce intensive products CPI data line up (within a year) with fluctuations in the domestic producer price index, which suggests that part of the volatility reflects the underlying price volatility of domestically manufactured goods.⁹

⁹For example, the annual growth rate of Japanese domestic producer price index fell from 0.0 in 2000 to -1.4 percent in 2003. It then rose to 1.3 percent in 2007 and 4.0 percent in 2008 before falling to -5.0 percent in 2009. By 2014, Domestic producer prices were rising at a 2.8 percent annual rate. Source <https://fred.stlouisfed.org/series/JPNPPDMMINMEI#0>. In the online appendix Figure A.2, we also present unweighted estimates, which show a similar pattern. The relatively low bottom-quartile point estimate for 1992 is driven by

Figure 2: Price Growth of Goods with High and Low E-Commerce Intensity



Note: The lines show the average rate of expenditure-weighted price increase for the goods in two groups: products with bottom quartile e-commerce sales intensity (thick line) and products with top quartile e-commerce intensity (thin line). Shaded areas show the 95-percent confidence bands. The red dashed line shows the average rate of price increase before and after the entry of e-commerce firms for goods with bottom quartile e-commerce sales intensity, and the blue dashed line shows that for goods with top quartile e-commerce sales intensity.

Before the entry of the Rakuten in 1997, the average rate of price increase for the types of goods that would ultimately be sold on e-commerce platforms was -1.8 percent per year, while the average annual rate of price increase for goods that did not sell much on these platforms was 1.0 percent per year. Thus, even before the entry of e-commerce firms, there was a 2.8 percentage point gap between the relative rates of price increase for

fruits and vegetable prices in 1992, as can be seen by looking at the online appendix Figure A.1. When we drop fruits and vegetables, which have very volatile price movements in Japan, the rates of price change of goods in the bottom and top quartiles look quite similar in 1992. [Ahn et al. \(2011\)](#) argue the 2007 spike in prices may have been related to credit constraints.

goods that would be sold intensively online relative to those never sold intensively online. These differences in rates of price increase may reflect the fact that the production of standardized, non-perishable goods, which tend to dominate e-commerce platforms, may benefit more from the cost reductions associated with modern manufacturing techniques.

It is also interesting to see what happened to this gap in rates of price change after the entry of e-commerce firms. While we do not see much change in pricing behavior in the first five years after the entry of Rakuten, by 2002, we see that the differences in the price growth rates between the two sets of goods widened significantly in subsequent years, when e-commerce firms became major players in Japanese retail. Goods in the top quartile of e-commerce sales intensity had an average rate of price growth from 1997 to 2016 of -3.0 percent per year: a 1.2 percent per year fall in the rate of price growth. By contrast, the rate of price growth for goods in the bottom quartile of e-commerce sales *rose* to 2.9 percent per year: an increase of 1.9 percent per year. Thus, the gap widened by 3.1 percentage points. Nevertheless, the fact that the rate of price increase for e-commerce intensive goods fell while that of non-intensive goods rose suggests that one must be cautious about interpreting this evidence in terms of e-commerce having an effect on aggregate inflation.

Turning to our differences-in-differences specification, we present the results from estimating equation (11) in Table 3. The first column presents the results from estimating equation (11) over the period 1992 to 2001. Consistent with what we observed in Figure 2, we do not find much of an effect from e-commerce in the first few years after the entry of Rakuten as evidenced by the insignificant coefficient on the post-e-commerce intensity interaction ($x_{i09}^E D_t$) term. However, as one can see in column 2, we do see a significant decline in the *relative* prices of goods available online when we include later years in the sample.

As we have argued earlier, one possible challenge to our identification strategy is that e-commerce firms are not likely to have chosen which sectors they are likely to have entered randomly. In order to deal with this en-

Table 3: Relative Price Changes and E-Commerce Intensity

	(1)	(2)	(3)	(4)
	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}
D_t	-0.0015 (0.0024)	0.0094 (0.0023)	-0.0036 (0.0035)	0.0108 (0.0030)
E-Commerce Intensity $\times D_t$	0.0037 (0.0251)	-0.0912 (0.0209)	0.0456 (0.0601)	-0.1158 (0.0384)
Fixed Effects	Product	Product	Product	Product
Estimation Period	1992-2001	1992-2016	1992-2001	1992-2016
Observations	272,594	581,803	272,594	581,803
R^2	0.03	0.03		
First-Stage F-Stat			27.21	30.26
Estimation Method	OLS	OLS	IV	IV

Note: The dependent variable is the one-year log price change; D_t is dummy variable that is one after 1996 (i.e., the post e-commerce period); and e-commerce intensity (x_{i09}^E) is defined in equation (10). The table shows relative price changes for goods sold online intensively relative to goods not sold online intensively before and after the entry of e-commerce firms. Column 1 and 3 are for 1992-2001, and column 2 and 4 are for 1992-2016. The first two columns show ordinary least-squares (OLS) estimates using e-commerce sales intensity, and the second two columns use catalog sales intensity as IV. The F-statistic corresponds to the Kleibergen-Paap rk Wald statistic. Standard errors in parentheses.

dogeneity, we use catalog intensity (x_{i99}^C) as an instrument for e-commerce intensity (x_{i09}^E). Table 4 presents the first-stage results, which show the strength of catalog sales intensity as an instrument for e-commerce intensity. As one can see from the F -statistic reported in the first two columns of the table, catalog sales intensity in 1999 is a strong instrument for e-commerce sales intensity in 2009. Sectors that on average were major channels for catalog sales also became major channels of e-commerce firms. In the third column, we simply regress the e-commerce intensity of sectors in 2009 on catalog intensity in 1999 to show that the relationship holds in the cross section.

We report the results from our instrumental variables (IV) estimation in columns 3-4 of Table 3. Instrumenting for e-commerce sales intensity does not affect the results significantly, which suggests that there is not a

Table 4: First Stage of Instrumental Variables Regression

	(1)	(2)	(3)
	E-Commerce Intensity $\times D_t$	E-Commerce Intensity $\times D_t$	E-Commerce Intensity
Catalog Intensity $\times D_t$	0.2900 (0.0556)	0.2913 (0.0530)	
D_t	0.0304 (0.0044)	0.0317 (0.0044)	
Catalog Intensity			0.3074 (0.0410)
Constant			0.0444 (0.0059)
Fixed Effects	Product	Product	None
Estimation Period	1992-2001	1992-2016	
Observations	272,594	581,803	385
R^2	0.07	0.07	0.13
Estimation Method	IV-First Stage	IV-First Stage	OLS

Note: E-Commerce intensity (x_{i09}^E) and catalog intensity (x_{i99}^C) are defined in equation (10); D_t is dummy variable that is one after 1996 (i.e., the post e-commerce period). The first two columns of the table present the first-stage regression results (i.e., separate observations for each city and year): column 1 shows results using data for 1992-2001, and column 2 uses data for 1992-2016. The last column presents an OLS regression using only the goods data for 2009. The F-statistic corresponds to the Kleibergen-Paap rk Wald statistic. Standard errors in parentheses.

serious endogeneity bias.¹⁰ As before, we do not see much of an effect of e-commerce on national pricing in the first few years after the entry of Rakuten and the other e-commerce firms, but we do see strong effects in subsequent years. Overall, our IV estimate of the impact of e-commerce intensity ($x_{i09}^E D_t$) on price increases is about 20 percent larger in magnitude than the OLS estimate in the full sample estimates (columns 2 and 4), but

¹⁰However, we do find that instrumenting does matter when we use the Rakuten definition of e-commerce sales intensity, which may reflect the fact that it may help eliminate attenuation bias when using measures of e-commerce intensity not based on random samples.

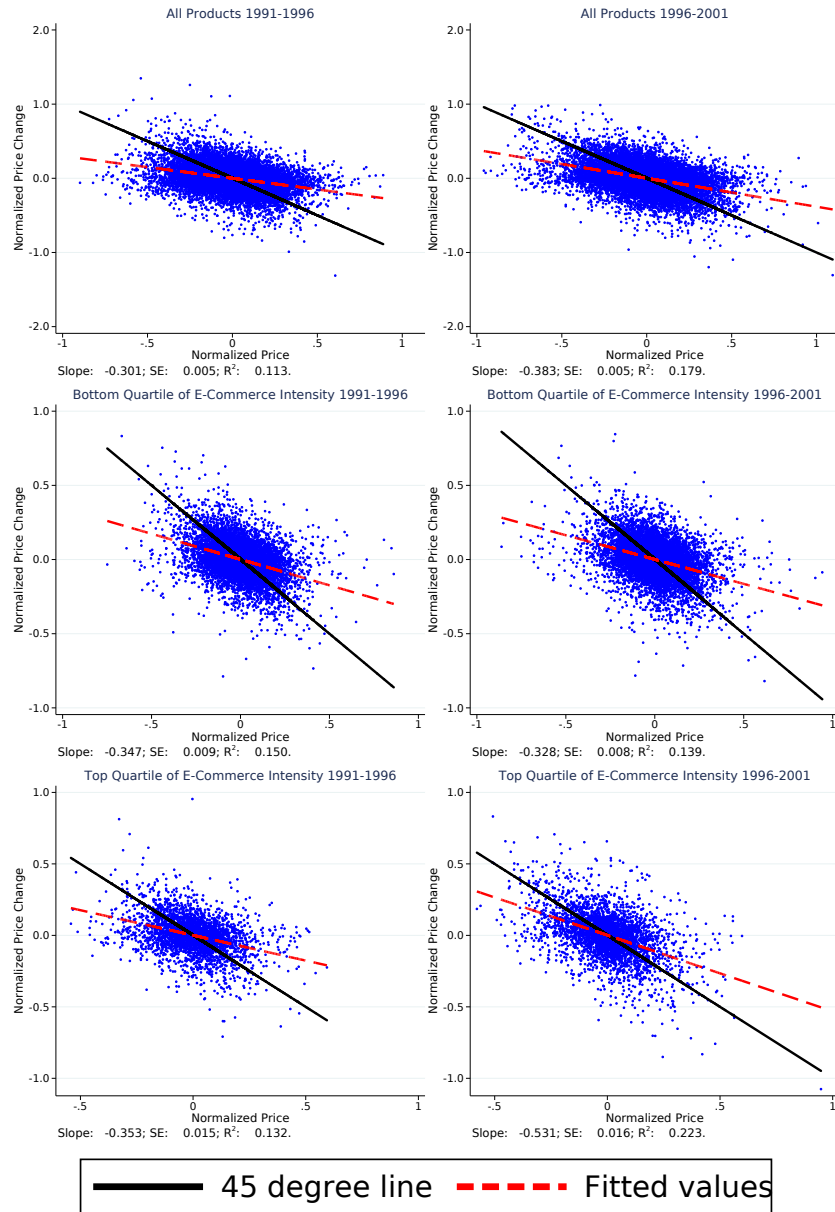
the difference is not significant. As we noted earlier, these results indicate that there is not significant endogeneity bias when thinking about aggregate prices. It is also consistent with our e-commerce variable not exhibiting much measurement error, which would also produce an attenuation bias in the OLS results. In terms of economic significance, the results in column 4 imply that the entry of e-commerce firms caused a good at the 90th percentile of internet sales intensity to have relative rates of price increase to fall by -2.2 percentage points per year.

4.2 Gains Due to Price Arbitrage

In this section, we explore an alternative mechanism through which e-commerce affected prices in Japan: reducing price dispersion across cities. We first consider two five-year periods: the first period (1991-1996) predates e-commerce and the second one (1997-2001) ends after Rakuten was a prominent, listed company. It is difficult to compare price changes across goods and cities in their raw form because rates of price change vary across goods and cities. We therefore normalize the data by regressing Δp_{ict} and p_{ict} on product and city fixed effects and construct normalized log price changes ($\Delta^5 p_{ict} - \hat{\alpha}_{it} - \hat{\beta}_{ct}$) and normalized log price levels ($p_{ic,t-5} - \hat{\alpha}'_{it-5} - \hat{\beta}'_{ct-5}$), where $\hat{\alpha}_{it}$ ($\hat{\alpha}'_{it-5}$) and $\hat{\beta}_{ct}$ ($\hat{\beta}'_{ct-5}$) are the estimated fixed effects from the regression of Δp_{ict} (p_{ict-5}) on product and city fixed effects. Thus, these normalized log prices remove the effect of any common price movements at the product or city level. Figure 3 presents plots of normalized five-year log changes in prices ($\Delta^5 p_{ict} - \hat{\alpha}_{it} - \hat{\beta}_{ct}$) against the normalized five-year lag of log prices in each city ($p_{ic,t-5} - \hat{\alpha}'_{it-5} - \hat{\beta}'_{ct-5}$).

The top panel shows how normalized log price changes vary with normalized log prices before and after the entry of e-commerce. There is a clear negative relationship between initial urban price deviations and future price growth, which indicates that goods that had high prices in cities tend to have lower rates of price increase than goods with low relative prices. This mean reversion is likely the product of price arbitrage. As one can see from these two plots, 30 percent of any relative price difference tends to

Figure 3: Normalized Price Change vs. Normalized Price



Note: This graph plots normalized price changes against normalized price levels. Normalized log price changes (log levels) equal the actual log price changes (log levels) less the fixed effects from a regression of log price changes (log levels) on product and city fixed effects. The left plots show normalized price changes before the entry of e-commerce, and the right panel shows them after the entry of e-commerce. The top-two plots are for all goods; the middle-two are for goods with e-commerce intensity lower than the bottom quartile, and the bottom-two are for goods with e-commerce intensity higher than the top quartile.

be eliminated within five years before the advent of e-commerce, and this number rises to 38 percent in the five years after e-commerce firms entered. These plots also speak to the relatively high quality of the Japanese data. For example, studies using U.S. data (c.f., [Parsley and Wei \(1996\)](#)) find no evidence of price convergence once one controls for city fixed effects.¹¹

The next two pictures show what is driving this increase in the intercity rate of price convergence. Here, we divide the sample into the set of goods with an internet sales intensity in the lowest first quartile of the distribution in 2009 ($x_{i09}^E < 0.014$) and the set of goods in the highest quartile of the distribution ($x_{i09}^E > 0.117$). As one can see from the middle panel in Figure 3, there was almost no change in the rate of convergence for goods not sold on the e-commerce. The slope of the line for goods not sold intensively online in the early period is -0.35, which is almost identical to the slope in the post e-commerce period (-0.33). In other words, the entry of e-commerce firms seems not to have affected the speed at which intercity price differentials converged for goods not sold intensively online. However, we see a very different pattern for goods with an e-commerce intensity in the upper quartile of the distribution. The slope steepens by 50 percent, rising in magnitude from -0.35 to -0.53. Thus, enabling consumers to shop online seems to have significantly reduced the ability of merchants to charge different prices in different cities for the same good. We now turn to exploring this result rigorously.

4.2.1 Estimating Convergence Rates

Table 5 presents the results of estimating equation (1) for five- and one-year intervals using 1999 catalog sales intensity as an instrument for e-commerce sales intensity. In the first two columns, we present separate regressions for 1996 and 2001, where we let the convergence rates vary across the two time periods as we did in the earlier plots. Comparing the first rows of columns

¹¹One plausible reason for the weaker evidence of price convergence in the U.S. is that the data used in [Parsley and Wei \(1996\)](#) is not based on purposive sampling, so price changes in cities are based on different goods in different locations (as shown in [Handbury and Weinstein \(2015\)](#)).

1 and 2 reveals the convergence rates for goods not suitable for e-commerce (i.e., those where $x_{i09}^E = 0$) were almost identical before and after the entry of e-commerce, which is the result that we saw in Figure 3. The coefficient on e-commerce intensity interacted with lagged prices ($x_{i09}^E p_{ic,t-5}$) in column 1 indicates that the rate of convergence for goods suitable for e-commerce sales was not significantly different than the convergence rate of other goods prior to the entry of e-commerce. However, the negative and significant coefficient on the triple interaction term ($D_t x_{i09}^E p_{ic,t-5}$) in the post-e-commerce sample indicates that goods sold intensively online exhibit significantly faster convergence rates after the entry of e-commerce firms, which formally confirms the result we saw in Figure 3.

Table 5: Price Convergence in Pre and Early Post E-Commerce Period

	(1)	(2)	(3)	(4)
	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}
Lagged Price	-0.292 (0.031)	-0.324 (0.037)	-0.292 (0.031)	-0.130 (0.014)
E-Commerce Intensity × Lagged Price	-0.174 (0.405)		-0.218 (0.383)	0.375 (0.168)
$D_t \times$ Lagged Price			-0.030 (0.029)	0.008 (0.012)
$D_t \times$ E-Commerce Intensity × Lagged Price		-1.154 (0.588)	-0.961 (0.472)	-0.598 (0.181)
t	{1996}	{2001}	{1996,2001}	Annual 1991-2001
Observations	26,265	27,693	51,880	272,594
First-stage F	29.90	33.99	17.01	16.51
Estimation	IV	IV	IV	IV

Note: E-Commerce intensity (x_{i09}^E) is defined in equation (10); D_t is dummy variable that is one after 1996 (i.e., the post e-commerce period). The table shows regression results of equation (1) using IV: e-commerce sales intensity in 2009 is instrumented using 1999 catalog sales intensity. The first column uses the five-year log differences in prices between 1991 and 1996 and the second column uses five-year differences from 1996 and 2001. The third column uses two five-year periods, 1991–1996 and 1996–2001. The dependent variable in the fourth column is the annual log differences in prices from 1991 to 2001. The F-statistic corresponds to the Kleibergen-Paap rk Wald statistic. Standard errors in parentheses.

In column 3 of Table 5, we estimate our baseline differences-in-differences specification of equation (1) using a five-year differences by letting t take on two values: 1996 and 2001. The most important result for our purpose is the estimate of the coefficient on the interaction term on the e-commerce intensity coefficient. As one can see from the table, the coefficient is negative and precisely measured. Not surprisingly, the estimated coefficient on $p_{ic,t-k}(\hat{\gamma})$ does not change much, and we continue to get a negative and significant coefficient on the e-commerce intensity interaction term in the five-year difference specification ($\hat{\delta}_3 = -0.961$) and in the one-year difference specification ($\hat{\delta}_3 = -0.598$). The estimate of δ_1 is insignificant, which means that before the entry of e-commerce firms, there was no clear difference in the rates of price convergence between goods that would ultimately be sold intensely online relative to those would not. This result is very much in line with the bottom two left-hand side plots in Figure 3.

We now turn to exploring the robustness of these results to alternative specifications. The first three columns of Table 6 replicate column 4 of Table 5 using the full sample of data and alternative base years to measure e-commerce intensity (i.e., we use $x_{ib}^E \in \{x_{i04}^E, x_{i09}^E, x_{i14}^E\}$). The results are all qualitatively similar. Interestingly, we find that the OLS results (reported in the online appendix section A.2) have the same sign as our IV results, but tend to have significantly smaller coefficients. For example, the point estimate on the triple interaction ($D_t x_{i09}^E p_{ic,t-5}$) is half as large (-0.237, with a standard error of 0.100) if we estimate the parameters just using OLS. In the next three columns, we replace our continuous measure of e-commerce intensity with a dummy in order to reduce the potential role played by outliers. “High E-Commerce Intensity” is one if x_{ib}^E is above the median. The results reported in columns 4-6 are qualitatively similar to the results based on the continuous measure of e-commerce intensity, which suggests that outliers are not driving our result.

The second concern that one might have with the results is that we may have a data measurement problem that is influencing the results. In order to make sure that some idiosyncratic component of the NSFIE survey method

Table 6: Price Convergence Over Alternative Periods

	(1)	(2)	(3)	(4)	(5)	(6)
	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}
Lagged Price	-0.132 (0.016)	-0.128 (0.014)	-0.133 (0.016)	-0.123 (0.011)	-0.123 (0.011)	-0.125 (0.013)
E-Commerce Intensity × Lagged Price	0.685 (0.362)	0.298 (0.160)	0.276 (0.152)			
$D_t \times$ Lagged Price	-0.014 (0.015)	-0.022 (0.013)	-0.014 (0.016)	-0.032 (0.013)	-0.031 (0.013)	-0.028 (0.015)
$D_t \times$ E-Commerce Intensity × Lagged Price	-1.315 (0.380)	-0.570 (0.151)	-0.493 (0.144)			
High E-Commerce Intensity × Lagged Price				0.023 (0.027)	0.021 (0.024)	0.029 (0.037)
$D_t \times$ High E-Commerce Intensity × Lagged Price				-0.043 (0.023)	-0.040 (0.020)	-0.056 (0.031)
t	Annual 1991-2016					
Observations	579,157	581,803	578,704	579,157	581,803	578,704
E-Commerce Intensity Year	2004	2009	2014	2004	2009	2014
First-stage F	20.53	16.91	8.09	29.50	54.98	13.58
Estimation	IV	IV	IV	IV	IV	IV

Note: E-Commerce and catalog intensity are defined in equation (10); D_t is dummy variable that is one after 1996 (i.e., the post e-commerce period). The first three columns of the table show the regression results of equation (1) using 1999 catalog sales intensity as an instrument for e-commerce sales intensity in 2004 (column 1), 2009 (column 2), and 2014 (column 3). The last three columns of table present regression results with discrete e-commerce intensity variable that indicates whether a good has an e-commerce intensity above the median. The F-statistic corresponds to the Kleibergen-Paap rk Wald statistic. Dependent variables are the annual log price differences from 1991 to 2016. Standard errors in parentheses.

is not driving our results, we replicate all our main results using measures of e-commerce intensity based on Rakuten sales data instead. We report the results from this exercise in Appendix A.3 Table C3 for over period 1991 - 2001 and Table C4 for 1991 - 2016, which shows that we obtain very similar results regardless of whether we measure e-commerce sales intensity using consumer expenditure data or Rakuten e-commerce sales data.

4.2.2 Gains from Arbitrage

Table 7: Counterfactual Welfare Gain due to Price Arbitrage

σ	ΔW_{14}
3	0.0021
4	0.0029
5	0.0036
6	0.0043
7	0.0050

Note: ΔW_t equals the change in welfare in each year as a share of expenditure in the previous year and is computed according to equation (9).

Aggregate consumer gains due to faster price convergence can be calculated from equation (9). One of the features of these equations is that the welfare gain is proportional to the choice of demand elasticity. In all cases, we base our estimates of the impact of e-retail on the rate of convergence in Table 6 column 2. Estimating these elasticities using our data is difficult because we do not have good instruments for the price of goods in each city. We, therefore, choose to rely on prior estimates of demand elasticities. Interestingly, estimates of the elasticity of substitution across goods available in retail stores and across retail merchants tend to be close to four, so we will adopt that estimate as our baseline one in the welfare calculations in this section and the next.¹² Table 7 presents that welfare gains arising from price

¹²For example, Dolfen et al. (2019) estimate the elasticity between online and offline merchants to be 4.3. Hottman (2019) estimates the elasticity of substitution across stores within a city and obtains a median estimate of 4.7. Hottman et al. (2016) estimate the elasticity of substitution across firms selling packaged goods for a wide range of products and

arbitrage for 2014. We find that if we adopt an elasticity of 4, we estimate that the gains from faster price arbitrage equals 0.3 percent of expenditure. The magnitude of the gain is proportional to the elasticity used in the calculation, but all reasonable estimates of the elasticity of substitution suggest substantial welfare gains through this channel.

4.3 Variety Gains

An alternative mechanism through which e-commerce might affect welfare is by enabling consumers to access new varieties as in [Brynjolfsson et al. \(2003\)](#). One of the challenges of estimating the gains from new varieties is that our data does not enable us to see which varieties became available. Fortunately, we do observe sufficient statistics that enable us to compute the welfare gain even in a world in which we do not see the underlying varieties. In order to do this, we adopt the framework of [Feenstra \(1994\)](#) to compute the change in the CES price index due to new varieties. Since there is no clear evidence that e-commerce affected average prices, we will assume that these prices were unaffected by e-commerce and only focus on varieties. In this case, Feenstra showed that the log change in the consumer price index due to new varieties can be written as:

$$\Delta \ln P_t = \frac{1}{\sigma - 1} \ln \left(\frac{\lambda_t}{\lambda_{t-k}} \right), \quad (13)$$

where $\lambda_t \in (0, 1]$ is the share of consumer expenditures in period t on varieties available in *both* periods t and $t - k$.

In order to take the theory to data, we need to make some additional assumptions. First, we assume that purchasing from an e-commerce merchant differs in some way from purchasing the good from a physical store and that this differentiation is captured in σ , which we continue to set equal 4 in our base case. Second, we assume that there was no change in the qual-

obtain a median value of 3.9. [Thomassen et al. \(2017\)](#) estimate that the median markup for supermarkets in the UK to be 1.30, which in a CES setup with monopolistically competitive firms would correspond to an elasticity of substitution across retailers of 4.3 (where we make use of the fact that in a monopolistically competitive setting with CES preferences the markup equals $\sigma / (\sigma - 1)$).

ity of the shopping experience in physical stores over this period. While there definitely was turnover in physical retailers over this time period, this is not a problem for our estimation if the stores that exited were replaced with stores of comparable quality, so that consumers' shopping experience in physical stores remained unchanged.

There is evidence that this assumption is reasonable. For example, when Rakuten entered the Japanese market in 1997, there were 1,015 square meters of retail floor space per capita, and in 2014, this number stood at 1,061 square meters. Similarly, the amount of physical retail space in Japan stayed almost constant at approximately 130 million square meters over the entire time period, and employment in retail trade stayed constant at 7.4 million people between 1994 and 2012.¹³ Thus, the rise of e-commerce is not associated with a fall in the amount of space used to display goods by physical merchants or the number of sales staff in these stores. Based on this, we think that it is fair to assume that consumers entering Japanese physical stores could continue to experience comparable shopping experiences. In this case, we can set λ_{t-k} and λ_t equal to the share of expenditures in physical retailers in periods $t - k$ and t .

The rise of e-commerce merchants is necessarily associated with a decline in expenditure shares in physical retailers. Based on equation (13), it is straightforward to see how e-commerce will affect prices. If we start in a simple case in which consumers face two choices for goods purchases—e-retail or other retail—and choose a base period ($t - k$) that predates e-commerce, it must be the case that $\lambda_{t-k} = 1$ because the initial share of purchases from e-commerce firms is zero by construction. However, the share purchased in period t (λ_t) will be less than one because consumers in the later period will only purchase a fraction of their goods from other retail stores. Thus, the price level in period t will fall because consumers now have access to new (online) varieties.

One problem with this approach is that it implicitly assumes that the

¹³Source: Japan Statistical Handbook, various years
<http://www.stat.go.jp/english/data/nenkan/index.html>

elasticity of substitution between physical retailers and e-commerce merchants (σ) is the same as that between e-commerce merchants and catalog merchants. However, it may be the case that e-commerce is more substitutable with catalog sales than it is with purchases in physical stores. If e-commerce is perfectly substitutable with catalog sales, then we should compute our λ ratios based on the purchases from physical stores divided by the sum of purchases from physical and telemarketing merchants (i.e., e-commerce and catalog merchants). In this case, the share of purchases from physical merchants in period $t - k$ (λ_{t-k}) will be less than one because consumers purchased some goods from catalogs before the entry of e-commerce firms. Similarly, we will also have $\lambda_t < 1$ because in the later period consumers purchase from both catalog and e-commerce merchants. Intuitively, the price index will fall as long as $\lambda_t/\lambda_{t-k} < 1$, which is equivalent to saying that a necessary condition for e-commerce to lower the price index is that its sales did not come completely at the expense of catalog merchants. We will refer to the price change based on this assumption as ΔP_t^T to reflect the fact that the change in the price level can be thought of as the consumer gain from telemarketing.

In order to implement these calculations, we need to first adjust the data to take into account that not all consumer expenditures occur through retailers. We first do the calculation under the assumption that we can treat all retail channels other than e-commerce as continuing to sell a constant set of varieties of constant quality. In the NSFIE data, the share of household expenditures purchased from retailers (χ) was 62 percent in 2014, with the remaining expenditures (e.g., utilities, travel, and education) being made from non-retail sources. In 2014, e-commerce expenditures on goods as a share of all retail expenditures, which we denote by ψ , was 0.0437. The share of household expenditures from non-e-commerce firms in 2014 was therefore $\lambda_{14} = (1 - \psi)\chi + 1 - \chi = 0.97$. Assuming an elasticity of substitution of 4, we can use equation (13) to obtain an estimate of the percentage price drop between 1996 and 2014 associated with e-commerce in Japan ($\Delta \ln P_{14}^E$) of 0.9 percent. We report this number in the first column of Table

8 along with a number of alternative estimates based on different plausible estimates of the trade elasticity. These consumer gains range from 0.5 percent to 1.4 percent in 2014.

Table 8: Price Change Due to E-Commerce Varieties


σ	$\Delta \ln P_{14}^E$	$\Delta \ln P_{14}^T$
3	-0.014	-0.015
4	-0.009	-0.010
5	-0.007	-0.008
6	-0.005	-0.006
7	-0.005	-0.005

Note: $\Delta \ln P_t^E$ is the implied log price change associated with the entry of e-commerce firms under the assumption that elasticity of substitution between e-commerce and physical retail is the same as between e-commerce and catalog merchants. $\Delta \ln P_t^T$ is the implied log price change associated with the entry of all telemarketing firms (e-commerce and catalog merchants).

The second column in Table 8 defines local sales as total expenditures less expenditures on products sold over the internet or through catalogs. These estimates will be smaller if e-commerce sales came at the expense of catalog sales, but will be larger if both e-commerce and catalog sales expanded relative to retail sales. Interestingly, we see that associated price reductions ($\Delta \ln P_t^T$) are larger when we allow for the fact that consumers purchased and continue to purchase goods through catalogs. This result may be due to the fact that all telecommunications prices fell, which benefited both catalog and e-commerce merchants at the expense of physical retailers.

The implicit gain due to varieties is larger than the estimate from the price arbitrage approach in Section 4.2.2 and reflects the impact of different modeling assumptions on the welfare estimates. In the price arbitrage approach, e-commerce may generate exports of products from one city to another (which might appear in the data as e-commerce transactions) but there is no variety gain. Instead, the welfare gain arises from the reduction in price dispersion across cities from sales of the same products. By contrast,

Figure 4: Price Change Due to New Varieties vs. Share of College Education



Note: Figure plots the log price change due to the increased variety ($\Delta \ln P_{14}^E$) from e-commerce in 2014 against the share of college education.

in the variety approach, consumers gain from the ability of shopping from e-commerce firms even if these merchants offered the same products at the same prices (perhaps because online shopping saves time). If one thinks the first approach applies to commonly available goods and the second one to new ways of purchasing goods, one could add the two estimates to obtain the total. However, in an [Eaton and Kortum \(2002\)](#) setup, there is no variety channel, so as [Arkolakis et al. \(2012\)](#) argue, in this case the “variety” channel would simply be measuring the gains from trade arising from prices converging to their free-trade levels. In order to be conservative, we will therefore assume that the two methods of measuring gains are not additive, so our baseline estimate of the gains due to e-commerce (using an elasticity of substitution of 4) ranges from 0.3 to 0.9 percent.

One of the characteristics of calculations of the variety gains due to e-commerce is that the magnitude of these gains is driven entirely by the increase in the share of sales by e-merchants. This share is strongly correlated with the share of college graduates in a prefecture as one can see in

Figure 4, which plots the implied log price change due to e-commerce at the prefectural level (ΔP_{ct}^E , where c denotes prefecture) against the share of that prefecture's population with a college degree.¹⁴ One plausible explanation for this high correlation between e-commerce use and education is that highly educated people are likely to be more comfortable using computers and other information technologies and also are more likely to use credit cards and non-cash payments. Whatever the cause, this result suggests that variety models imply a digital divide in which regions with a large share of highly educated people benefit more than regions with fewer highly educated people. These differences are economically significant. In column 1 of Table 9, we regress the price change due to e-commerce on the share of people with a college education in each prefecture. We estimate that the e-commerce gains for a city like Tokyo with a share of college-educated people equaling 0.35 is 1.2 percent, while a city with a college share of only 0.175, like Fukushima, is only half as large.¹⁵

One obvious concern with this plot is that the share of college-educated people might be correlated with other factors that matter for internet purchases. For example, [Dolfen et al. \(2019\)](#) document that e-commerce in the U.S. is positively associated with city size. This is also true in Japan as well, where Tokyo is both the largest city and the city with the second-highest share of college-educated people. Alternatively, it may be the case that income or age is associated with e-commerce intensity. In order to understand the robustness of our regional results to controlling for these factors, we regress the welfare gain on population (which is a proxy for urban vs. rural prefectures), prefectural income per capita, and average age and report the results in columns 2-6 of Table 9. We find that none of these variables are significant once we control for the share of college-educated people in a prefecture. When we include the share of secondary-school graduates, we find that it is significant in one specification, but it has a positive sign, which re-

¹⁴Japanese prefectures are roughly the size of US counties.

¹⁵Whether we include or exclude catalog sales does not matter substantively for our prefectural results. We do not find a similar pattern when computing arbitrage gains because gains are not highly correlated with the share of e-commerce purchases in a locale.

Table 9: Correlates of Implicit Price Drop Due to E-Commerce

	Dependent Variable: $\Delta \ln(P_{14}^E)$					
	(1)	(2)	(3)	(4)	(5)	(6)
Share of College Educated	-0.033 (0.005)	-0.031 (0.007)	-0.037 (0.006)	-0.030 (0.007)	-0.020 (0.008)	-0.025 (0.009)
ln(Population)		-0.000 (0.000)				0.000 (0.000)
ln(Income per Capita)			0.003 (0.002)			0.003 (0.002)
ln(Average Age)				0.005 (0.009)		-0.002 (0.010)
Share of Secondary Educated					0.015 (0.007)	0.016 (0.008)
Constant	0.000 (0.001)	0.002 (0.005)	-0.036 (0.033)	-0.019 (0.035)	-0.009 (0.005)	-0.044 (0.047)
Observations	47	47	47	47	47	47
R^2	0.460	0.461	0.474	0.464	0.506	0.527

Note: The table shows how prefectural welfare gains due to increased variety relate to characteristics of the prefecture - the share of the college education, population, income per capita, and the share of secondary education. Standard errors in parentheses.

inforces our earlier point that it is highly educated people that are the main users of e-commerce.¹⁶

5 Conclusion

This paper makes use of a unique Japanese data set covering hundreds of products over close to three decades to examine the impact of the e-retail on Japanese relative prices and welfare. While we find that at the na-

¹⁶This may explain why [Fan et al. \(2018\)](#) find no link between education and internet sales intensity. Chinese education levels are much lower than in Japan, which means that very few people have gone to college in their sample. The average number of years of education in [Fan et al. \(2018\)](#) is only 8.8 years, whereas the average in our sample of Japanese cities is 11.9 years.

tional level, the price increases for goods sold intensively online are lower than those sold principally in physical stores, we show that this result was present even before the advent of e-commerce. Nevertheless, the entry of e-commerce firms is associated with a widening of this gap, which is consistent with e-commerce affecting relative price increases.

At the local level, we find strong evidence that the rate at which inter-city price differences converge rose significantly for goods sold intensively online after e-commerce sales became common in Japan, but goods sold predominantly offline experienced no change in their rate of price convergence. This provides evidence that information technology significantly reduced the ability of stores to charge different prices in different locations. Analyzing the impact of this faster rate of price convergence through the lens of the [Jensen \(2007\)](#) model indicates that the welfare gains due to e-commerce were sizable: Japanese welfare in 2014 was 0.3 percent higher as a result of e-commerce.

We also explored the consumer gains due to e-commerce through the lens of models featuring variety gains. These models suggest that the entry of e-commerce firms lowered variety-adjusted prices by 0.9 percent and that these gains benefited highly educated regions more than less educated regions. Although a feature of new-trade theory models is that no location can be made worse off as a result of trade liberalization, the estimated welfare gains in relatively rich cities like Tokyo are two times higher than in most small cities. This result arises from the fact that higher-educated consumers buy substantially more online than less-educated consumers.

References

- Ahn, J., M. Amidi, and D. E. Weinstein (2011). "Trade Finance and the Great Trade Collapse". *AER Papers and Proceedings Session 101*(3), 298–302.
- Aker, J. C. (2010). "Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger". *American Economic Journal: Applied Economics* 2, 46–59.
- Allen, T. (2014). "Information Frictions in Trade". *Econometrica* 82(6), 2041–

- 2083.
- Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). "New Trade Models, Same Old Gains?". *American Economic Review* 102(1), 94–130.
- Bergin, P. R., R. Glick, and J.-L. Wu (2017). "Conditional PPP" and Real Exchange Rate Convergence in the Euro Area". *Journal of International Money and Finance* 73, 78–92.
- Broda, C. and D. E. Weinstein (2008). "Understanding International Price Differences Using Barcode Data". *NBER Working Paper* 14017.
- Brown, J. R. and A. D. Goolsbee (2002). "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry". *Journal of Political Economy* 110(3), 481–507.
- Brynjolfsson, E., Y. J. Hu, and M. D. Smith (2003). "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers". *Management Science* 49(11), 1580–1596.
- Cavallo, A. (2017). "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers". *American Economic Review* 107(1), 283–303.
- Cavallo, A. (2018). "More Amazon Effects: Online Competition and Pricing Behaviors". *Jackson Hole Economic Symposium Conference Proceedings (Federal Reserve Bank of Kansas City)*.
- Cecchetti, S. G., N. C. Mark, and R. J. Sonora (2002). "Price Index Convergence Among United States Cities". *International Economic Review* 43(4), 1081–1099.
- Couture, V., B. Faber, Y. Gu, and L. Liu (2018). "E-Commerce Integration and Economic Development: Evidence from China". *NBER Working Paper* 24384 82(6), 2041–2083.
- Crucini, M. J., M. Shintani, and T. Tsuruga (2015). "Noisy information, distance and law of one price dynamics across US cities". *Journal of Monetary Economics* 74, 52–66.
- Dolfen, P., L. Einav, P. J. Klenow, B. Klopach, J. D. Levin, L. Levin, and W. Best (2019). "Assessing the Gains from E-Commerce". *NBER Working Paper* 25610.
- Eaton, J. and S. Kortum (2002). "Technology, Geography and Trade". *Econometrica* 70(5), 1741–1779.
- Fan, J., L. Tang, W. Zhu, and B. Zou (2018). "The Alibaba Effect: Spatial Con-

- sumption Inequality and the Welfare Gains from E-commerce". *Journal of International Economics* 114, 1580–1596.
- Feenstra, R. C. (1994). "New Product Varieties and the Measurement of International Prices". *American Economic Review* 84(1), 157–177.
- Freund, C. L. and D. Weinhold (2004). "The effect of the Internet on international trade". *Journal of International Economics* 62, 171–189.
- Goldmanis, M., A. Hortaçsu, C. Syverson, and Ö. Emre (2010). "E-Commerce and the Market Structure of Retail Industries". *The Economic Journal* 120, 651–682.
- Goolsbee, A. D. and P. J. Klenow (2018). "Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce". *NBER Working Paper* 24649.
- Goyal, A. (2010). "Information, Direct Access to Farmers, and Rural Market Performance in Central India". *American Economic Journal: Applied Economics* 2, 22–45.
- Handbury, J. and D. E. Weinstein (2015). "Goods Prices and Availability in Cities". *Review of Economic Studies* 82, 258–296.
- Hottman, C. J. (2019). "Retail Markups, Misallocation, and Store Variety across U.S. Cities". *Working paper*.
- Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). "Quantifying the Sources of Firm Heterogeneity". *The Quarterly Journal of Economics*, 1291–1364.
- Jensen, R. (2007). "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector". *The Quarterly Journal of Economics* 122(3), 879–924.
- Olson, P. (2012). "Mission Impossible: How Rakuten Billionaire Hiroshi Mikitani Plans To 'Beat Amazon'". *Forbes*.
- Parsley, D. C. and S.-J. Wei (1996). "Convergence to the Law of One Price Without Trade Barriers or Currency Fluctuations". *The Quarterly Journal of Economics* 111(4), 1211–1236.
- Steinwender, C. (2018). "Real Effects of Information Frictions: "When the States and the Kingdom became United". *American Economic Review* 108(3), 657–696.
- Thomassen, Ø., H. Smith, S. Seiler, and P. Schiraldi (2017). "Multi-category

Competition and Market Power: A Model of Supermarket Pricing". *American Economic Review* 107(8), 2308–51.