The Impact of E-Commerce on Relative Prices and Consumer Welfare*

Yoon J. Jo
Columbia University†

Misaki Matsumura
Columbia University‡

David E. Weinstein
Columbia University and NBER§

July 3, 2019

Download the latest version here

Abstract

Using Japanese data on intercity prices and expenditures by retail outlet type, we find that the entry of e-commerce firms reduced the rate of price increase for goods sold intensively online relative to other goods, and significantly reduced intercity price dispersion of goods sold intensively online but had no effect on other goods. We overcome endogeneity issues by using historical catalog sales as an instrument for e-commerce sales intensity and estimate that reductions in price dispersion raised welfare by 0.4 percent. E-Commerce also lowered variety adjusted prices on average by 0.6 percent, and more in cities with highly educated populations.

JEL CLASSIFICATION: F11, F14, L86, R32
KEYWORDS: e-commerce, trade, prices, arbitrage, variety, welfare

*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 David Nagy, Amine Ouazad, Roberto Rigobon, Takashi Unayama, and Tsutomu Watanabe for many useful comments on early drafts. We also thank Yingjie Wu for providing excellent research assistance.

†420 W. 118th Street, MC 3308, New York, NY 10027. Email: yj2304@columbia.edu.
‡420 W. 118th Street, MC 3308, New York, NY 10027. Email: mm4387@columbia.edu.
§420 W. 118th Street, MC 3308, New York, NY 10027. Email: david.weinstein@columbia.edu.

1 Introduction

How has e-commerce affected prices and consumers? One of the challenges in answering this question is that researchers typically only have short time series that do not allow them to compare pricing dynamics before and after the advent of e-commerce. Thus, while we can observe how the pricing dynamics of goods sold intensively online differs from those not sold online, it is difficult to assess whether any differences arise due to the advent of e-commerce or because of inherent differences in the pricing behavior of the goods themselves. This issue is particularly relevant because the types of products sold intensively online—books, clothing, electronics, and hardware—are also the types of goods that used to be sold intensively through catalogs. Thus, evidence about different pricing dynamics for these types of products is not necessarily evidence that e-commerce *caused* these pricing dynamics.

In order to resolve these issues, this paper makes use of a unique Japanese data set covering price quotes for the set of goods that make up the Japanese consumer price index over the period 1991 to 2016 to examine the impact of the internet on Japanese relative prices and consumer welfare. We merge these data with Japanese government survey data documenting the share of consumption expenditures occurring through each retail channel—catalog, e-commerce, and physical store—for each of these product categories. The long time series enables us to control for important pre-trends in the pricing dynamics of the types of goods available from online merchants. Second, we are also able to correct for an important endogeneity bias arising from the fact that e-commerce firms tend to enter sectors where they anticipate high markups by using historical catalog sales as an instrument for e-commerce sales.

We find that goods sold intensively online have significantly lower rates of price increase than goods not sold much online. While one might be tempted to attribute this 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, this basic differential pattern in pricing behavior seems to be a characteristic of the types of goods amenable to online sales rather than to the advent of e-commerce. Nevertheless, we document that after the entry of e-commerce merchants, the difference in rates of price increase between goods not sold intensively online and those sold intensively rose, suggesting that e-commerce may have exacerbated the difference in relative rates of price increase.

Second, we document that e-commerce had important impacts on rates of intercity price differentials. Following Cavallo (2018), we argue that e-commerce is a technology that promotes uniform pricing across locations. As such, we should expect to see the rate of intercity price arbitrage rise for goods sold intensively online but not for goods sold principally in physical stores. This is exactly what we observe in the data. While we find that prior to e-commerce intercity price differentials dissipated a similar rates for the sets of goods that would eventually be sold online compared to goods not available online, after the advent of the internet, we find that intercity price differentials dissipated rapidly for goods available online but not for goods sold mostly in physical stores.

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 consumers using the model developed in Jensen (2007). Jensen (2007) showed how one can measure welfare gains for consumers from an information technology that reduces price dispersion. While he applied this to cell phones in India, we adapt his framework to e-commerce by modeling it as a technology that reduces intercity price dispersion by enabling consumers to purchase the cheapest version of a good from any merchant selling online. We estimate the consumer gains due to e-commerce to be 0.4 percent of consumption expenditure in 2014. While price arbitrage necessarily produces aggregate welfare gains in this setup, the model also predicts that consumers in low-priced cities lose while consumers in high-priced cities gain. This result is consistent with DellaVigna and Gentzkow (2017) who argue that uniform

regional pricing by chain stores is likely to benefit high-income locations, where demand is likely to be more inelastic and hurt poorer consumers who benefit from the lower prices associated with price discrimination.

In addition to the gains arising from improved price arbitrage across cities, we also compute the gains due to new varieties. We model goods purchased online as new varieties either because they were not available locally before the advent of e-commerce or because an online shopping experience differs in important ways (convenience, service, etc.) from an offline shopping experience. As Brynjolfsson et al. (2003) and Einav et al. (2017) have argued, these variety channels are likely to be quite important. In order to address these concerns, we also use the approach developed in Feenstra (1994) and Arkolakis et al. (2012) to compute the gains due to varieties. The gains due to varieties are similar to those estimated using the Jensen (2007) approach: about a 0.6 percent welfare gain. Interestingly, we also find evidence of an important digital divide in terms of which consumers benefit the most from new varieties. To the extent that access to new varieties can be considered an alternative source of welfare gain, these gains should be seen as an additional source of benefits. Since e-commerce expenditure shares are highly correlated in Japan with college education, we estimate that the gains due to e-commerce in a new-trade theory setup are four times higher in cities with highly educated populations like Tokyo than in cities with low shares of college-educated people like Akita.

1.1 Related Literature

Our results are related to a number of papers related to how information technology has affected pricing and welfare. A large literature 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 also show that e-retail serves to reduce price dispersion. However, our work differs 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. The 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. This fact helps motivate our assumption that local retailers with high prices should face stiff competition from online retailers.

Our paper is also related to studies of the impact of e-commerce on welfare. Many of these studies have 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 Einav et al. (2017) estimate the gains from e-retail due to shopping convenience and new varieties in US. An important difference between these studies and ours is that we make use of household survey data to measure e-commerce sales shares for all goods in Japanese expenditures, and we deal with the endogeneity of e-commerce entry by using historical catalog sales.

Our paper also relates to studies of 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 PPP convergence regressions. Parsley and Wei (1996) were the first to document that differences in convergence coefficients across cities was linked to trade costs, an insight that we build upon in this paper. We estimate that intercity convergence rates for Japan pre-Rakuten 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 prior to the advent of e-commerce, the half-lives for price differentials across Japanese cities were only 4.5 years. 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. Moreover, we find that after the entry of e-commerce firms the half lives of goods sold intensively online collapsed to just a few months whereas goods not sold much online experienced no similar change. 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.

The remainder of the paper is organized as follows. Section 2 introduces the the estimation strategy and provides the theory for the welfare calculation. Section 3 presents the data and provides some stylized facts about e-commerce ssuitability. 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 Sections 4.2.. Section 4.3 presents our calibration of 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 $x_i \in [0, 1]$ be the "e-commerce sales intensity" of a good, where zero indicates it is not suitable for e-commerce and one indicates that it is the most suitable good for e-retail. Let D_t be an indicator variable that is one if e-commerce are positive 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 we introduce a modification that allows the rate of price converge for goods available online to change, i.e.,

$$\Delta^k p_{ict} = \alpha_{it} + \beta_{ct} + (\gamma + \delta_1 x_i + \delta_2 D_t x_i)\, p_{ic,t-k} + \epsilon_{ict}, \tag{1}$$

where α_{it} is a 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 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 price shocks arising from period t local supply-and-demand conditions for an item in a city that are not shared by all items in the city and are uncorrelated with past prices.

In this specification, a critical parameter is the rate of convergence given by $(\gamma + \delta_1 x_i + \delta_2 D_t x_i)$, which we expect to be between -1 and 0. A value of -1 means that 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 $(\gamma + \delta_1 x_i + \delta_2 D_t x_i) = 0$, which implies that the price of that good i in city c follows a random walk with a drift term given by $(\alpha_{it} + \beta_{ct})$. In intermediate cases where $(\gamma + \delta_1 x_i + \delta_2 D_t x_i) \in (-1, 0)$, price differences across cities can persist for more than k years.

Similarly, we can write the approximate half-life of any price deviation from the steady-state price (measured in intervals of length k) as

$$H_t \equiv \frac{\ln(0.5)}{\ln\left(1 + \hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D_t x_i\right)}.$$

As one can see from the above formula, the change in the rate of convergence depends on all of the estimated convergence parameters, therefore there is not a simple mapping from changes in δ_t into rates of convergence. Thus, the impact of e-commerce on the rate of convergence for any good i can be written as:

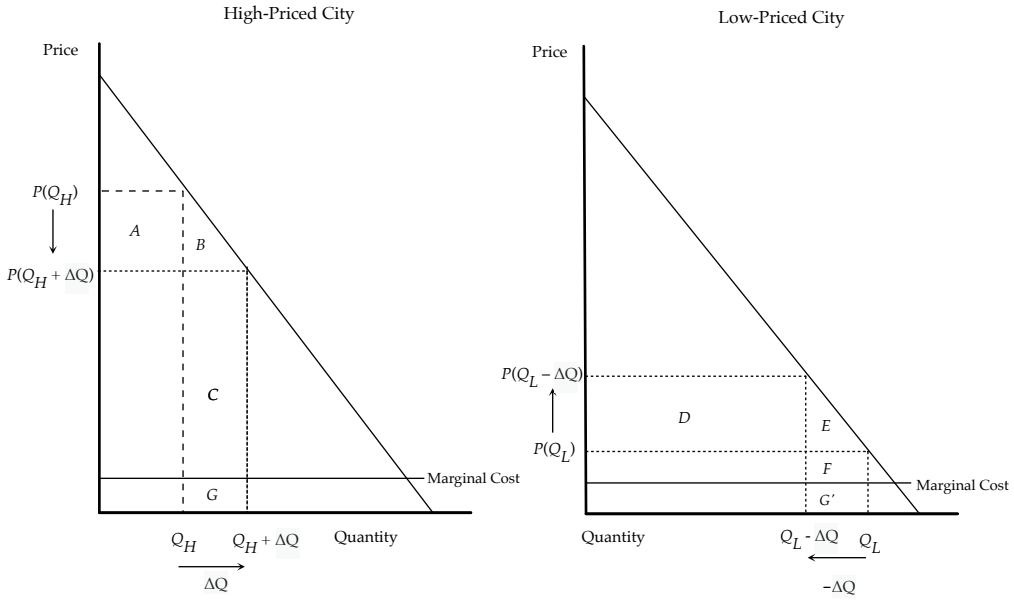
$$\Delta H_t \equiv \frac{\ln(0.5)}{\ln\left(1 + \hat{\gamma} + (\hat{\delta}_1 + \hat{\delta}_2) x_i\right)} - \frac{\ln(0.5)}{\ln\left(1 + \hat{\gamma} + \hat{\delta}_1 x_i\right)}. \tag{2}$$

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 price in H is at least as large as the price in the region L after arbitrage (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 heights 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 costs of production, then $G = G^{'}$, and we can still compute the welfare gain as $(B + C + G) - \left(E + F + G^{'}\right) = (B + C) - (E + F)$.[1]

[1]The 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.

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 will assume that the economy has deviated from a steady state equilibrium but experiences different rates of price convergence and this convergence process does not change the aggregate demand for goods as in Jensen (2007). 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. Given a distribution of prices in period $T-1$ and a steady-state price level $p_{i,}^{*}$, we know that the evolution of prices will follow:

$$\widehat{\Delta p}_{icT}(D) = \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D x_i\right)\left(p_{ic,T-1} - p_i^*\right). \tag{3}$$

The price level in period T can be written as

$$\Delta P_{icT}(D) \equiv P_{icT}(D) - P_{ic,T-1}(D),$$

where (using a log approximation) $P_{icT}(D)$ can be written as

$$P_{icT}(D) = P_{ic,T-1}\left[1 + \widehat{\Delta p}_{icT}(D)\right].$$

If we assume a constant elasticity of demand system, the demand for good i in city c in time t in counterfactual D is given by

$$Q_{ict}(D) = \frac{(P_{ict}(D)/\varphi_{ic})^{-\eta}}{[P_{ct}(D)]^{1-\eta}} E_c, \tag{4}$$

where $P_{ct}(D)$ denotes the city price index whose formula is given by

$$P_{ct}(D) \equiv \left[\sum_i [P_{ict}(D)/\varphi_{ic}]^{1-\eta}\right]^{\frac{1}{1-\eta}}.$$

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

Following Jensen, we consider a set of price changes that are consistent with equation (3) but imply no net quantity changes. The aggregate movement in quantity is simply given by the sum of quantity changes for each city, i.e., $\Delta Q_{iT}(D) = \sum_c \Delta Q_{icT}(D)$. The counterfactual change in urban consumption in any city c in time T is given by $\widehat{\Delta Q}_{icT}(D) \equiv Q_{icT}(D) - Q_{ic,T-1}$. Using the log change as an approximation for the percentage change, we can rewrite this as

$$Q_{icT}(D) = Q_{ic,T-1}\left[1 + \widehat{\Delta q}_{icT}(D)\right], \tag{5}$$

where equation (4) implies that

$$\widehat{\Delta q}_{icT}(D) = (\eta - 1)\Delta p_{cT}(D) - \eta\widehat{\Delta p}_{icT}(D) + \Delta \ln E_{cT},$$

If we make the partial equilibrium assumption that when we use counterfactual price changes ($\widehat{\Delta p}_{icT}(D)$), aggregate urban expenditures, and prices are unaffected by price arbitrage, this equation reduces to

$$\widehat{\Delta q}_{icT}(D) = -\eta\widehat{\Delta p}_{icT}(D). \tag{6}$$

Substituting equation (6) into equation (5) and summing produces

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

If we then substitute equation (3) into equation (7), we obtain

$$\sum_c Q_{ic,T-1}\left(\hat{\gamma}+\hat{\delta}_1 x_i+\hat{\delta}_2 D x_i\right)\left(p_{ic,T-1}-p_i^*\right)=0,$$

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

$$p_i^*=\frac{\sum_c Q_{ic,T-1}\left(\hat{\gamma}+\hat{\delta}_1 x_i+\hat{\delta}_2 D x_i\right)p_{ic,T-1}}{\sum_c Q_{ic,T-1}\left(\hat{\gamma}+\hat{\delta}_1 x_i+\hat{\delta}_2 D x_i\right)}.$$

The partial equilibrium welfare gain arising from prices moving from their values in $T-1 \quad (P_{icT-1})$ towards their steady state values (P_i^*) can be written as

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

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 good's prices not equaling their steady state levels equals

$$\Delta W_{iT}\left(D\right)=\frac{1}{2}\sum_c\left(2P_{ic16}+\Delta P_{icT}\left(D\right)\right)\Delta Q_{icT}\left(D\right)-\underbrace{m_i\sum_c\Delta Q_{icT}\left(D\right)}_{=0}.$$

We can rewrite this in terms of welfare gain as a share of expenditures:

$$\Delta\hat{W}_{iT}\left(D\right)=\frac{\Delta W_{iT}\left(D\right)}{E_{iT-1}\left(D\right)}.$$

The welfare gain in the economy is given by

$$\Delta W_T\left(D\right)=\sum_i\Delta\hat{W}_{iT}\left(D\right)\frac{E_{iT-1}\left(D\right)}{E_{T-1}\left(D\right)}.$$

The gain due to e-commerce for each item 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

$$\Delta\hat{W}_{iT} = \Delta\hat{W}_{iT}\left(D=1\right) - \Delta\hat{W}_{iT}\left(D=0\right),$$

and the economy-wide welfare gain due to arbitrage is given by

$$\Delta W_T = \Delta W_T\left(D=1\right) - \Delta W_T\left(D=0\right). \tag{8}$$

3 Data

A major advantage of using Japanese data is that one can obtain measures of consumer expenditures by product and type of sales merchant. The National Survey of Family Income and Expenditures (NSFIE) is a representative survey of households with two or more members that records expenditures by product from each major retail outlet store type: small retail, supermarket, convenience, department, club, discount, catalog, internet, and "other".[2] 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 due to questions about which retail outlet they used. Fortunately, the Ministry of Economy Trade and Industry (METI) reports very reliable aggregate estimates of sales by e-commerce and other retailers by surveying 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 also conduct a robustness check for all of our main results using data from Rakuten, available from Online Appendix C. Rakuten is the largest e-retail company in Japan, who provided us with 2010 internet sales data (aggregated across buyers and merchants) for each of approximately

[2]Services and goods not sold by retailers (e.g., cars) are not included in the survey.

40,000 goods categories or "genres." In that year, Rakuten had a 30 percent market share of all Japanese e-commerce.[3] We then matched these genres to the expenditure categories in the 2010 Japanese Family Income and Expenditure Survey (FIES) that are used in to construct the Japanese consumer price index. This generated a matched sample in which we have 312 goods in a typical year, which we use in our main specifications.

We construct e-commerce intensity of an expenditure category by comparing the average household total expenditure on that category with the average household's online expenditure on it. We measure total expenditure share e_i on category i by using national average expenditures per household in 2009 taken from the Family Income and Expenditure Survey ("FIES", which forms the basis of the Japanese CPI and is distinct from the NSFIE). We denote online expenditure share in category i from NSFIE by s_i. We then define e-commerce intensity x_i of category i by taking the ratio of the online to total expenditures share, normalized by the maximum value of this ratio:

$$x_i = \frac{s_i}{e_i} / \max_j \left(\frac{s_j}{e_j} \right).$$

Thus, our measure of e-commerce intensity takes on a value of zero if there are no online transactions involving an expenditure category and a value of 1 if the online expenditure relative to those in the economy is the highest among all categories of goods. Expressing e-commerce intensity this way makes our e-commerce intensity (x_i) invariant to the size of sector i.

In order to see how e-commerce intensity varies across products, we aggregated the FIES codes into some broader categories in Table 1 so that we could 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 ($E_\ell \equiv \sum_{i \in \Omega^\ell} e_i / \sum_j e_j \times 100$), where Ω^ℓ is the set of items in some more aggregated category ℓ. In the second column,

[3] Rakuten, Inc. (2010) *Annual Report.*

we report the percentage of online expenditure in 2009 that corresponds to that category (($S_\ell \equiv \sum_{i \in \Omega^\ell} s_i / \sum_j s_j \times 100$), where s_i is online expenditure from NSFIE). The third column reports the "e-commerce intensity" in 2009, which we define at the aggregate level (i.e., $x_\ell \equiv S_\ell / E_\ell / [\max_k \{S_k / E_k\}]$).

Table 1 makes clear some basic stylized facts of our data. First, within goods categories we see that there are no zeros except newspapers and magazines in the table indicating that at this level of aggregation all categories of goods were available online in 2009. Second, 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), restricted/time-sensitive items (e.g., medicine and physical newspapers), and high weight-to-value items (non-perishable groceries) 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 furniture and furnishings—are sold very intensively online. Interestingly, we see that domestic utensils, household consumables (which includes non-durable household supplies like paper products and cleaning agents), and recreational goods (which includes items like sports equipment and gardening supplies) are sold very intensively online as well. Third, 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.

One of the other striking features of the table is 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 interesting to note that goods sold intensively online tend

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

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

Data source: FIES, NSFIE, Rakuten, and authors' calculation. Notes: The first column is from FIES. Column 2- 6 are from NSFIE and the last column is from Rakuten. Notes: 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 is calculated as $x_i = \frac{s_i}{e_i}/max_j(\frac{s_i}{e_i})$.

to have characteristics that are similar to those goods historically available in catalogs. For example, newspapers, meat, dairy, products, and fruits and vegetables do not sell much online, but also did not sell much in catalogs in 1999. Similarly, electronics was one of the products that was sold intensively in catalogs.

Although e-commerce was small in Japan in 1999, the situation changed radically over the next two years. Although there is no data on when city obtained broadband internet, every prefecture in Japan had DSL and cable TV connections available by 2001, with the number of household DSL contracts rising from 70,655 in the first quarter of 2001 to 13,325,408 by the end of 2004.[4] By April of 2000, when Rakuten announced its initial public offering and a year before the entry of Amazon into Japan, Rakuten 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.[5] 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. By 2017, e-commerce firms accounted for 5.8 percent of Japanese retail sales or about ¥16.5 trillion (about $149 billion). Despite the explosive growth, as one can see in the last column of Table 1, the set of goods selling well on the Rakuten platform remained remarkably similar to those that sold well in catalogs. Moreover, the Rakuten sales intensities are highly correlated ($\rho = 0.57$) with the e-commerce sales intensities we obtained from the NSFIE data, which suggests that these datasets are in broad agreement as to what

[4]Source: Ministry of Public Management, Home Affairs, Posts and Telecommunications. We thank Takashi Unayama for providing us with these numbers. In 2005, there were 49,063 thousand households in Japan, and 17,220 thousand among them are elderly households in which every member of households are 65 years of age or over.

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

goods sell well online.

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 of the products in the FIES categories. These prices are sampled in all cities that are either a prefectural government or have population of 150,000 or more, which gives us the ability to not only tracking product prices across time but also across space. This information typically identifies the brand of an item or 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, there is a tendency to select the same products by, for example, always picking the largest selling item within a sampling frame if available. Thus, while historical US CPI data is based on different baskets of goods in different cities, the JSB's "purposive" sampling generate samples in which the same good or very similar goods are sampled in different cities. Therefore, it is reasonable to believe that intercity prices are informative about true price differences across locations.[6]

One problem in the data is that we have periodic product substitutions that arise as goods are added to or dropped from the CPI sample. Fortunately, we have official quality-adjusted price quotes for Tokyo computed by the JSB[7], which we use to adjust the prices in other cities. This procedure is equivalent to assuming that the quality change associated with a product substitution in the CPI is identical across cities.

Another 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

[6]In order to further clean the data, we drop all observations in which the item only appears in one city. We also trimmed 3 smallest and 3 largest price quotes within an item-year observation. Finally, we dropped the bottom and top 1% of log price changes.

[7]Source: http://www.e-stat.go.jp/SG1/estat/List.do?bid=000001033703&cycode=0, accessed on April 5th, 2017.

quality in bar-code data is 0.9, so we should expect sampling problems to produce greater levels of price dispersion in JSB data is likely to be larger than that in barcode data because variation in JSB data 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.

Table 2 reports the sample statistics for our data. As one can see from the table, we have more than 100,000 price quotes in each of our five-year periods since e-retail has become available in Japan in 1997. The first line of the table shows the reports information on the e-commerce intensity of the goods in our sample (x_i). The values of x_i across goods tell us about the relative importance of online sales. Here we see that goods in the the upper 90^{th} percentile of the distribution have an e-commerce sales intensity of 0.157 over the full sample period, which is more than seven times higher than a good with the median intensity. Moreover, at the upper tail of the distribution, we observe goods with an e-commerce intensity that is more than 20 times higher than that of the median good. These summary statistics reflect 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. The second line of Table 2 presents catalog sales intensity in 1999. Similar to e-retail sales intensity, the distribution of catalog intensity shows left skewness. From the third line of the table, it shows average annual rate of inflation across the 5-year

sample period from 1991 to 2016. As one can see, on average goods prices fell before 2011, which reflects the deflation that can be observed in Japan over this time period.

Table 2: Summary Statistics for the Sample of Goods

	Mean	Standard Deviation	Min	p10	p50	p90	Max
$x_{i(t=2009)}$	0.057	0.073	0.000	0.004	0.022	0.157	0.456
$xcat_{i(t=1999)}$	0.078	0.119	0.000	0.004	0.035	0.207	0.995
Period: 1991 to 1996							
$\Delta^1 p_{ict}$	-0.004	0.099	-0.858	-0.101	0.000	0.090	0.953
Observations	74,732						
Period: 1996 to 2001							
$\Delta^1 p_{ict}$	-0.008	0.106	-1.124	-0.117	-0.001	0.084	1.165
Observations	109,486						
Period: 2001 to 2006							
$\Delta^1 p_{ict}$	-0.009	0.114	-1.798	-0.127	-0.003	0.104	1.679
Observations	163,473						
Period: 2006 to 2011							
$\Delta^1 p_{ict}$	-0.001	0.124	-1.695	-0.127	0.000	0.126	1.556
Observations	164,029						
Period: 2011 to 2016							
$\Delta^1 p_{ict}$	0.014	0.100	-1.276	-0.084	0.010	0.122	1.092
Observations	168,241						

Data source: RPS, NSFIE, and authors' calculation. Notes: This table shows summary statistics of of price changes, e-commerce intensity, and catalog intensity for five five-year-periods from 1991 to 2016. Prices are in natural log. $\Delta^1 p_{ict}$ is the one-year log difference in prices;$x_{i(t=2009)} = [s_i/e_i]/[\max_j(s_j/e_j)]$ shows e-commerce intensity in 2009 using NSFIE and $xcat_{i(t=1999)}$ indicates catalog sales intensity in 1999.

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 as opposed to absolute price declines of online goods 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 inflation rates that were about 1.3 percentage point lower than goods traded offline. Here, we present that goods traded online have lower price increase rates than goods not available online. We also find that differential rates of price increase were present long before the entry of e-commerce firms, became more pronounced after the entry of e-commerce merchants, and 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}) on good (α_i) and city (β_c) fixed effects along with an indicator variable, D_t, that is one starting in 1997 (the year Rakuten opened) and zero before as well as the e-commerce intensity of the good interacted with this dummy ($x_i D_t$):

$$\Delta p_{ict} = \alpha_i + \beta_c + \phi D_t + \theta x_i D_t + \epsilon_{ict}, \tag{9}$$

where α_i is a parameter to capture any pre-trends in the data that might arise if goods available online exhibit have different price increase trends than goods not available online. The coefficient on D_t (ϕ) tells us whether there was any differential trend in price increase for goods available online

after the entrance of e-commerce firms and θ, the coefficient on the e-commerce intensity interaction term ($x_i D_t$) tells us about the differential rate of price change for goods traded online after the entry of e-commerce firms. We do this for two time periods (1992-2001) and (1992-2016) to see if there is any difference in the results we obtain by looking at the period immediately after the entry of e-commerce firms versus the full time period.

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 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_i) in the bottom quartile, and 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:

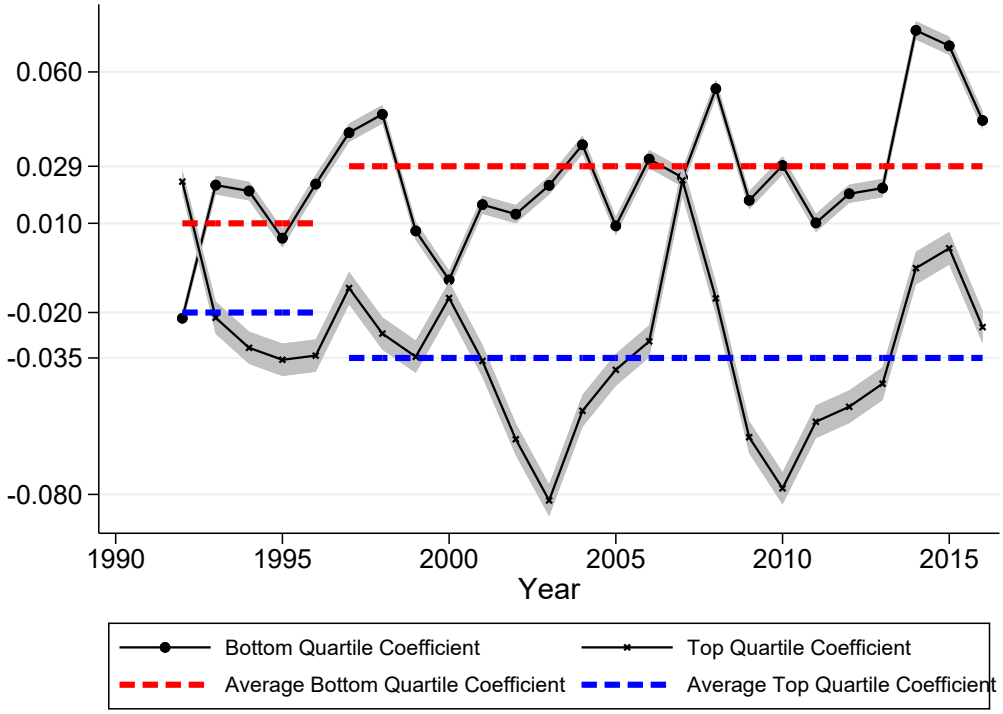
$$\Delta p_{ict} = \theta_t + \epsilon_{ict}, \tag{10}$$

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. As the figure makes clear, there are unmistakable pre-trends in the data.[8] 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 -2.0 percent per year, while the average annual rate of price increase for goods that not sold much on these platforms was 1.0 percent per year. Thus, even before the entry of e-commerce firms, there was a 3 percentage point gap

[8]In online appendix Figure A2, we also present unweighted estimates, which show a similar pattern. The relatively low bottom-quartile point estimate for 1992 is driven by fruits and vegetable prices in 1992 as can be seen by looking at online appendix Figure A1. 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.

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



Data source: RPS, NSFIE, and authors' calculation. Notes: This black line shows the average rate of expenditure-weighted price increase for the goods in two groups: products with bottom quartile e-commerce sales intensity (black line with dot) and products with top quartile e-commerce intensity (black line with symbol x). 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.

between the relative rates of price increase for 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 the mid-2000s, when e-commerce firms became major players in Japanese retail. The relative inflation rate for goods sold heavily online relative to goods not sold much online fell substantially. Goods in the top quartile of e-commerce sales intensity had an average rate of price growth from 1997 to 2016 of -3.5 percent per year: a 1.5 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.

Table 3: Relative Price Changes

	(1) $\Delta^1 p_{ict}$	(2) $\Delta^1 p_{ict}$	(3) $\Delta^1 p_{ict}$	(4) $\Delta^1 p_{ict}$
D_t	-0.0001	0.0097***	-0.0001	0.0142***
	(0.0021)	(0.0022)	(0.0038)	(0.0030)
E-Commerce Intensity	-0.0069	-0.0906***	-0.0049	-0.1636***
$\times D_t$	(0.0246)	(0.0200)	(0.0668)	(0.0411)
Sample	Goods	Goods	Goods	Goods
Fixed Effects	Product	Product	Product	Product
Estimation Period	1992-2001	1992-2016	1992-2001	1992-2016
Observations	152,416	393,246	150,417	387,918
R^2	0.04	0.03		
E-Commerce Intensity Year	2009	2009	2009	2009
First-Stage F-Stat			57.33	58.26
Estimation Method	OLS	OLS	IV	IV

Data source: RPS, NSFIE, and author's calculation. Notes: 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 OLS estimates using e-commerce sales intensity and the second two columns use catalog sales intensity as IV.

Turning to our differences-in-differences specification, we present the results from estimating equation (9) in Table 3. The first column presents the results from estimating equation (9) 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. However, as one can see in column 2, we do see a significant decline in the *relative* prices of goods available online as evidenced by the coefficient of -0.09 on the post-e-commerce e-commerce intensity interaction ($x_i D_t$) term.

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 endogeneity, we construct a variable, catalog intensity, which is constructed analogously to e-commerce intensity except that we use catalog sales instead of e-commerce sales. Unfortunately, the earliest year for which we have catalog sales data is 1999, but since Rakuten was only two-years old in 1999 and still a small company and major players in the e-commerce market like Amazon Japan had not even entered the Japanese market, we think it plausible to argue that the distribution household catalog purchases in 1999 were unlikely very different than those before the entry of Rakuten.

Table 4 reports the validity of our instrumental variable, catalog sales intensity, 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. This establishes that as long as historical catalog sales were not being driven by the anticipation of e-commerce, we have a powerful instrument for e-commerce sales intensity.

We report the results from our instrumental variables (IV) estimation in columns 3-4 of Table 3. 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 ($x_i D_t$) on price increases doubles in magnitude in the full sample estimates (columns 2 and 4). The fact that the OLS estimates are attenuated

Table 4: First Stage Regression

	(1) E-Commerce Intensity $\times D_t$	(2) E-Commerce Intensity $\times D_t$	(3) E-Commerce Intensity
Catalog Intensity$\times D_t$	0.7363***	0.7304***	
	(0.0972)	(0.0957)	
D_t	0.0237***	0.0241***	
	(0.0043)	(0.0043)	
Catalog Intensity			0.7457***
			(0.0745)
Constant			0.0269***
			(0.0051)
Sample	Goods	Goods	Goods
Fixed Effects	Product	Product	None
Estimation Period	1992-2001	1992-2016	
Observations	150,417	387,918	306
R^2	0.26	0.26	0.25
E-Commerce Intensity Year	2009	2009	2009
First-Stage F-Stat	57.33	58.26	
Estimation Method	IV-First Stage	IV-First Stage	OLS

Data source: NSFIE and author's calculation. The table shows the first stage regression results.

implies that e-commerce firms tended to enter sectors where prices were rising, perhaps because these markets were likely to be more profitable. This pattern of behavior would explain why estimates that do not control for the endogeneity of market entry are likely to underestimate the the relative impact of e-commerce on pricing. In terms of economic significance, the results in column 4 imply that a good at the 90^{th} percentile of internet sales intensity had rates of price increase that were 2.4 percentage points per year lower than goods not sold online after the entry of e-commerce firms.

4.2 Gains Due to Price Arbitrage

As the last section made clear, while there is strong evidence that the rise of e-commerce caused the relative prices of goods sold online to decline

in Japan, there is little evidence that this caused the overall price level to fall because the the lower relative rate of price increase for goods sold intensively by e-commerce firms was in part due to higher rates of price increase for goods sold principally by physical merchants. However, there is an alternative mechanism through which e-commerce might affect prices along the lines suggested by Jensen (2007) and DellaVigna and Gentzkow (2017): namely e-commerce might force retailers to adopt uniform pricing across regions. This effect would be manifest in our data by an acceleration of price arbitrage across cities.

In order to visualize whether this is likely to be important, we first consider two five-year periods. The first five-year period (1991-1996), predates the formation of e-commerce by at least a year, so we can call this period the "pre-e-commerce period." We start the second period in 1996 because we assume that in 1996, the distribution of prices was reflective of a world without e-commerce but by 2001, Rakuten was already a prominent, listed company, with tens of millions of hits and thousands of stores selling on its platform.

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 normalized the data by regressing Δp_{ict} and p_{ict} on product and city fixed effects and construct normalized price changes ($\Delta^5 p_{ict} - \hat{\alpha}_{it} - \hat{\beta}_{ct}$) and normalized 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 prices remove the effect of any common price movements at the product or city level. Figure 3 presents plots of normalized five-year change in prices ($\Delta^5 p_{ict} - \hat{\alpha}_{it} - \hat{\beta}_{ct}$) against the normalized five-year lag of prices in each city ($p_{ic,t-5} - \hat{\alpha}'_{it-5} - \hat{\beta}'_{ct-5}$).

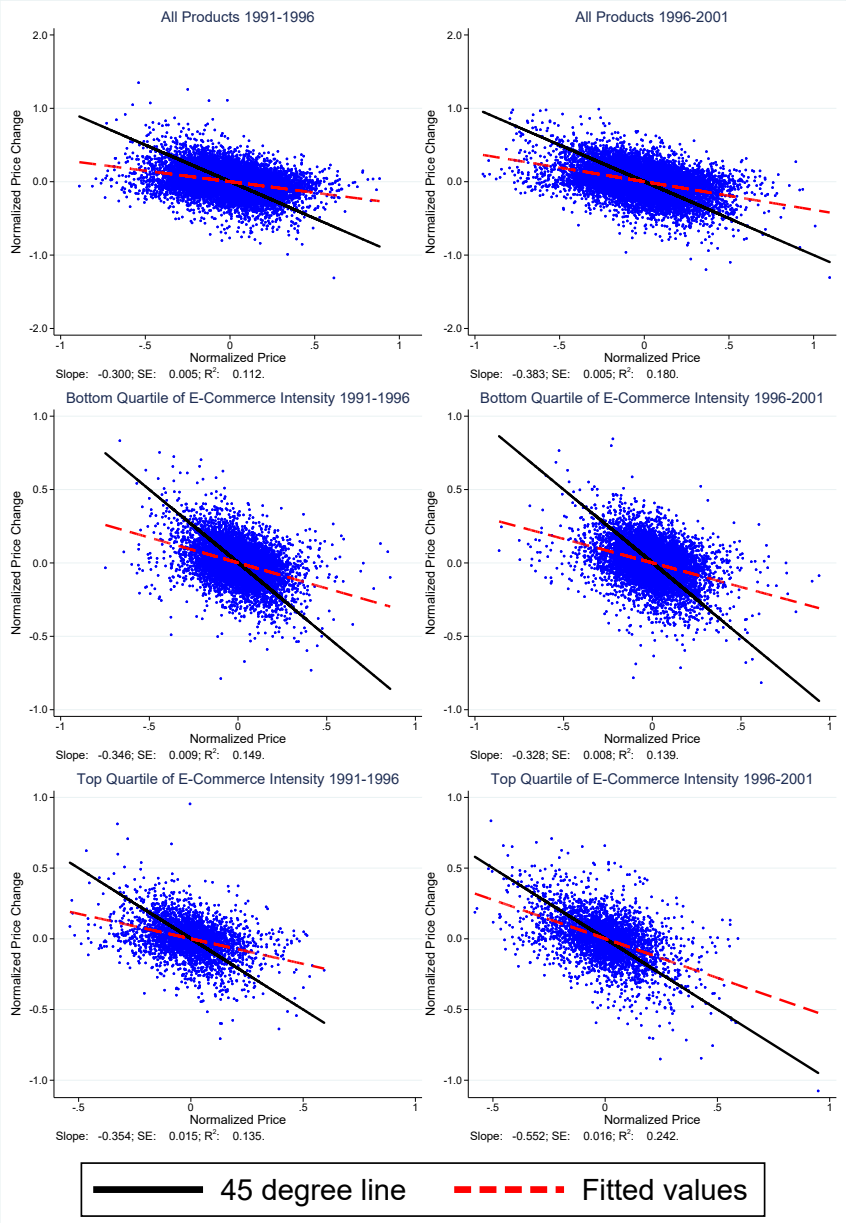
The first panel shows how normalized price changes vary with normalized prices before and after the entry of e-commerce. There is a clear negative relationship between initial urban price deviations and future price growth, which indicates that goods that had high prices in cities

tended to have lower rates of price increases 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 be eliminated within five years before the advent of e-commerce and this number rose 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.[9]

The next two pictures show what was 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 ($x_i < 0.013$) and the set of goods in the highest quartile of the distribution ($x_i > 0.076$). As one can see from the second 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 -0.29, which is almost identical to the slope in the pre-e-commerce period (-0.30). 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 66 percent, rising in magnitude from -0.29 to -0.48. 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 in Section 2.1.

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

Figure 3: Normalized Price Change vs. Normalized Price



Data source: RPS, NSFIE, and authors' calculation. Notes: This graph plots normalized price changes against normalized price levels. The left panel shows normalized price changes before the entry of e-commerce and the right panel shows them after the entry of e-commerce. The first panel plots for all goods, the second panel plots for goods with e-commerce intensity lower than the bottom quartile, and the third panel shows for goods with e-commerce intensity higher than the top quartile.

4.2.1 Estimating Convergence Rates

Following Rogoff (1996), we can test for whether we observe absolute price convergence or relative price convergence by estimating equation (1) and testing whether the estimated city-time fixed effects are jointly zero. If they are, then the data suggests that the prices of goods are converging to the same price across cities. Otherwise, it implies that the prices of goods converge to different levels in different cities. We can use an *F*-test to reject the hypothesis that the city-year fixed effects are zero, which suggests that absolute price convergence fails, so average price levels of goods do not converge to exactly the same level in all cities. We therefore city-time fixed effects in our main specifications.[10]

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 1 and 2 reveals the convergence rates for goods not suitable for e-commerce (i.e., those where $x_i = 0$) 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 ($x_i p_{ic,t-5}$) 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 to the entry of e-commerce. However, the negative and significant coefficient on the interaction term ($D_t x_i p_{ic,t-5}$) 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. In column 3, we estimate our baseline differences-

[10] Although we cannot statisitcally reject the hypothesis that the prices of goods on average converge to the same level in every city, our results also show that the estimated values of β_{ct} are very close to zero. We also report results without city-year fixed effects in an online appendix as a robustness check to show that their inclusion does not matter qualitatively for our results.

Table 5: Estimates Over Period 1991-2001

Dependent Variable	(1) Δp^5_{ict}	(2) Δp^5_{ict}	(3) Δp^5_{ict}	(4) Δp^1_{ict}
Lagged Price	-0.292***	-0.324***	-0.309***	-0.129***
	(0.032)	(0.037)	(0.031)	(0.013)
E-Commerce Intensity	-0.170		-0.039	0.032
(t=2009) × Lagged Price	(0.413)		(0.410)	(0.151)
E-Commerce Intensity		-1.145*	-1.292***	-0.509***
× Lagged Price × Post Rakuten		(0.592)	(0.315)	(0.130)
t	{1996}	{2001}	{1996,2001}	Annual 1992-2001
Observations	25,848	27,407	51,012	152,416
E-Commerce Intensity Year	2009	2009	2009	2009
First-Stage F-Stat	29.82	33.98	17.81	16.42
Estimation Method	2SLS	2SLS	2SLS	2SLS

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (1) using 2SLS: e-commerce sales intensity in 2009 is instrumented using 1999 catalog sales intensity. The first column uses the five-year log differences in prices from 1991 - 1996 and the second column uses that from 1996 - 2001. The third column uses two five-year periods, 1991 - 1996 and 1996 - 2001. The fourth column use the annual log differences in prices from 1992 to 2001. The OLS regression results are available in Appendix B from Table A1.

in-differences specification of equation (1) using a five-year difference by letting t take on two values: 1996 and 2001. The most important result for our purposes 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 ($\hat{\delta}_2 = -1.292$). The estimate of δ_1 is close to zero, which means that before the entry of e-commerce firms, there was no 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.

Table 6: Estimates Over Alternative Periods

Dependent Variable	(1) $\Delta^5 p_{ict}$	(2) $\Delta^5 p_{ict}$	(3) $\Delta^5 p_{ict}$	(4) $\Delta^5 p_{ict}$	(5) $\Delta^1 p_{ict}$
Lagged Price	-0.378***	-0.450***	-0.373***	-0.373***	-0.144***
	(0.028)	(0.032)	(0.030)	(0.026)	(0.013)
E-Commerce Intensity	0.676	1.337***	0.518	0.519	0.172
(t=2009) × Lagged Price	(0.415)	(0.439)	(0.423)	(0.403)	(0.165)
E-Commerce Intensity	-1.846***	-3.168***	-1.943***	-1.741***	-0.775***
× Lagged Price × Post Rakuten	(0.354)	(0.316)	(0.373)	(0.247)	(0.102)
t	{1996,2006}	{1996,2011}	{1996,2016}	{1996,2001, 2006,2016}	Annual 1992-2016
Observations	51,845	43,256	42,555	87,515	393,246
E-Commerce Intensity Year	2009	2009	2009	2009	2009
First-Stage F-Stat	15.03	23.61	19.13	18.02	27.17
Estimation Method	2SLS	2SLS	2SLS	2SLS	2SLS

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (1) using 1999 catalog sales intensity as an instrument for e-commerce sales intensity. First three columns compare the five-year period of log price differences before (1991-1996) and after the entry of e-commerce firms (2001-2006, 2006-2011, and 2011-2016). The fourth column uses 4 pooled five-year-period of log price differences from 1991 to 2016. The last column uses the annual frequency of log price changes. The OLS regression results are available in Appendix B Table A2.

We now turn to exploring the robustness of these results to alternative specifications. The estimates in Table 5 are likely to understate the impact of e-commerce because e-commerce firms were relatively small before 2001. In order to deal with this concern, Table 6 presents results in which we use alternative time periods. In the first three columns, we do a differences in differences based comparing the five years prior to the entry of e-commerce firms (1991-1996) with three alternative non-overlapping periods: 2001-2006, 2006-2011, and 2011-2016. These results demonstrate a similar to the one we observed in Table 3; the estimated effects of e-commerce on pricing are stronger after e-commerce firms has a chance to expand operations. The coefficient on the e-commerce triple interaction ($\hat{\delta}_2$) rises by over 40 percent if we compare rates of convergence between 2006 and 2011 with those of 1991-1996 instead of 1997-2001 with 1991-1996. We obtain a similar result when we repeat the estimation over the full period (1992-2016) at the annual

frequency.

Our results are economically significant as well. If we use the estimates in column 5 of Table 6 as a benchmark, we find that the half life for a relative price difference for a good not traded online is 4.5 years. By contrast, the half life for a good with maximal internet sales intensity is six months. The decreased half-life due to faster convergence rate of a good with the 90^{th} percentile of e-commerce sales intensity is 3 years, based on the equation (2). Thus, our estimates imply that the advent of e-commerce have significantly altered the ability of retailers to charge different prices in different cities.

The second concern that one might have with the 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 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 C Table A6 for over period 1991 - 2001 and Table A7 for alternative periods, which shows that we obtain very similar results regardless of whether we measure internet sales intensity using consumer expenditure data or Rakuten e-commerce sales data.

4.2.2 Welfare Gain

Aggregate consumer gains due to faster price convergence can be calculated from the equation (8). One of the interesting 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 on Table 6 column 5. 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. Hottman et al. (2016) obtain smaller estimates when working at the firm level—their 10-90 range is 2.6 to 7.3 with a median value of 3.9.[11] Table 7 presents that welfare gain arising from price arbitrage in 2013

[11] Broda and Weinstein (2010) estimate elasticities across different brands in 122 Nielsen retail product groups and obtain somewhat higher elasticities (a median of 7.5), but there

towards their steady state is 0.3 percent when the demand elasticity is -4. [12]

Table 7: Counterfactual Welfare Gain due to Price Arbitrage

Eta	ΔW_{14}	ΔW_{17}
3	0.0020	0.0017
4	0.0026	0.0023
5	0.0033	0.0029
6	0.0040	0.0035
7	0.0046	0.0041

4.3 Gains in "New Trade Models"

An alternative channel 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 Arkolakis et al. (2012). Suppose that the correct model of how retail operates is described by a Krugman (1980) or Melitz (2003) model in which firms in their model correspond to retail merchants who sell locally through physical stores or at a distance through catalogs or e-commerce operations. In this simple extension of Melitz (2003), we assume that manufactured goods are produced locally using only labor using a constant returns-to-scale technology, so the cost of producing a manufactured good in region j is $w_j a$, where w_j corresponds the wage in region j and a is the unit labor requirement. Manufacturers sell to local retailers who then can sell these goods locally or at at a distance. In order for a retailer to sell locally, it needs

is a lot of overlap. We chose the the lower elasticity estimates to be conservative about the magnitude of the welfare gains.

[12] To get some sense of how large this is, we can compare the gain to Brynjolfsson et al. (2003) estimate of the gains due to Amazon's entry into U.S. book market. That paper estimated a gain of less than 1 billion dollars in 2000—only 0.015 percent of U.S. personal consumption expenditures in that year.

to pay a fixed cost and incurs a marginal cost of sales equal to $(a/\varphi)\, w_j$, where φ is the productivity of the retailer. Similarly, in order to sell in a different location, a retailer needs to pay an addition "export" fixed cost and an iceberg transportation cost between regions i and j of τ_{ij}. It is immediately obvious that this cost structure is exactly that of Melitz (2003). Moreover, we can think of e-commerce and catalog sales as a technology that reduces the cost of trade between regions (τ_{ij}).[13]

If we think about e-commerce as a trade-facilitating technology, we can use the result in Arkolakis et al. (2012) to write the log change in welfare following a trade liberalization $\Delta W_t = \frac{1}{\epsilon} \ln(\lambda_t/\lambda_{t-k})$, where $\lambda_t \in (0, 1]$ is the share of consumer of expenditures on sales from retailers other than e-commerce firms in period t, and ϵ equals the "trade elasticity." To understand how this formula works, imagine that in the initial period $(t-k)$ consumers only purchase products locally, so $\lambda_{t-k} = 1$, but after the advent of e-commerce, consumers purchase ten percent of their goods online, so $\lambda_t = 0.9$. If we use a standard estimate of the trade elasticity of -3, we will obtain a welfare gain of 3.5 percent (=$-\ln(0.9)/3$). Since we observe e-commerce sales at the prefectural level in Japan, we can also use regional e-commerce sales shares to compute gains for each prefecture. Moreover, the formula can easily be adapted to account for catalog sales. In order to account for catalog sales, we simply define λ_t to be the share of expenditures on products not sold by catalogs or e-commerce firms, i.e., local physical stores and use the same formula. This is a very simple way to incorporate the fact that e-commerce gains may higher or lower once we take into account catalog sales.

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. Based on the NSFIE data, we know that the share of household

[13] In order for the counterfactual to be exact, we also need to assume that the advent of e-commerce does not cause labor to move across regions. However, given the abundant evidence of sluggish migration across regions even in the presence of large shocks and the likely small impact of e-commerce on relative wages, we think that this assumption is a reasonable approximation.

expenditures purchased from all retailers (χ) was 0.62 in 2014, with the remaining expenditures covering utilities, education, and other expenditure items that we will assume are not affected by e-commerce's entry into the goods sectors. In 2014, e-commerce expenditures on goods as a share of all retail expenditures, which we denote by s, was 0.0437. The share of household expenditures from non-e-commerce firms in 2014 is therefore $\lambda = (1-s)\,\chi + 1 - \chi = 0.97$. Assuming a trade elasticity of -3, this gives us an estimate of the welfare gain from e-commerce in Japan in 2014 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 welfare gains range from 0.5 percent to 1.4 percent in 2014 and from 0.6 percent to 1.8 percent in 2017. The higher numbers in later years reflect the fact that e-commerce sales have continued to expand rapidly in Japan.

Table 8: National Welfare gain due to increase variety

Epsilon	ΔW_{14}^{E}	ΔW_{14}^{C}	ΔW_{14}^{T}	ΔW_{17}^{E}	ΔW_{17}^{C}	ΔW_{17}^{T}
-2	0.014	0.001	0.015	0.018	0.001	0.020
-3	0.009	0.001	0.010	0.012	0.001	0.013
-4	0.007	0.001	0.008	0.009	0.001	0.010
-5	0.005	0.001	0.006	0.007	0.001	0.008
-6	0.005	0.000	0.005	0.006	0.000	0.007

Data source: NSFIE, JSB, MIETI and authors' calculation. Notes: The first three columns show welfare gain due to increased variety from e-commerce, catalog, and both e-commerce and catalog in 2014 along with plausible trade elasticities. The last three columns show them in 2017.

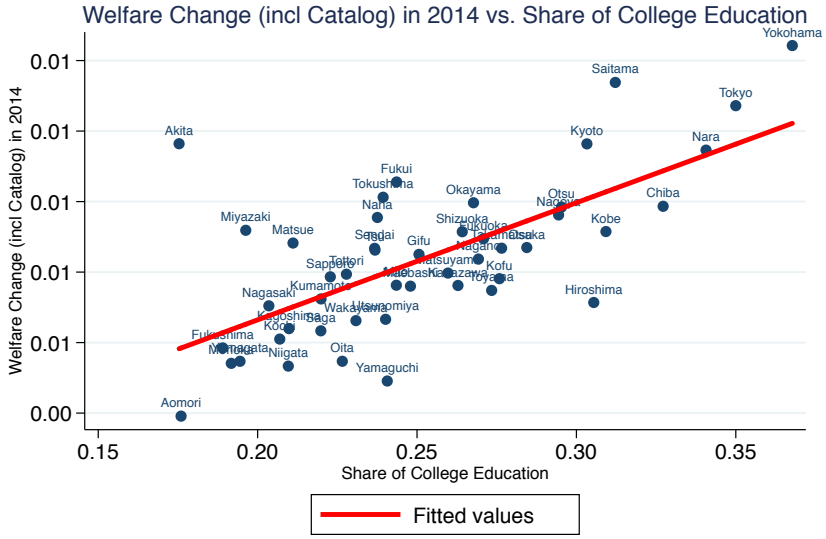
The second two columns in Table 8 make define local sales as total expenditures less expenditures on products sold over the internet or through catalogs. Interestingly, we see that aggregate welfare gains, shown in the third column,appear to be higher when we allow for the fact that consumers purchased and continue to purchase goods through catalogs. The mechanical reason for this result is that the share of consumer expenditures through catalogs actually grew slightly between 1996 and

2014. While one might have thought that the growth of e-commerce would have led to lower catalog sales because e-commerce is a good substitute for catalogs, there are a number of reasons why they may have grown together. First, the remarkable reduction in the costs of data transmission through the internet also occurred at a time when it became significantly cheaper to obtain and use phones. Thus, reductions in telecommunications costs may have benefited both catalog and e-commerce merchants. Second, e-retailers often advertise their wares in catalogs, and as e-commerce firms grew, they may have expanded catalog mailings, which may have caused both to rise. Nevertheless, the rise in catalog sales was quite small, so adjusting for it raises the welfare gains by around 0.2 percentage points.

This number is substantially larger than the structural approach we applied in Section 4.2.2 and reflects the impact of different modeling assumptions. A major advantage of the Arkolakis et al. (2012) approach is that it corresponds exactly to the gains implied by "core" trade models and the calculation takes into account general equilibrium effects. A disadvantage of this approach is that these models impose a number of assumptions that may not exactly fit the data: iceberg transportation costs, balanced trade, no labor mobility, monopolistic or perfect competition, aggregate profits being a constant share of revenues, and so on. A second potential problem with variety-based approaches to modeling the internet is that they are "hard-wired" to produce welfare gains as long as e-commerce shares (or e-commerce and catalog shares) rise everywhere. The major advantage of the Jensen (2007) approach used in Section 4.2.2 is that it not based on all of these identifying assumptions and as a result produces different estimates and allows for the possibility that e-commerce may not benefit everyone, but the main disadvantage is that it does not take into account general equilibrium forces that might also matter for welfare. Since it is difficult to say which approach is most plausible, we simply note that reasonable estimates of the percentage welfare gains from e-commerce in 2014 range from 0.5 percent to 0.76 percent,when demand elasticity is -4.

Now let see how these gains have affected individual prefectures. Since

Figure 4: Welfare gain due to increased variety vs. Share of College Education



Slope: 0.033; SE: 0.005; R^2: 0.464.

Data source: NSFIE, JSB, MIETI and authors' calculation. Notes: Figure plots the welfare gain due to the increased variety from both e-commerce and catalog sales in 2014 against the share of college education. The number of observations is 47.

there are 47 prefectures in Japan (which are similar in size to U.S. counties), we do not present for prefectures individually, but instead look for patterns in the data. One of the strongest patterns arises from the fact that in Japan, the share of online purchases is strongly associated with college education, which means that the share of prefectural expenditures on e-commerce has a strong positive correlation with the share of college educated people in a prefecture. This produces a strong positive correlation (0.33) between the welfare gain in a prefecture and the share of college-educated people in the prefecture. We can see this clearly in Figure 4, where we plot the welfare gain (including both e-commerce and catalog sales in our definition of non-local sales). This result suggests a digital divide in which regions with a large share of highly educated people benefit more than regions with fewer highly educated people is present.[14]

[14]Whether we include or exclude catalog sales does not matter substantively for our prefectural results. If we define λ_{p14} without counting catalog sales as non-local

Table 9: Prefecture Welfare Change

	(1) ΔW_{14}	(2) ΔW_{14}	(3) ΔW_{14}	(4) ΔW_{14}	(5) ΔW_{14}	(6) ΔW_{14}
Share of College Educated	0.0325***	0.0278***	0.0357***	0.0300***	0.0205**	0.0240***
	(0.0053)	(0.0068)	(0.0064)	(0.0068)	(0.0078)	(0.0087)
Population		0.0000				0.0000
		(0.0000)				(0.0000)
Income per Capita			-0.0000			-0.0000
			(0.0000)			(0.0000)
Average Age				-0.0001		0.0000
				(0.0002)		(0.0002)
Share of Secondary Educated					-0.0145**	-0.0140
					(0.0071)	(0.0087)
Constant	-0.0002	0.0006	0.0008	0.0057	0.0088*	0.0087
	(0.0013)	(0.0016)	(0.0018)	(0.0100)	(0.0046)	(0.0101)
Observations	47	47	47	47	47	47
R^2	0.460	0.474	0.469	0.464	0.506	0.528

Data source: NSFIE, JSB, MIETI and authors' calculation. The table shows how prefectural welfare gains due to increased variety relate to characteristics of prefecture - share of the college education, population, income per capita, and share of secondary education.

One obvious concern with these results is that the share of college educated people might be correlated with other factors that matter for internet purchases. For example, Einav et al. (2017) document that e-commerce in the U.S. is positively associated with city size. Alternatively, it may be the case that income or age may be associated with e-commerce intensity. In order to understand the importance of these factors, we regressed 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 Table 9. We find that none of these variables are significant once we control for the share of college educated people in a

expenditures, we obtain the same correlation up to two significant digits, and the plot looks very similar.

prefecture. When we include the share of secondary-school graduates, we find that it is significant in one specification, but it has a negative sign, which reinforces our earlier point that it is highly educated people that are the main users of e-commerce. In fact, most of the coefficients are precisely estimated zeros.[15] These differences are economically significant. The gain for Tokyo (the prefecture with the highest share of college-educated people) is four times that of Aomori, which has the lowest gain and has a share of college-educated people that is half that of Tokyo.

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 consumer welfare. While we find that at the national 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. However, part of the reason for the increasingly different price trends is due to the fact that the rate of price increase for goods sold intensively in physical stores rose, which underscores the difficulty in interpreting the relatively low rates of price increase for goods sold online as providing information about how e-commerce affects relative prices

At the local level, we find strong evidence that the rate at which intercity price differences disappear significantly for goods sold intensively online after e-commerce sales became common in Japan. 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:

[15]This 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.

Japanese welfare in 2014 was 0.3 percent higher as a result of e-commerce.

When we examine the robustness of these results by calibrating new-trade theory models which control for general equilibrium forces and consider welfare gains through variety expansion, we find a similar overall gains from e-commerce—a welfare rise of 0.5 percent, and we find the same pattern of highly educated regions benefiting 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 four times higher than in small cities. This result arises from the fact that higher-educated consumers buy substantially more online than less-educated consumers. Thus, while the level of the gains varies depending on the modeling framework adopted, the core result that e-commerce has a differentially positive effect for cities with a large share of high-income or highly-educated people remains.

References

Jenny C. Aker. Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics*, 2:46–59, 2010.

Treb Allen. Information Frictions in Trade. *Econometrica*, 82(6):2041–2083, 2014.

Costas Arkolakis, Arnaud Costinot, and Andrés Rodríguez-Clare. New Trade Models, Same Old Gains? *American Economic Review*, 102(1):94–130, 2012.

Paul R. Bergin, Reuven Glick, and Jyh-Lin Wu. "Conditional PPP" and Real Exchange Rate Convergence in the Euro Area. *Journal of International Money and Finance*, 73:78–92, 2017.

Christian Broda and David E. Weinstein. Understanding International Price Differences Using Barcode Data. *NBER Working Paper 14017*, 2008.

Christian Broda and David E. Weinstein. Product Creation and Destruction: Evidence and Price Implications. *American Economic Review*, 100(3):691–723, 2010.

Jeffrey R. Brown and Austan D. Goolsbee. Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry. *Journal of Political Economy*, 110(3):481–507, 2002.

Erik Brynjolfsson, Yu (Jeffrey) Hu, and Michael D. Smith. Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers. *Management Science*, 49(11):1580–1596, 2003.

Alberto Cavallo. Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers. *America Economic Review*, 107(1):283–303, 2017.

Alberto Cavallo. More Amazon Effects: Online Competition and Pricing Behaviors. *NBER Working Paper Series, W25138*, 2018.

Stephen G. Cecchetti, Nelson C. Mark, and Robert J. Sonora. Price Index Convergence Among United States Cities. *International Economic Review*, 43(4):1081–1099, 2002.

Victor Couture, Benjamin Faber, Yizhen Gu, and Lizhi Liu. E-Commerce Integration and Economic Development: Evidence from China. *NBER Working Paper 24384*, 82(6):2041–2083, 2018.

Stefano DellaVigna and Matthew Gentzkow. Uniform Pricing in US Retail Chains. *NBER Working Paper Series, W23996*, 2017.

Liran Einav, Peter J. Klenow, Benjamin Klopack, Jonathan D. Levin, Larry Levin, and Wayne Best. Assessing the Gains from E-Commerce. *Stanford University, mimeo*, 2017.

Jingting Fan, Lixin Tang, Weiming Zhu, and Ben Zou. The Alibaba Effect: Spatial Consumption Inequality and the Welfare Gains from E-commerce. *Journal of International Economics*, 114:1580–1596, 2018.

Robert C. Feenstra. New Product Varieties and the Measurement of International Prices. *American Economic Review*, 84(1):157–177, 1994.

Caroline L. Freund and Diana Weinhold. The effect of the Internet on international trade. *Journal of International Economics*, 62:171–189, 2004.

Maris Goldmanis, Ali Hortaçsu, Chad Syverson, and Önsel Emre. E-Commerce and the Market Structure of Retail Industries. *The Economic Journal*, 120:651–682, 2010.

Austan D. Goolsbee and Peter J. Klenow. Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce. *NBER Working Paper*

24649, 2018.

Aparajita Goyal. Information, Direct Access to Farmers, and Rural Market Performance in Central India. *American Economic Journal: Applied Economics*, 2:22–45, 2010.

Jessie Handbury and David E. Weinstein. Goods Prices and Availability in Cities. *Review of Economic Studies*, 82:258–296, 2015.

Colin J. Hottman, Stephen J. Redding, and David E. Weinstein. Quantifying the Sources of Firm Heterogeneity. *The Quarterly Journal of Economics*, pages 1291–1364, 2016.

Robert Jensen. The Digital Provide: Information (Technology), Market Performance, and Welfate in the South Indian Fisheries Sector. *The Quarterly Journal of Economics*, 122(3):879–924, 2007.

Paul Krugman. Scale Economies, Product DiÂ§erentiation, and the Pattern of Trade. *American Economic Review*, 5(70):950–959, 1980.

Marc J. Melitz. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, 6(71):1695–1725, 2003.

Parmy Olson. Mission Impossible: How Rakuten Billionaire Hiroshi Mikitani Plans To 'Beat Amazon'. *Forbes*, September 2012.

David C. Parsley and Shang-Jin Wei. Convergence to the Law of One Price Without Trade Barriers or Currency Fluctuations. *The Quarterly Journal of Economics*, 111(4):1211–1236, 1996.

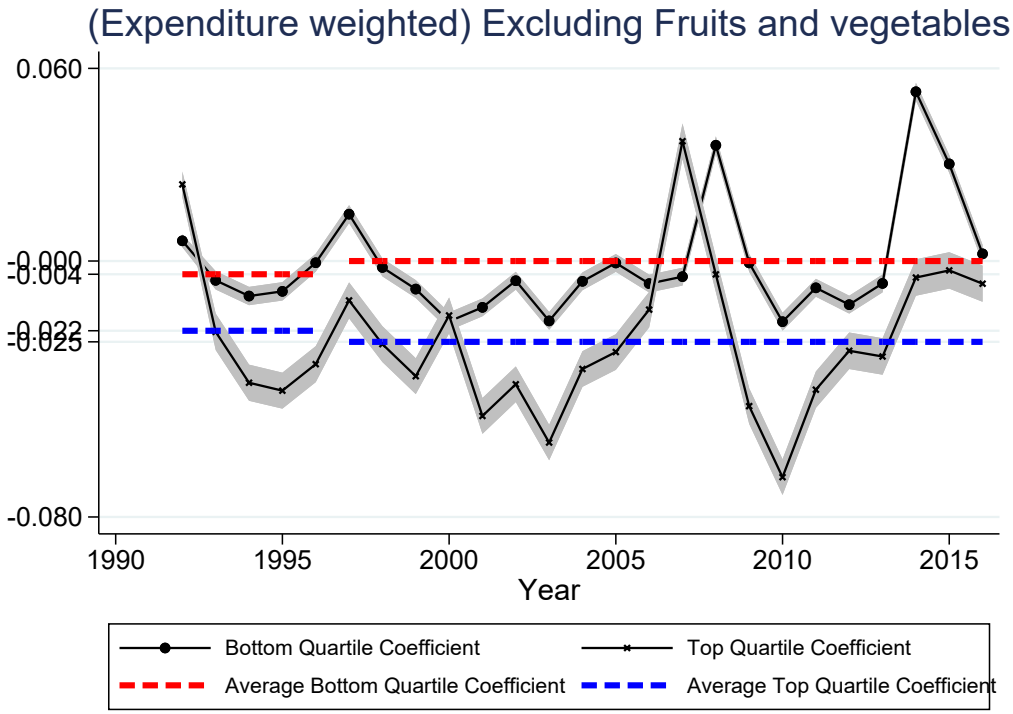
Kenneth Rogoff. The Purchasing Power Parity Puzzle. *Journal of Economic Literature*, 34(2):647–668, 1996.

Claudia Steinwender. Real Effects of Information Frictions: "When the States and the Kingdom became United". *American Economic Review*, 108 (3):657–696, 2018.

Appendix

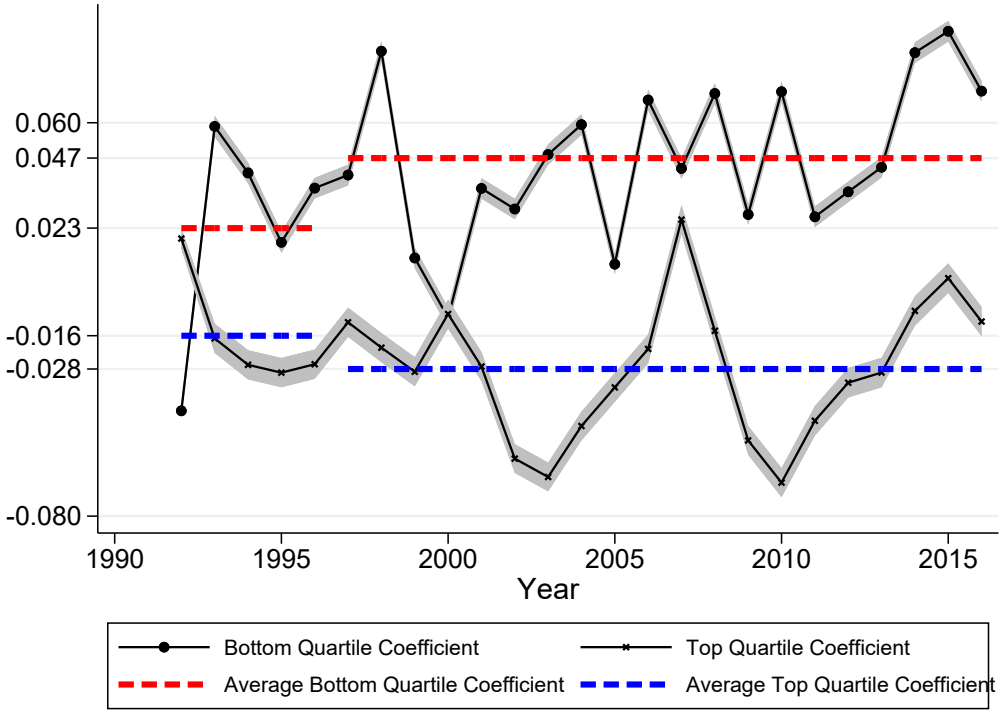
A Trimmed Sample

Figure A1: Price Growth of of Goods With High and Low E-Commerce Intensity excluding Fruits and vegetables



Data source: RPS, NSFIE, and authors' calculation. Notes: This black line shows time fixed effect $\hat{\theta}_t$ from equation (10), which tells the average rate of expenditure-weighted price increase for the goods excluding the sector "Fruits and vegetables" in two groups: products with bottom quartile e-commerce sales intensity (black line with dot) and products with top quartile e-commerce intensity (black line with symbol x). 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.

Figure A2: Price Growth of of Goods With High and Low E-Commerce Intensity



Data source: RPS, NSFIE, and authors' calculation. Notes: This black line shows time fixed effect $\hat{\theta}_t$ from equation (10), which tells the average rate of unweighted price increase for the goods in two groups: products with bottom quartile e-commerce sales intensity (black line with dot) and products with top quartile e-commerce intensity (black line with symbol x). 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.

B OLS regression results

Table A1: OLS Estimates Over Period 1991-2001

Dependent Variable	(1) Δp_{ict}	(2) Δp_{ict}	(3) Δp_{ict}	(4) Δp_{ict}
Lagged Price	-0.292***	-0.325***	-0.311***	-0.127***
	(0.026)	(0.027)	(0.024)	(0.011)
E-Commerce Intensity	-0.179		-0.033	-0.040
(t=2009) $\times$ Lagged Price	(0.269)		(0.251)	(0.114)
E-Commerce Intensity		-1.126***	-1.222***	-0.413***
$\times$ Lagged Price $\times$ Post Rakuten		(0.321)	(0.329)	(0.101)
t	{1996}	{2001}	{1996,2001}	Annual 1992-2001
Observations	25,848	27,407	51,012	152,416
R^2	0.52	0.52	0.52	0.51
E-Commerce Intensity Year	2009	2009	2009	2009

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (1) using OLS estimation method. The first column uses the five-year log differences in prices from 1991 - 1996 and the second column uses that from 1996 - 2001. The third column uses two-five year period 1991 - 1996 and 1996 - 2001. The 2SLS regression results are available from Table 5.

Table A2: OLS Estimates Over Alternative Periods

Dependent Variable	(1) Δp_{ict}	(2) Δp_{ict}	(3) Δp_{ict}	(4) Δp_{ict}	(5) Δp_{ict}
Lagged Price	-0.391***	-0.460***	-0.389***	-0.390***	-0.156***
	(0.020)	(0.026)	(0.019)	(0.016)	(0.009)
E-Commerce Intensity	0.638**	1.076***	0.477**	0.485**	0.171
(t=2009) $\times$ Lagged Price	(0.275)	(0.277)	(0.233)	(0.213)	(0.109)
E-Commerce Intensity	-1.319***	-2.458***	-1.380***	-1.297***	-0.553***
$\times$ Lagged Price $\times$ Post Rakuten	(0.250)	(0.270)	(0.321)	(0.211)	(0.086)
t	{1996,2006}	{1996,2011}	{1996,2016}	{1996,2001, 2006,2016}	Annual 1992-2016
k	5	5	5	5	
Observations	51,845	43,256	42,555	87,515	393,246
R^2	0.53	0.58	0.63	0.60	0.46
E-Commerce Intensity Year	2009	2009	2009	2009	2009

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (1) using OLS estimation method. First three columns compare the five-year period of log price differences before (1991-1996) and after the entry of e-commerce firms (2001-2006, 2006-2011, and 2011-2016). The last column uses the annual frequency of log price changes. The 2SLS regression results are available from Table A2.

C Online Appendix: Results using e-commerce intensity using Rakuten

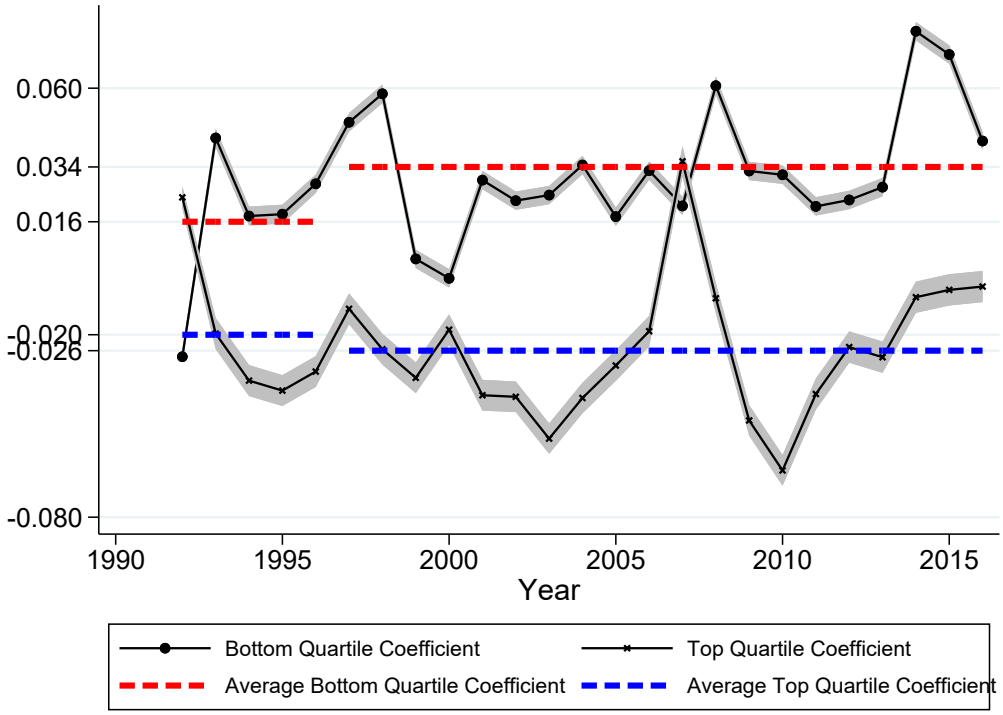
C.1 Summary Statistics

Table A3: Summary Statistics for the Sample of Goods

	Mean	Standard Deviation	Min	p10	p50	p90	Max
$x_{i(t=2010)}$	0.047	0.087	0.000	0.000	0.013	0.143	1.000
$xcat_{i(t=1999)}$	0.072	0.112	0.000	0.004	0.033	0.170	0.995
midrule Period: 1991 to 1996							
$\Delta^1 p_{ict}$	-0.003	0.109	-1.077	-0.109	0.000	0.099	1.201
Observations	141,957						
Period: 1996 to 2001							
$\Delta^1 p_{ict}$	-0.006	0.109	-1.124	-0.116	0.000	0.089	1.165
Observations	150,303						
Period: 2001 to 2006							
$\Delta^1 p_{ict}$	-0.010	0.115	-1.798	-0.129	-0.003	0.104	1.679
Observations	164,011						
Period: 2006 to 2011							
$\Delta^1 p_{ict}$	-0.001	0.125	-1.695	-0.129	0.000	0.126	1.556
Observations	164,531						
Period: 2011 to 2016							
$\Delta^1 p_{ict}$	0.014	0.100	-1.276	-0.085	0.010	0.121	1.092
Observations	168,592						

Data source: RPS, Rakuten, and authors' calculation. Notes: This table shows summary statistics of of price changes, e-commerce intensity, and catalog intensity for five five-year-periods from 1991 to 2016. Prices are in natural log. $\Delta^1 p_{ict}$ is the one-year log difference in prices; $x_{i(t=2010)} = \frac{s_i/e_i}{\max_j(s_j/e_j)}$ shows e-commerce intensity in 2010 using Rakuten and $xcat_{i(t=1999)}$ indicates catalog sales intensity in 1999.

Figure A3: Price Growth of of Goods With High and Low E-Commerce Intensity



Data source: RPS, Rakuten, and authors' calculation. Notes: This black line shows time fixed effect $\hat{\theta}_t$ from equation (10), which tells the average rate of expenditure weighted price increase for the goods in two groups: products with bottom quartile e-commerce sales intensity (black line with dot) and products with top quartile e-commerce intensity (black line with symbol x). The red dashed line shows the average rate of price increase before and after the entry of Rakuten 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.

Table A4: Relative Price Changes

	(1) $\Delta^1 p_{ict}$	(2) $\Delta^1 p_{ict}$	(3) $\Delta^1 p_{ict}$	(4) $\Delta^1 p_{ict}$
D_t	-0.0009	0.0065***	-0.0004	0.0178***
	(0.0021)	(0.0021)	(0.0045)	(0.0042)
E-Commerce Intensity	-0.0085	-0.0467***	-0.0156	-0.2601***
$\times D_t$	(0.0143)	(0.0140)	(0.0893)	(0.0739)
Sample	Goods	Goods	Goods	Goods
Fixed Effects	Product	Product	Product	Product
Estimation Period	1992-2001	1992-2016	1992-2001	1992-2016
Observations	269,742	574,509	266,202	566,374
R^2	0.03	0.03		
E-Commerce Intensity Year	2010	2010	2010	2010
First-Stage F-Stat			17.08	18.28
Estimation Method	OLS	OLS	IV	IV

Data source: RPS, Rakuten, and author's calculation. Notes: 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. For the first two columns, OLS estimates are shown with e-commerce sales intensity and the second two columns use catalog sales intensity as IV.

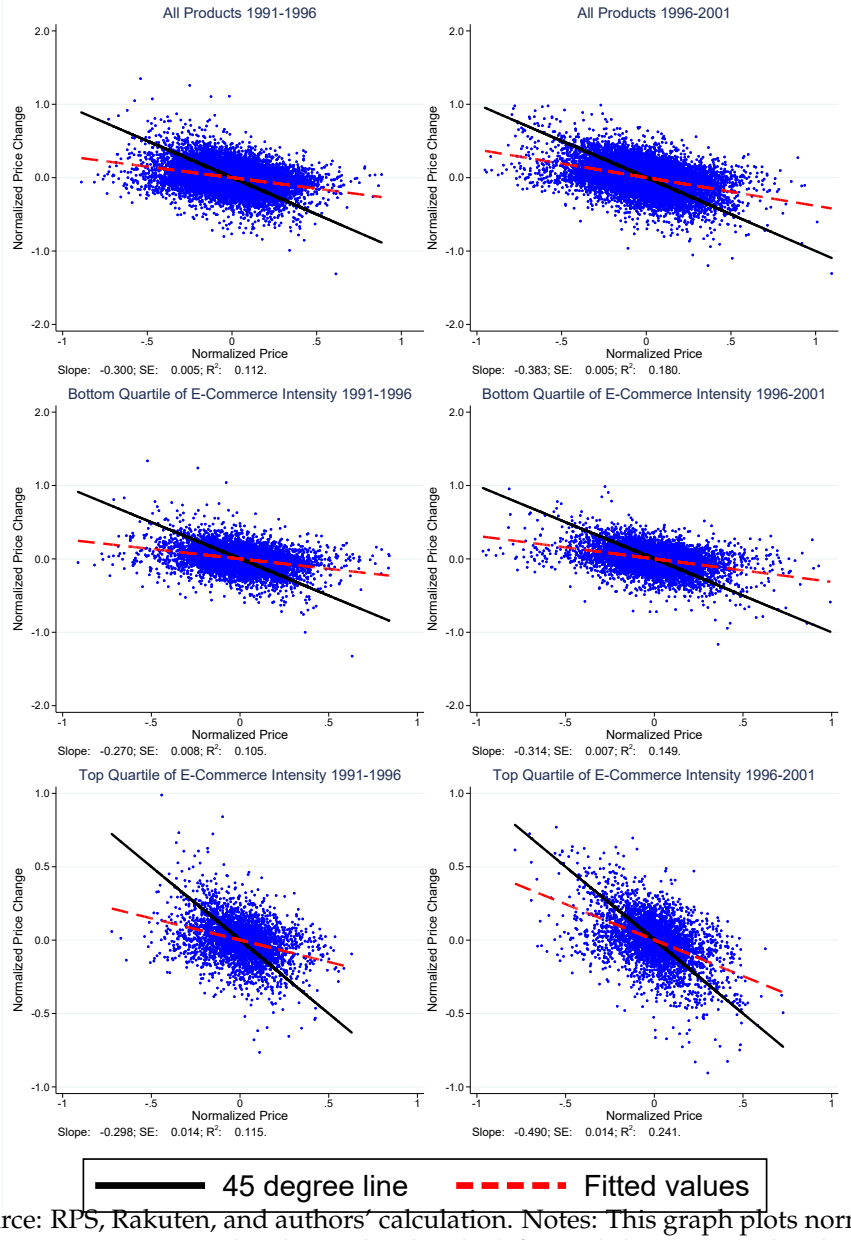
C.2 E-commerce and National Prices

Table A5: Relative Price Changes

	(1) E-Commerce Intensity $\times D_t$	(2) E-Commerce Intensity $\times D_t$	(3) E-Commerce Intensity
Catalog Intensity$\times D_t$	0.4387***	0.4455***	
	(0.1062)	(0.1042)	
D_t	0.0278***	0.0295***	
	(0.0051)	(0.0052)	
Catalog Intensity			0.4672***
			(0.1279)
Constant			0.0408***
			(0.0089)
Sample	Goods	Goods	Goods
Fixed Effects	Product	Product	None
Estimation Period	1992-2001	1992-2016	
Observations	266,202	566,374	307
R^2	0.07	0.06	0.04
E-Commerce Intensity Year	2010	2010	2010
First-Stage F-Stat	17.08	18.28	
Estimation Method	IV-First Stage	IV-First Stage	OLS

Data source: Rakuten and author's calculation. Notes: Table shows the first stage regression results.

Figure A4: Normalized Price Change vs. Normalized Price



Data source: RPS, Rakuten, and authors' calculation. Notes: This graph plots normalized price changes against normalized price levels. The left panel shows normalized price changes before the entry of e-commerce and the right panel shows them after the entry of e-commerce. The first panel plots for all goods, the second panel plots for goods with e-commerce intensity lower than the bottom quartile, and the third panel shows for goods with e-commerce intensity higher than the top quartile.

C.3 Estimation of Price Convergence

Table A6: Estimates Over Period 1991-2001

	(1)	(2)	(3)	(4)
Dependent Variable	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}
Lagged Price	-0.289***	-0.308***	-0.299***	-0.127***
	(0.035)	(0.042)	(0.035)	(0.016)
E-Commerce Intensity	-0.234		-0.169	0.420
(t=2010) $\times$ Lagged Price	(0.519)		(0.522)	(0.258)
E-Commerce Intensity		-1.568**	-1.565***	-0.629***
$\times$ Lagged Price $\times$ Post Rakuten		(0.788)	(0.378)	(0.146)
t	{1996}	{2001}	{1996,2001}	Annual 1992-2001
Observations	25,962	27,550	51,261	269,742
R^2	0.11	0.18	0.15	0.06
E-Commerce Intensity Year	2010	2010	2010	2010
First-stage F	.	.	.	.
Estimation	2SLS	2SLS	2SLS	2SLS

Source: RPS, Rakuten, and authors' calculation. Notes: Table shows regression results of equation (1) using 2SLS: e-commerce sales intensity in 2010 is instrumented using 1999 catalog sales intensity. The first column uses the five-year log differences in prices from 1991 - 1996 and the second column uses that from 1996 - 2001. The third column uses two five-year periods, 1991 - 1996 and 1996 - 2001.

Table A7: Estimates Over Alternative Periods

Dependent Variable	(1) Δp_{ict}	(2) Δp_{ict}	(3) Δp_{ict}	(4) Δp_{ict}	(5) Δp_{ict}
Lagged Price	-0.370***	-0.438***	-0.367***	-0.359***	-0.143***
	(0.034)	(0.036)	(0.032)	(0.031)	(0.015)
E-Commerce Intensity	0.696	1.435**	0.541	0.445	0.558**
(t=2010) × Lagged Price	(0.569)	(0.598)	(0.546)	(0.542)	(0.275)
E-Commerce Intensity	-2.293***	-3.759***	-2.214***	-2.095***	-0.966***
× Lagged Price × Post Rakuten	(0.529)	(0.481)	(0.431)	(0.324)	(0.134)
t	{1996,2006}	{1996,2011}	{1996,2016}	{1996,2001, 2006,2016}	Annual 1992-2016
Observations	52,055	43,417	42,711	87,875	574,509
R^2	0.17	0.23	0.20	0.20	0.08
E-Commerce Intensity Year	2010	2010	2010	2010	2010
First-stage F	.	.	.	.	.
Estimation	2SLS	2SLS	2SLS	2SLS	2SLS

Source: RPS, Rakuten, and authors' calculation. Notes: Table shows regression results of equation (1) using 1999 catalog sales intensity as an instrument for e-commerce sales intensity. First three columns compare the five-year period of log price differences before (1991-1996) and after the entry of e-commerce firms (2001-2006, 2006-2011, and 2011-2016). The last column uses the annual frequency of log price changes.

C.4 Gains Due to Price Arbitrage

Table A8: Counterfactual Welfare Gain due to Price Arbitrage

Eta	ΔW_{14}	ΔW_{17}
3	0.0021	0.0019
4	0.0028	0.0026
5	0.0035	0.0032
6	0.0042	0.0039
7	0.0049	0.0045