# The Impact of E-Commerce on Urban Prices and Welfare*

Yoon J. Jo
Columbia University†

Misaki Matsumura
Columbia University‡

David E. Weinstein
Columbia University and NBER§

June 3, 2019

Download the latest version here

**Abstract**

This paper estimates the impact of the e-commerce on Japanese prices and welfare. We find that goods sold intensively online have always had lower relative rates of price increase than goods sold mainly in physical stores, but the gap in inflation rates rose after the advent of e-commerce. This happened in part because goods sold offline began experiencing faster rates of price increase. Second, we compute the welfare gains generated by e-commerce through reducing intercity price differentials and by increasing available varieties. While we show the national gains were substantial, we also find that welfare rose much more for residents of high-income cities with highly educated populations and may have fallen for residents of other cities.

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 welfare? 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 prices and 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.

Following pioneering work by Goolsbee and Klenow (2018), we also find that goods sold intensively online have significantly lower rates of price increase than goods not sold much online. However, we differ in that we also show that this pattern was also true before e-commerce firms entered the Japanese market. Moreover, we show that while

e-commerce appears to have increased the relative difference in goods price inflation between the two sets of goods, an important reason for the increased differential is that the rate of price increase of goods not sold on e-commerce platforms rose. This may reflect a mechanism in which e-commerce eliminated low-cost physical stores, which reduced the amount of competition faced by stores that do not compete directly with online merchants.

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 welfare using the model developed in Jensen (2007). We estimate the welfare gains due to e-commerce to be 0.3 percent of consumption expenditure in 2017. 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.

The Jensen (2007) model has a number of potential shortcomings that we address in the final section of the paper. First, it is a partial equilibrium model that does not take

into account how e-commerce might have affected wages. Second, the approach is suitable for modeling the impact of e-commerce on relative prices, but it does not provide a framework that enables us to understand the impact of e-commerce on welfare that might arise because it enables consumers to access new varieties of products. 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 adopt the approach developed in Arkolakis et al. (2012) to compute the general equilibrium gains due to varieties that would have occurred if e-commerce acted like a trade technology that reduced the cost of purchasing products at a distance and thereby allowed consumers purchase more varieties.

Our results for this exercise suggest that the general equilibrium gains that arise from the main class of new trade theory models (e.g., Melitz (2003) or Eaton and Kortum (2002)) are slightly larger thanthose estimated using the Jensen (2007) approach: about a 0.7 percent welfare gain. These smaller gains are surprising given that calibrated new-trade models are hard-wired to produce welfare gains as long as e-commerce shares are nonzero. Interestingly, despite the lower aggregate gain, we also find evidence of the digital divide in this setup as well. 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 and new varieties. An important difference between these studies and ours is that we make use of household survey data to measure e-commerce sales shares and control for pre-trends and 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, Glick, and Wu (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 suitability. Section 4.1 presents our results on national prices. We present our main estimates for the impact of Rakuten on price convergence and welfare in Sections 4.2, 4.2.1, and 4.2.2. 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. Estimating the impact of e-commerce on average prices and in New Trade Theory is very straightforward following Arkolakis et al. (2012), so we will skip the theoretical discussion of how to do this and just deal with the estimation issues in Sections 4.1 and 4.3.

## 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 $\Delta^k$ 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 $\alpha_{it}$ is a item-time fixed effect; $\beta_{ct}$ is a city-time fixed effect; $\gamma$ is a parameter that captures the rate of intercity price convergence for goods not available online; $\delta_1$ is a parameter that captures the rate of price convergence for goods available online prior to the entry of e-commerce firms; $\delta_2$ captures the increase in rate of price convergence for online goods after the entry of e-commerce firms; and $\epsilon_{ict}$ 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 ($\alpha_{it}$), a common local market premium ($\beta_{ct}$), and an iid error term that is not persistent. In this case, any idiosyncratic price shock to a good in a city ($\epsilon_{ict}$) 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[1] of any price deviation from the steady-state price (measured in intervals of length $k$) as

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

As one can see from this 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 $\delta_t$ into rates of convergence. Thus, the impact of e-commerce on the rate of convergence for any good $i$ can be written as:

$$\Delta H_t \equiv \frac{\ln(0.5)}{\ln\left(1 + \hat{\gamma} + (\hat{\delta}_1 + \hat{\delta}_2) x_i\right)} - \frac{\ln(0.5)}{\ln\left(1 + \hat{\gamma} + \hat{\delta}_1 x_i\right)}. \tag{3}$$

[1] As Goldberg and Verboven (2005) note, this formula is only correct for AR1 processes.

## 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.

Jensen (2007) considered a case in which the marginal cost of supplying a market 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 Figure1, if we assume constant and equal marginal costs of production, then $G = G'$, and we can still compute the welfare gain as $(B + C + G) - \left(E + F + G'\right) = (B + C) - (E + F)$.[2]

We can use our estimates of the impact of e-commerce on price convergence to calibrate the Jensen model. In order to compute the partial equilibrium welfare gain due to e-commerce, we will consider the difference in implied gains in two equilibria: the first is one in which consumers do not have access to e-commerce, and the second is one in

[2]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.

Figure 1: Welfare Gains from Arbitrage in the Jensen Model

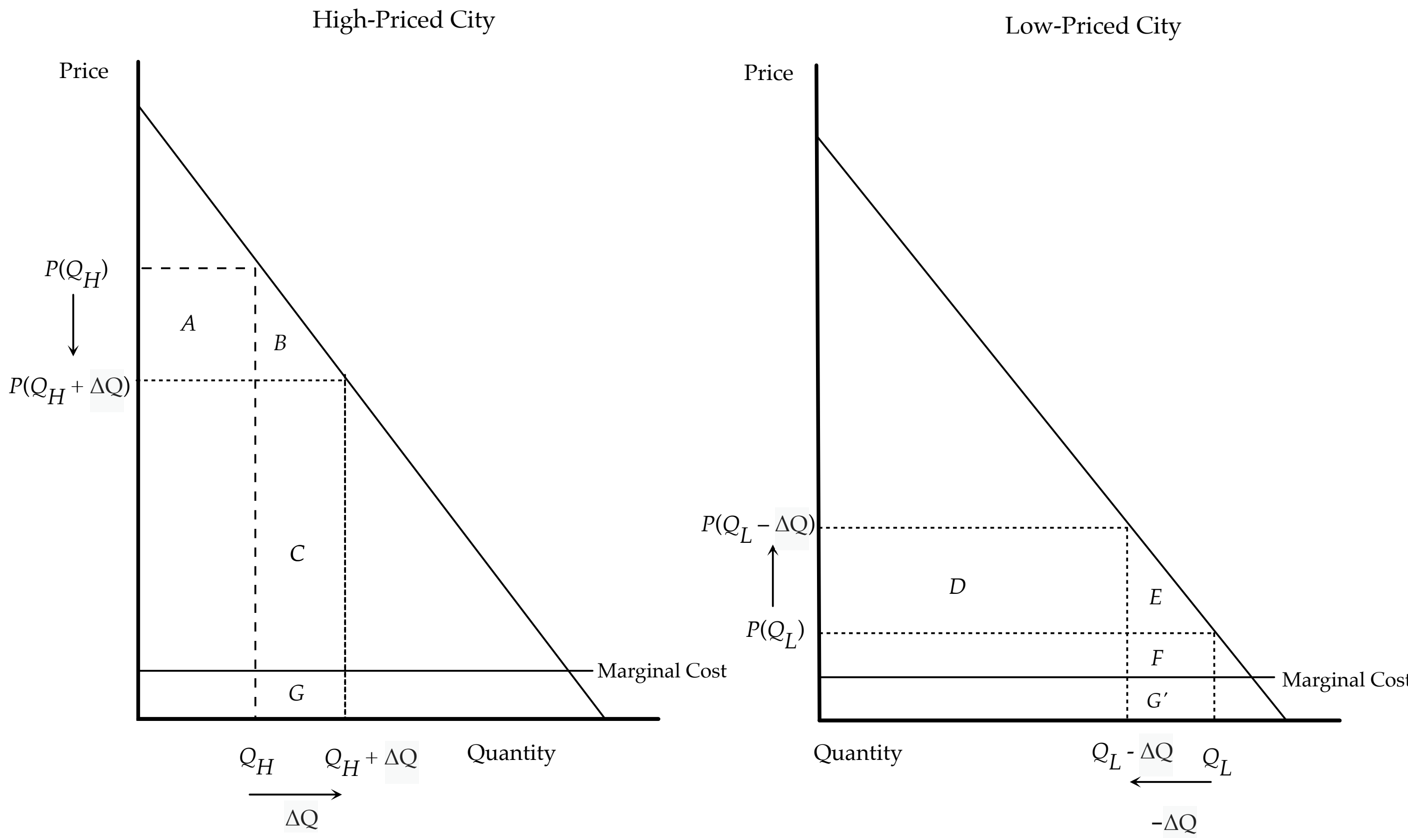


which they do have access. In each case, we will assume that the economy has deviated from a steady state equilibrium but experiences different rates of price convergence differ. We then can compute the welfare gains associated with the different convergence rates. Let $D$ denote the counterfactual, 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

$$\widehat{\Delta p_{icT}}\,(D) = \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D x_i\right)(p_{ic,T-1} - p^*_i)\,, \tag{4}$$

The price level in period $T$ therefore can be written as

$$\Delta P_{icT}\left(D\right) \equiv P_{icT}\left(D\right) - P_{ic,T-1}\left(D\right), \tag{5}$$

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

$$P_{icT}(D) = P_{ic,T-1}\left[1 + \widehat{\Delta p}_{icT}\left(D\right)\right]. \tag{6}$$

The demand for good $i$ in city $c$ in time $t$ in counterfactual $D$ is given by

$$Q_{ict}\left(D\right) = \frac{\left(P_{ict}\left(D\right)/\varphi_{ic}\right)^{-\eta}}{\left[P_{ct}\left(D\right)\right]^{1-\eta}}E_{c}, \tag{7}$$

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

$$P_{ct}\left(D\right) \equiv \left[\sum_{i}\left[P_{ict}\left(D\right)/\varphi_{ic}\right]^{1-\eta}\right]^{\frac{1}{1-\eta}}. \tag{8}$$

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

Following Jensen, we consider a set of price changes that are consistent with equation (4) 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., $\Delta Q_{iT}\left(D\right) = \sum_{c}\Delta Q_{icT}\left(D\right)$. The counterfactual change in urban consumption in any city $c$ in time $T$ is given by $\widehat{\Delta Q}_{icT}\left(D\right) \equiv Q_{icT}\left(D\right) - Q_{ic,T-1}$. Using the log change as an approximation for the percentage change, we can rewrite this as

$$Q_{icT}(D) = Q_{ic,T-1}\left[1 + \widehat{\Delta q}_{icT}\left(D\right)\right], \tag{9}$$

where equation (7) implies that

$$\widehat{\Delta q}_{icT}\left(D\right) = \left(\eta - 1\right)\Delta p_{cT}\left(D\right) - \eta\widehat{\Delta p}_{icT}\left(D\right) + \Delta\ln E_{cT}, \tag{10}$$

If we make the make the partial equilibrium assumption that when we use counter-

factual price changes ($\widehat{\Delta p}_{icT}(D)$) aggregate urban expenditures and prices are unaffected by price arbitrage, this equation reduces to

$$\widehat{\Delta q}_{icT}(D) = -\eta \widehat{\Delta p}_{icT}(D), \tag{11}$$

Substituting equation (11) into equation (9) and summing produces

$$\sum_c Q_{ic,T-1} \widehat{\Delta p}_{icT}(D) = 0 \tag{12}$$

If we then substitute equation (4) into equation (12) we obtain

$$\sum_c Q_{ic,T-1} \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D x_i\right)(p_{ic,T-1} - p_i^*) = 0 \tag{13}$$

or

$$p_i^* = \frac{\sum_c Q_{ic,T-1} \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D x_i\right) p_{ic,T-1}}{\sum_c Q_{ic,T-1} \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D x_i\right)} \tag{14}$$

In order to calibrate this model, we set $P_{ic,T-1} = P_{ic16}$, and $Q_{ic,T-1} = E_{ic16}/P_{ic16}$, where $E_{ic16}$ is expenditures in city $c$ on good $i$ in 2016. This lets us obtain the steady state price $p_i^*$ for each good.

The partial equilibrium welfare gain arising from prices moving from their values in 2016 ($P_{ic16}$) towards their steady state values ($p_i^*$) can be written as

$$\Delta W_{icT} = \frac{1}{2}\left(2P_{ic16}(D) + \Delta P_{icT}(D)\right)\Delta Q_{icT}(D) - m_i \Delta Q_{icT}(D), \tag{15}$$

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

The welfare change associated with price not equaling their steady state levels equals

$$\Delta W_{iT}(D) = \frac{1}{2}\sum_c (2P_{ic16} + \Delta P_{icT}(D))\,\Delta Q_{icT}(D) - \underbrace{m_i \sum_c \Delta Q_{icT}(D)}_{=0}. \tag{16}$$

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

$$\Delta \hat{W}_{iT}(D) = \frac{\Delta W_{iT}(D)}{E_{i16}(D)}. \tag{17}$$

The welfare loss in the economy is given by

$$\Delta W_T(D) = \sum_i \Delta \hat{W}_{iT}(D)\,\frac{E_{i16}(D)}{E_{16}(D)}, \tag{18}$$

and the gain due to e-commerce can then be written as

$$\Delta \hat{W}_{iT} = \Delta \hat{W}_{iT}(D=1) - \Delta \hat{W}_{iT}(D=0). \tag{19}$$

$$\Delta W_T = \Delta W_T(D=1) - \Delta W_T(D=0). \tag{20}$$

In other words, the partial-equilibrium gain due to e-commerce arises from the relative increase in price convergence arising from the ability to arbitrage prices across space.

# 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”. 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, the largest e-retail company in Japan, who provided us with 2010 internet sales data (aggregated across buyers and merchants) for each of approximately 40,000 product 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 tradable 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:

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

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

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. As most of services are not available online, we will focus on e-commerce's impact on goods prices for all of our main results. 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 $\ell$ 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 $\Omega^\ell$ is the set of items in some more aggregated category $\ell$. In the second column, 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 to be the ratio of the two previous columns divided it by the maximum value of $S_\ell/E_\ell$ (i.e., $x_i \equiv S_i/E_i/\left[\max_j \left\{S_j/E_j\right\}\right]$). Thus, our measure of e-commerce intensity takes on a value of zero if there are no 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 $i$.

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 $x_i = \frac{s_i}{e_i}/max_j(\frac{s_i}{e_i})$.

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.

As one can also see from the table, 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 (Olsen 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 have characteristics that are similar to those goods historically available in catalogs—i.e., goods that are non-perishable, low weight-to-value, standardized, and storable.

Although e-commerce was small in Japan in 1999, the situation changed radically over the next two years. 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.[4] The following year sales on the Rakuten platform exceeded ¥52 billion (about $430 million). Thus, within five years, Japanese consumers in any city went from only being able to buy locally or from catalogs to being able to purchase goods from thousands of merchants located across Japan. Rakuten's growth was part of a broader e-commerce boom in Japan. 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 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"). While the data is not sufficiently detailed to always pin down the exact barcode, the data leaves limited scope for unobserved quality differences to affect intercity price differentials. For example, Imai and Watanabe (2015) find that it is sufficiently detailed to rule out approximately 85 percent of all bar codes in a CPI product category. Moreover, 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 US CPI data typically is based on different baskets of goods in different cities, the JSB's "purposive" sampling generate samples in

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

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.[5] 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[6], 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.

[5]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.

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

Table 2: Summary Statistics for the Sample of Goods

| | Mean | Standard Deviation | Min | p10 | p50 | p90 | Max |
|---|---|---|---|---|---|---|---|
| Period: 1991 to 1996 | | | | | | | |
| $\Delta^1 p_{ict}$ | -0.004 | 0.099 | -0.858 | -0.101 | 0.000 | 0.090 | 0.953 |
| $x_{i(t=2009)}$ | 0.057 | 0.073 | 0.000 | 0.004 | 0.022 | 0.157 | 0.456 |
| $x_{i(t=2009)} \times p_{ic(t-1)}$ | 0.456 | 0.679 | 0.000 | 0.021 | 0.142 | 1.340 | 4.107 |
| $xcat_{i(t=1999)}$ | 0.078 | 0.119 | 0.000 | 0.004 | 0.035 | 0.207 | 0.995 |
| $xcat_{i(t=1999)} \times p_{ic(t-1)}$ | 0.592 | 1.022 | 0.000 | 0.019 | 0.219 | 1.662 | 9.725 |
| 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 |
| $x_{i(t=2009)}$ | 0.053 | 0.070 | 0.000 | 0.004 | 0.022 | 0.151 | 0.456 |
| $x_{i(t=2009)} \times p_{ic(t-1)}$ | 0.417 | 0.646 | 0.000 | 0.021 | 0.125 | 1.215 | 4.075 |
| $xcat_{i(t=1999)}$ | 0.073 | 0.112 | 0.000 | 0.004 | 0.034 | 0.170 | 0.995 |
| $xcat_{i(t=1999)} \times p_{ic(t-1)}$ | 0.541 | 0.934 | 0.000 | 0.022 | 0.212 | 1.376 | 9.826 |
| 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 |
| $x_{i(t=2009)}$ | 0.052 | 0.074 | 0.000 | 0.004 | 0.022 | 0.146 | 1.000 |
| $x_{i(t=2009)} \times p_{ic(t-1)}$ | 0.400 | 0.690 | 0.000 | 0.021 | 0.115 | 1.169 | 11.690 |
| $xcat_{i(t=1999)}$ | 0.071 | 0.107 | 0.000 | 0.004 | 0.033 | 0.170 | 0.995 |
| $xcat_{i(t=1999)} \times p_{ic(t-1)}$ | 0.508 | 0.873 | 0.000 | 0.020 | 0.202 | 1.243 | 9.703 |
| 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 |
| $x_{i(t=2009)}$ | 0.053 | 0.081 | 0.000 | 0.004 | 0.022 | 0.146 | 1.000 |
| $x_{i(t=2009)} \times p_{ic(t-1)}$ | 0.411 | 0.751 | 0.000 | 0.020 | 0.114 | 1.197 | 10.684 |
| $xcat_{i(t=1999)}$ | 0.070 | 0.104 | 0.000 | 0.004 | 0.032 | 0.162 | 0.995 |
| $xcat_{i(t=1999)} \times p_{ic(t-1)}$ | 0.497 | 0.845 | 0.000 | 0.019 | 0.197 | 1.235 | 9.411 |
| 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 |
| $x_{i(t=2009)}$ | 0.052 | 0.080 | 0.000 | 0.004 | 0.020 | 0.143 | 1.000 |
| $x_{i(t=2009)} \times p_{ic(t-1)}$ | 0.398 | 0.725 | 0.000 | 0.020 | 0.109 | 1.175 | 10.300 |
| $xcat_{i(t=1999)}$ | 0.068 | 0.104 | 0.000 | 0.004 | 0.029 | 0.162 | 0.995 |
| $xcat_{i(t=1999)} \times p_{ic(t-1)}$ | 0.485 | 0.838 | 0.000 | 0.016 | 0.165 | 1.215 | 9.081 |
| 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; $x_{i(t=2009)} = \frac{s_i/e_i}{\max_j(s_j/e_j)}$ shows e-commerce intensity in 2009 using NSFIE and $xcat_{i(t=1999)}$ indicates catalog sales intensity in 1999.

One obvious concern with these data is that they are not as good as barcode data. However, we can test how problematic they are by computing some simple sample statistics. Hottman, Redding, and Weinstein (2016) show that the correlation between price and quality in bar-code data is 0.9, so we should expect sampling problems to produce high levels of price dispersion in our sample. Thus, if there is substantial quality variation within the goods used in the Japanese CPI sample, we should expect to see a lot of intercity price dispersion for the same items. In order to check for this, we compute the price of each good in each city less the average price of that good across all cities and then taking the standard deviation of this difference. When we do this, we find that the standard deviation of intercity price differences for the same good in Japan is about 15 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 average annual rate of inflation across the sample period. 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. The second line 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 0.146 over the full sample period, which is more than six 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 45 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.

# 4 Estimation

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 in Section 4.2.1. Finally, in Section 4.2.2, we present our estimates of the welfare gain from e-retail.

## 4.1 E-commerce and National Prices

Goolsbee and Klenow (2018) find that goods traded online have inflation rates that were about one percentage point lower than goods not available online. Here, we extend this work to show that these 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 ($\Delta p_{ict}$) on good ($\alpha_i$) and city ($\beta_c$) 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 ($x_i D_t$):

$$\Delta p_{ict} = \alpha_i + \beta_c + \phi D_t + \theta x_i D_t + \epsilon_{ict}, \tag{21}$$

where $\alpha_i$ 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 $D_t$ ($\phi$) tells us whether there was any differential trend in price inflation for goods available online after the entrance of e-commerce firms and $\theta$, 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 ($\alpha_i$) that might confound specifications that compare growth rates of goods available online with those not sold online. In order to understand 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:

$$\Delta p_{ict} = \theta_t + \epsilon_{ict}, \tag{22}$$

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

We plot these estimates and the 95-percent confidence bands in Figure 2. As the figure makes clear, there are unmistakable pre-trends in the data. Before the entry of the Rakuten

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

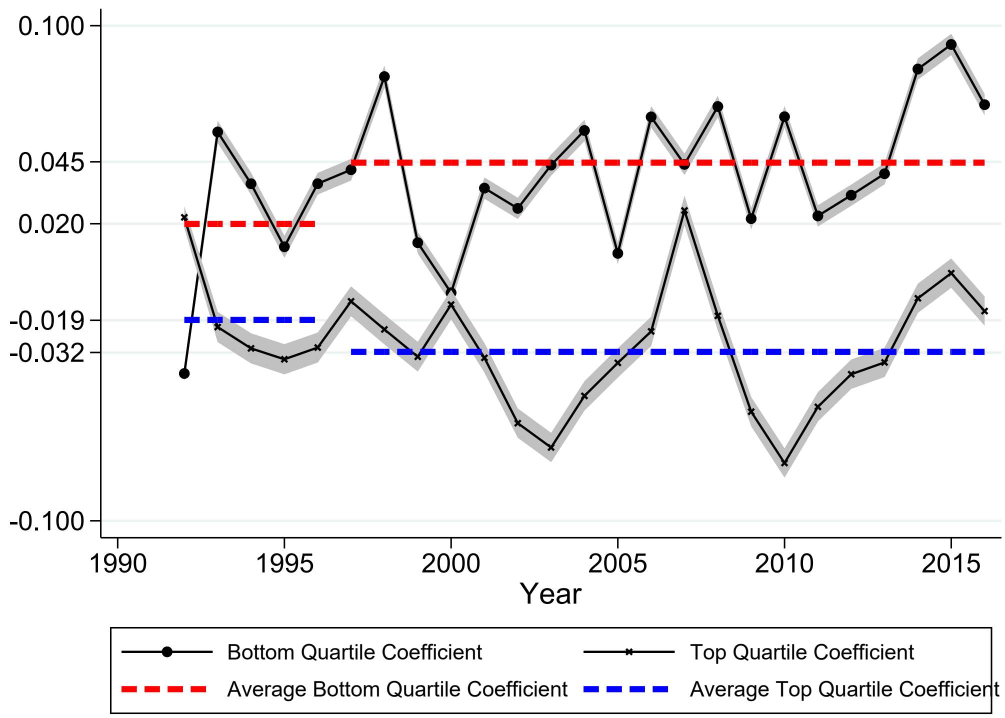


Data source: RPS, NSFIE, and authors' calculation. Notes: This black line shows time fixed effect $\hat{\theta}_t$ from equation (22), which tells the average rate of price increase for the goods in two groups: products with bottom quartile e-commerce sales intensity (black line with dot) and products 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.

in 1997, the average rate of price increase for the types of goods that would ultimately be sold on e-commerce platforms was -1.9 percent per year, while the average annual rate of price increase for goods that not sold much on these platforms was 2.0 percent per year. Thus, there was a 3.9 percentage point gap between the relative inflation rates of goods would be sold intensively online relative to those would not be sold intensively online even in the early 1990s. These differences in inflation rates 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 inflation rates 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. 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 -3.2 percent per year: a 1.3 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 4.5 percent per year: an increase of 2.5 percent per year.[7]

[7]The cause of this price increase is not clear, but the higher prices charged by merchants for goods not sold intensively online may be related to the impact that e-commerce firms have had on large national retailers. DellaVigna and Gentzkow (2017) argue that chain stores typically offer uniform pricing across cities and this works to depress regional price dispersion. This implies that a negative shock to chain stores that sell standardized products might be associated with a rise in prices in stores selling products not available online. For example, suppose mass-merchandisers traditionally sell groceries and electronics. If e-commerce firms eliminate mass-merchandise stores, local grocery stores may find themselves with fewer physical competitors and more able to raise prices.

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. 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.

Turning to our differences-in-differences specification, we present the results from estimating equation (21) in Table 3. The first column present the results from estimating equation (21) 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 columns 2 and 3, 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. The coefficient implies that a good at the $90^{th}$ percentile of internet intensity experienced a 1.3 percent per year drop in its rate of annual price growth relative to goods not available online.

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 results of our instrumental variables estimation. 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.

Table 4: Relative Price Changes

| | (1)<br>E-Commerce Intensity $\times D_t$ | (2)<br>E-Commerce Intensity $\times D_t$ | (3)<br>E-Commerce Intensity |
|---|---|---|---|
| Catalog Intensity | 0.7363*** | 0.7304*** | |
| $\times D_t$ | (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. Notes: Table shows the first stage regression results.

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 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), pre-dates 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 different goods exhibit different average price changes in different years. We therefore normalized the data by regressing $\Delta p_{ict}$ 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 ($p_{ic,t-5} - \hat{\alpha}'_{it} - \hat{\beta}'_{ct}$), where $\hat{\alpha}_{it}$ ($\hat{\alpha}'_{it}$) and $\hat{\beta}_{ct}$ ($\hat{\beta}'_{ct}$) are the estimated fixed effects from the regression of $\Delta p_{ict}$ ($p_{ict-5}$) 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 be-

Figure 3

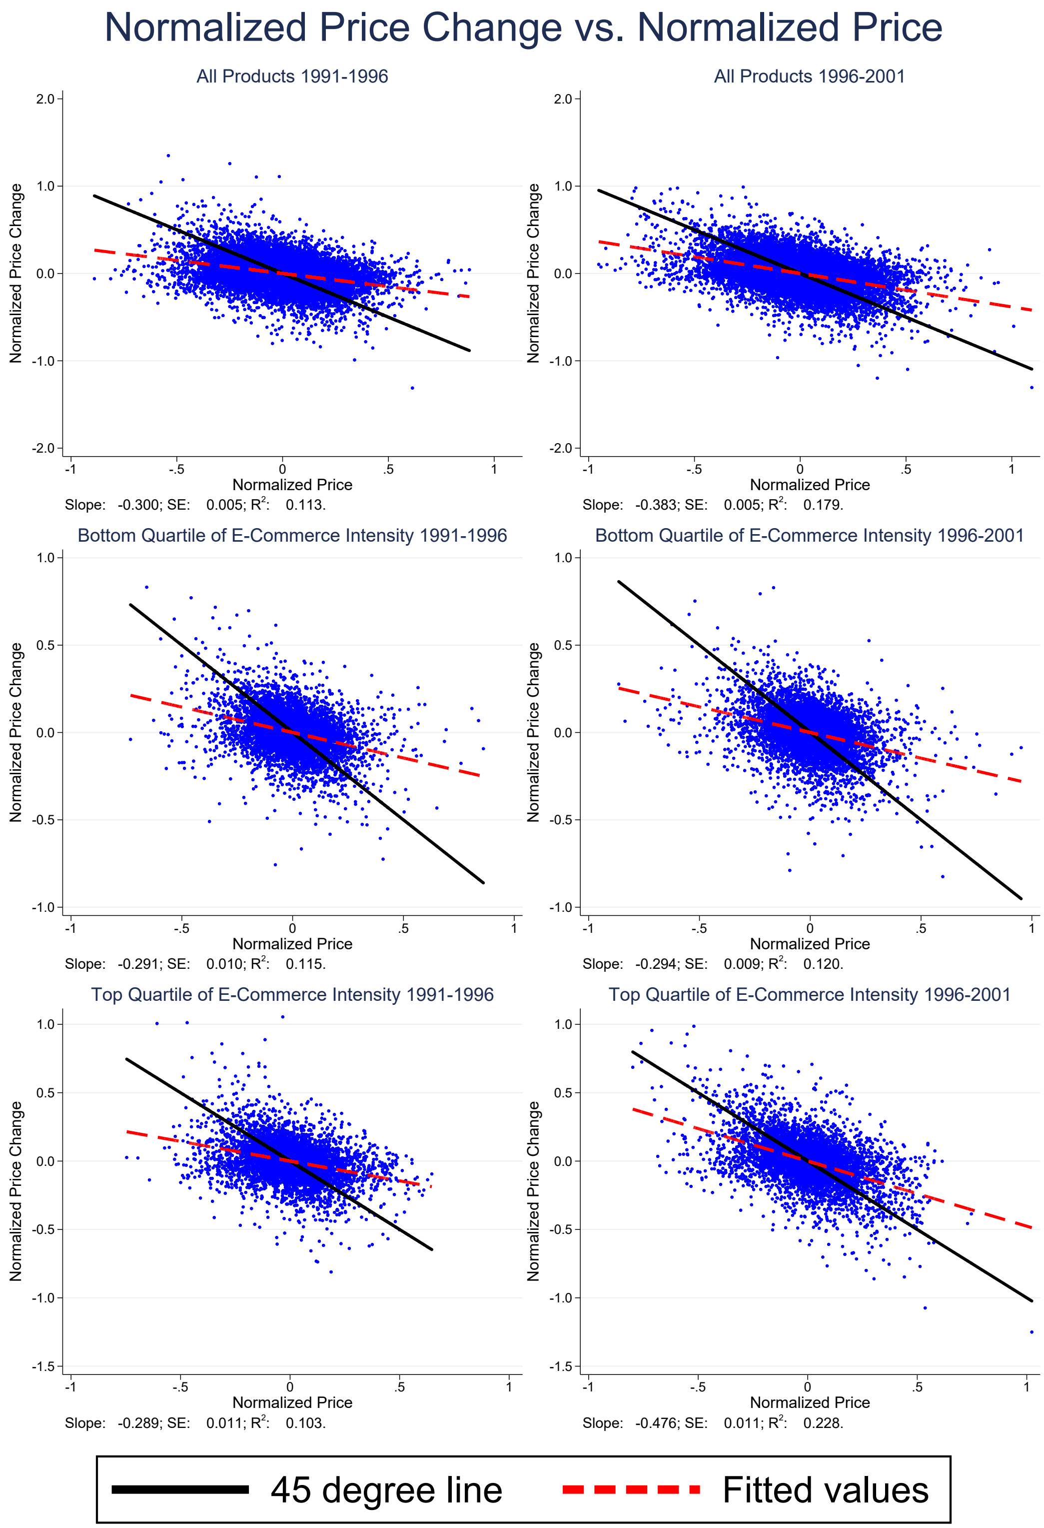


Data source: RPS, NSFIE, and authors' calculation. Notes: This graph plots normalized price changes against normalized price levels. The left panel shows normalized price changes before the entry of e-commerce and the right panel shows them after the entry of e-commerce. The first panel plots for all 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.

fore 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 inflation 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 controlled for city fixed effects.[8]

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 ($x_i < 0.076$) and the set of goods in the highest quartile of the distribution ($x_i > 0.13$). 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 much 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.

[8]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)).

### 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. 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 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 (where we dropped the $x_i p_{ic,t-5}$ term from the specification because we do not have any pre-e-commerce observations) indicates significantly faster convergence rates for goods available online after the entry of e-commerce firms also confirming the result we saw in we saw in Figure 3.

Table 5: Estimates Over Period 1991-2001

| Dependent Variable | (1) $\Delta p^5_{ict}$ | (2) $\Delta p^5_{ict}$ | (3) $\Delta p^5_{ict}$ | (4) $\Delta p^1_{ict}$ |
|---|---|---|---|---|
| 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) $\times$ Lagged Price | (0.413) | | (0.410) | (0.151) |
| E-Commerce Intensity | | -1.145* | -1.292*** | -0.509*** |
| $\times$ Lagged Price $\times$ 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 OLS regression results are available first three columns OLS regression results are available in Appendix A from Table A1.

In column 3, we estimate our baseline differences-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 ($\hat{\delta}_2 = -1.292$). Interestingly, the estimate of $\delta_1$, the differential rate of price convergence for e-commerce intensive goods remains close to zero, indicating that e-commerce appears to have no impact on intercity price convergence rates of goods not sold intensively online.

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 last column uses the annual frequency of log price changes. The OLS regression results are available in Appendix A Table A2.

One additional concern is that our measure of e-commerce intensity imposes a parametric structure on how shifts in e-commerce sales shares translate into convergence rates. Since it may not be the case that a doubling in e-commerce sales intensity produces a doubling in convergence rates, we divided our internet intensity into bins depending on which percentiles of the distribution they fell into. The first column of Table A3 reports the results using two bins and the second column reports the results using three bins. In each case, the dummy variable $I_{by}$ is an indicator variable that is one when an observation falls into bin $y$ in a specification with $b$ bins. In all specifications, a greater value of $y$ corresponds to a bin of observations that have larger e-commerce intensities than those in a bin with a smaller value of $y$. As before, we instrument $I_{by}$ by using the catalog sales intensity bin for the same product. As one can see from Table A3, products that fall into higher percentiles of the e-commerce sales intensity tend to converge at faster rates,

and this result holds whether we use two bins or three bins. Thus, our results are not qualitatively dependent on the particular parameterization of the relationship between convergence speed and e-commerce sales intensity.

Table 7: Estimates with Indicators

| | (1) $\Delta p_{ict}$ | (2) $\Delta p_{ict}$ |
|---|---|---|
| Lagged Price | -0.119*** | -0.120*** |
| | (0.011) | (0.013) |
| I22 × Lagged Price | 0.014 | |
| | (0.021) | |
| I21 × Lagged Price × Post E-Commerce | -0.030** | |
| | (0.013) | |
| I22 × Lagged Price × Post E-Commerce | -0.073*** | |
| | (0.011) | |
| I32 × Lagged Price | | -0.032 |
| | | (0.062) |
| I33 × Lagged Price | | 0.052** |
| | | (0.024) |
| I31 × Lagged Price × Post E-Commerce | | -0.024* |
| | | (0.014) |
| I32 × Lagged Price × Post E-Commerce | | -0.063*** |
| | | (0.022) |
| I33 × Lagged Price × Post E-Commerce | | -0.086*** |
| | | (0.015) |
| $t$ | Annual | Annual |
| Observations | 574,509 | 574,509 |
| $R^2$ | 0.08 | 0.08 |
| E-Commerce Intensity Year | 2009 | 2009 |
| First-stage F | 184.58 | 7.98 |
| Estimation | 2SLS | 2SLS |

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (XX) using 1999 catalog sales intensity as an instrument for e-commerce sales intensity. Both columns uses the annual frequency of log price changes in the period 1992-2016. The OLS regression results are available in Appendix A Table A3.

There are a number of potential problems with the evidence that we have just presented. A first concern is that these results may 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 confirm what we saw 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 ($\hat{\delta}_2$) approximately doubles if we compare periods ten or more years after the entry of Rakuten with the period before. 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. Similarly, goods at the $90^{th}$ percentile of e-commerce sales intensity have a half life of price dispersion of 2.6 years. Thus, our estimates imply that the advent of e-commerce seems to 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 our result using measures of e-commerce intensity based on Rakuten sales data instead. We report the results from this exercise in Appendix Table A5, 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 (**??**). One of the interesting features of these equations is that the welfare gain is proportional to the choice of demand elasticity. Since this elasticity has been estimated in other papers, we calibrate a demand elasticity of 6 and simply note that the welfare gain

using any other elasticity ($\eta$) equals the welfare gain we report multiplied by $\eta/6$. If we use the convergence estimates based on column 1 of Table A3, we obtain a welfare gain of 0.0029 or 0.27 percent due to increased price arbitrage. Similarly, convergence estimates based on column 2 produce a gain of 0.32 percent.

### 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 $j$ 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 to pay a fixed cost and incurs a marginal cost of sales equal to $(a/\varphi)\, w_j$, where $\varphi$ 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 $\tau_{ij}$. 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 ($\tau_{ij}$). [9]

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 $\epsilon$ 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 -5, we will obtain a welfare gain of 2.1 percent (=$-\ln(0.9)/5$). 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 $\lambda_t$ 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 expenditures purchased from all retailers ($\chi$) 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 -5, his gives us an estimate of the welfare gain from e-commerce in Japan

[9] 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.

in 2014 of 0.5 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 0.4 percent to 1.0 percent in 2014 and from 0.5 percent to 1.2 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 Change

| Epsilon | $\Delta W_{14}$ | $\Delta W_{17}$ | $\Delta W^c_{14}$ | $\Delta W^c_{17}$ |
|---|---|---|---|---|
| -3 | 0.009 | 0.012 | 0.010 | 0.013 |
| -5 | 0.005 | 0.007 | 0.006 | 0.008 |
| -7 | 0.004 | 0.005 | 0.004 | 0.006 |

Data source: NSFIE, JSB, MIETI and authors' calculation. Notes: The first two column shows welfare gain due to new varieties from e-commerce in 2014 and 2017 along with plausible trade elasticities. The last two column shows welfare gain due to the increased variety from e-commerce and catalog sales in 2014 and 2017.

The second two columns 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 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, etcetera. 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 0.3 to 0.5 percent.

As in Section 4.2.2, it is also interesting to see how these gains have affected individual prefectures. In order to do this, we computed the welfare gains for each prefecture. Since 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 ($\lambda_{p14}$) has a strong negative correlation with the share of college educated people in a prefecture. This produces a strong positive correlation (0.68) between the welfare gain in a prefecture and the share of college-educated people in the

Figure 4: $\Delta W_{pt}^{c} : \epsilon = -5$ vs. Share of College Education

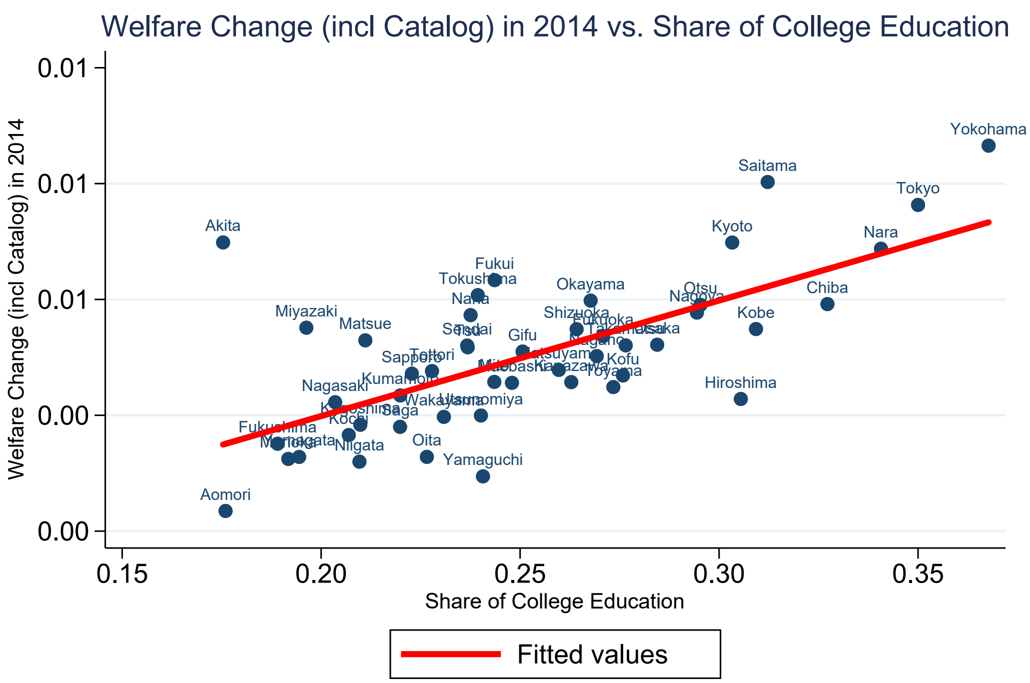


Slope: 0.020; SE: 0.003; $R^2$: 0.464.

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

prefecture. We can see this clearly in Figure 4, where we plot the welfare gain (including catalog sales in our definition of non-local sales). This result suggests that our earlier result about 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 even if we shift our methodology of computing welfare gains.[10]

Table 9: Prefecture Welfare Change

| | (1)<br>$\Delta W_{14}$ | (2)<br>$\Delta W_{14}$ | (3)<br>$\Delta W_{14}$ | (4)<br>$\Delta W_{14}$ | (5)<br>$\Delta W_{14}$ | (6)<br>$\Delta W_{14}$ |
|---|---|---|---|---|---|---|
| Share of College Educated | 0.0195***<br>(0.0032) | 0.0167***<br>(0.0041) | 0.0214***<br>(0.0039) | 0.0180***<br>(0.0041) | 0.0123**<br>(0.0047) | 0.0144***<br>(0.0052) |
| Population | | 0.0000<br>(0.0000) | | | | 0.0000<br>(0.0000) |
| Income per Capita | | | -0.0000<br>(0.0000) | | | -0.0000<br>(0.0000) |
| Average Age | | | | -0.0001<br>(0.0001) | | 0.0000<br>(0.0001) |
| Share of Secondary Educated | | | | | -0.0087**<br>(0.0043) | -0.0084<br>(0.0052) |
| Constant | -0.0001<br>(0.0008) | 0.0004<br>(0.0009) | 0.0005<br>(0.0011) | 0.0034<br>(0.0060) | 0.0053*<br>(0.0028) | 0.0052<br>(0.0061) |
| 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,

[10]Whether we include or exclude catalog sales does not matter substantively for our prefectural results. If we define $\lambda_{p14}$ without counting catalog sales as non-local expenditures, we obtain the same correlation up to two significant digits, and the plot looks very similar.

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 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.[11] 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 internet on Japanese prices and 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 online rose, which underscores the difficulty in interpreting

[11] 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.

the relatively low rates of price increase for goods sold online as providing information about how e-commerce affects aggregate inflation.

At the local level, we find strong evidence that the rate at which intercity price differences disappear rose 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 0.3 percentin 2017.

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 larger overall gains from e-commerce—a welfare rise of 0.7 percent. This may reflect the fact that these general equilibrium models allow for e-commerce to affect welfare not only through relative price movements, but also through increasing the set of varieties. These models also let us measure regional gains. We find that highly educated regions benefited more than less educated regions. 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.

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.

Jonathan Eaton and Samuel Kortum. Technology, Geography and Trade. *Econometrica*, 5 (70):1741–1779, 2002.

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.

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

Pinelopi Goldberg and Frank Verboven. Market integration and convergence to the Law of One Price: evidence from the European Car Market. *Journal of International Economics*, 65(1):49–73, 2005.

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.

Satoshi Imai and Tsutomu Watanabe. Replicating Japan's CPI Using Scanner Data. *CARF Working Paper Series, CARF-F-364*, 2015.

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.

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 OLS regression results

Table A1: OLS Estimates Over Period 1991-2001

| Dependent Variable | (1)<br>$\Delta p_{ict}$ | (2)<br>$\Delta p_{ict}$ | (3)<br>$\Delta p_{ict}$ | (4)<br>$\Delta p_{ict}$ |
|---|---|---|---|---|
| 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<br>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)<br>$\Delta p_{ict}$ | (2)<br>$\Delta p_{ict}$ | (3)<br>$\Delta p_{ict}$ | (4)<br>$\Delta p_{ict}$ | (5)<br>$\Delta p_{ict}$ |
|---|---|---|---|---|---|
| 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,<br>2006,2016} | Annual<br>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.

Table A3: OLS Estimates with Indicators

| | (1) $\Delta p_{ict}$ | (2) $\Delta p_{ict}$ |
|---|---|---|
| Lagged Price | -0.139*** | -0.147*** |
| | (0.012) | (0.014) |
| I22 $\times$ Lagged Price | 0.050*** | |
| | (0.016) | |
| I21 $\times$ Lagged Price $\times$ Post E-Commerce | -0.031*** | |
| | (0.011) | |
| I22 $\times$ Lagged Price $\times$ Post E-Commerce | -0.073*** | |
| | (0.008) | |
| I32 $\times$ Lagged Price | | 0.061*** |
| | | (0.020) |
| I33 $\times$ Lagged Price | | 0.056*** |
| | | (0.018) |
| I31 $\times$ Lagged Price $\times$ Post E-Commerce | | -0.021* |
| | | (0.012) |
| I32 $\times$ Lagged Price $\times$ Post E-Commerce | | -0.080*** |
| | | (0.009) |
| I33 $\times$ Lagged Price $\times$ Post E-Commerce | | -0.071*** |
| | | (0.009) |
| $t$ | Annual | Annual |
| Observations | 574,509 | 574,509 |
| $R^2$ | 0.42 | 0.42 |
| E-Commerce Intensity Year | 2009 | 2009 |

Source: RPS, NSFIE, and authors' calculation. Notes: Table shows regression results of equation (1) using OLS estimation method. Both columns uses the annual frequency of log price changes in the period 1992-2016. The 2SLS regression results are available from Table 7.

# B Results using e-commerce intensity using Rakuten

## B.1 E-commerce and National Prices

Table A4: Relative Price Changes

| Dependent Variable | (1)<br>$\Delta p_{ict}$ | (2)<br>$\Delta p_{ict}$ | (3)<br>$\Delta p_{ict}$ | (4)<br>$\Delta p_{ict}$ | (5)<br>$\Delta p_{ict}$ | (6)<br>$\Delta p_{ict}$ | (7)<br>$\Delta p_{ict}$ |
|---|---|---|---|---|---|---|---|
| $D_t$ | -0.0014** | 0.0002 | 0.0002 | 0.0047*** | 0.0064*** | 0.0064*** | 0.0030*** |
| | (0.0005) | (0.0006) | (0.0006) | (0.0005) | (0.0005) | (0.0005) | (0.0004) |
| Internet Intensity | 0.0047 | -0.0156*** | -0.0155*** | -0.0211*** | -0.0421*** | -0.0421*** | -0.0279*** |
| $\times D_t$ | (0.0035) | (0.0050) | (0.0050) | (0.0018) | (0.0043) | (0.0043) | (0.0040) |
| Constant | -0.0040*** | -0.0121*** | -0.0097** | -0.0041*** | -0.0166*** | -0.0164*** | -0.0138*** |
| | (0.0004) | (0.0040) | (0.0045) | (0.0004) | (0.0027) | (0.0030) | (0.0028) |
| Sample | Goods | Goods | Goods | Goods | Goods | Goods | Goods and Service |
| Fixed Effects | | Product | Product and City | | Product | Product and City | Product and City |
| $t$ | 1992-2001 | 1992-2001 | 1992-2001 | 1992-2016 | 1992-2016 | 1992-2016 | 1992-2016 |
| Observations | 152,958 | 152,958 | 152,958 | 394,663 | 394,663 | 394,663 | 459,279 |
| $R^2$ | 0.000 | 0.040 | 0.040 | 0.001 | 0.038 | 0.038 | 0.038 |

## B.2 Gains Due to Price Arbitrage

Table A5: Estimates Over Alternative Periods

| Dependent Variable | (1)<br>$\Delta p_{ict}$ | (2)<br>$\Delta p_{ict}$ | (3)<br>$\Delta p_{ict}$ | (4)<br>$\Delta p_{ict}$ | (5)<br>$\Delta p_{ict}$ |
|---|---|---|---|---|---|
| Lagged Price | -0.399*** | -0.466*** | -0.391*** | -0.401*** | -0.169*** |
| | (0.005) | (0.005) | (0.005) | (0.003) | (0.001) |
| Internet Intensity | 0.562*** | 0.816*** | 0.359*** | 0.416*** | 0.189*** |
| $\times$ Lagged Price | (0.073) | (0.078) | (0.072) | (0.073) | (0.025) |
| Internet Intensity | -0.950*** | -1.863*** | -1.142*** | -1.011*** | -0.348*** |
| $\times$ Lagged Price $\times$ Post Rakuten | (0.091) | (0.091) | (0.085) | (0.078) | (0.026) |
| $t$ | {1996,2006} | {1996,2011} | {1996,2016} | {1996,2001,<br>2006,2016} | Annual<br>1992-2016 |
| $k$ | 5 | 5 | 5 | 5 | 1 |
| Observations | 52,017 | 43,388 | 42,683 | 87,818 | 394,663 |
| $R^2$ | 0.55 | 0.60 | 0.64 | 0.61 | 0.46 |

Table A6: Robustness Check Using All Goods and Services

| Dependent Variable | (1) $\Delta p_{ict}$ | (2) $\Delta p_{ict}$ | (3) $\Delta p_{ict}$ |
|---|---|---|---|
| Lagged Price | -0.297*** | -0.353*** | -0.141*** |
| | (0.004) | (0.003) | (0.001) |
| Internet Intensity | -0.052 | 0.145** | 0.025 |
| $\times$ Lagged Price | (0.065) | (0.070) | (0.024) |
| Internet Intensity | -1.061*** | -1.013*** | -0.329*** |
| $\times$ Lagged Price $\times$ Post Rakuten | (0.083) | (0.076) | (0.025) |
| $t$ | {1996,2001} | {1996,2001 | Annual |
| | | ,2006,2016} | 1992-2016 |
| Observations | 59,466 | 102,102 | 459,279 |
| $R^2$ | 0.52 | 0.60 | 0.45 |