The Impact of E-Commerce on Prices and Welfare*

Yoon J. Jo Misaki Matsumura David E. Weinstein

Columbia University† Columbia University‡ Columbia University and NBER§

November 26, 2018

Preliminary and Incomplete: Please Do Not Cite or Quote Without Permission

Abstract

This paper estimates the impact of the internet on Japanese prices and welfare. First, we consider the possibility that e-commerce may have lowered prices by driving down the average prices of goods available online. Second, we compute the welfare gains due to the ability of e-commerce to enable consumers to purchase goods from other regions. Third, we compute the gains that arise through e-commerce's ability to arbitrage intercity price differences. We find that all three channels produced welfare gains in Japan, but our estimates suggest that the first and second channels are by far the most important, with welfare gains through these channels being eleven to sixteen times larger than through the price arbitrage channel. . Overall, we find that increased inter-city arbitrage raised Japanese welfare by 0.06 percent, the gains due to new varieties available through online shopping raised welfare 0.7 percent, and the gains due to overall price reductions for goods available online raised welfare by 1 percent.

JEL CLASSIFICATION: F11, F14, L86, R32

KEYWORDS: comparative advantage, trade, prices, quality, variety

*We are especially grateful to the Rakuten Institute of Technology for providing us with data and assistance for this project. 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? Economists have suggested a number of approaches to answering this question. In early work, Brynjolfsson et al. [2003] suggested that the internet raises welfare by giving consumers better access to new varieties. More recent work, has considered impacts through impacts on average prices (Goolsbee and Klenow [2018]) or through the ability of information technology to generate price convergence (c.f. Brown and Goolsbee [2002], Jensen [2007], and Steinwender [2018]). One of the limitations of prior work has been that it has focused on small samples of goods (e.g., books), had no data on price trends prior to the creation of the internet (which makes assigning the impact of e-commerce difficult to measure), or focused on different information technologies (cell phones or telegraphs).

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 using a number of popular modeling techniques. First, we consider the possibility that e-commerce may have lowered prices by driving down the average prices of goods available online. Second, following Brynjolfsson et al. [2003] and Arkolakis, Costinot, and Rodriguez-Clare (2012), we compute the welfare gains due to the ability of e-commerce to enable consumers to purchase goods from other regions. Third, following Jensen [2007], we compute the gains that arise through e-commerce's ability to arbitrage intercity price differences. We find that all three channels produced welfare gains in Japan, but our estimates suggest that the first and second channels are by far the most important, with welfare gains through these channels being eleven to sixteen times larger than through the price arbitrage channel. Overall, we find that increased inter-city arbitrage raised Japanese welfare by 0.06 percent, the gains due to new varieties available through online shopping raised welfare 0.7 percent, and the gains due to overall price reductions for goods available online raised welfare by 1 percent.

However, it has proven much harder to document how the most prominent e-commerce

technology for consumers—the online market places that make up much of e-retail—have affected prices and welfare in modern developed countries. In many countries, e-retailers have disrupted traditional physical retailers by giving consumers access to a vast number of stores in many locations. As a result local merchants must now compete not only within their own geographic area but with merchants located anywhere in a country, reducing their ability to charge high markups.

We make use use of two unique datasets. The first covering all sales transacted by Japan's largest e-retailer, Rakuten, and the second, a quinquennial government survey of households that reports the purchase locations of Japanese consumers. We merge these data with the price data underlying the Japanese consumer price index (CPI) , covering 30,000 price quotes per year in physical stores across Japan over the period 1991 to 2016. We also use these estimates to compute the welfare gains arising from the ability of e-retail to arbitrage away regional price differences.

In order to identify the impact of e-commerce, we exploit the fact that the growth of e-retail happened rapidly in Japan, so it is relatively easy to divide recent history into clear pre and post e-retail periods. Rakuten was one of the earliest entrants to the e-retail business, starting in February of 1997. Unlike Amazon, Rakuten's business model was not to sell goods directly to consumers, but rather to serve as a platform through which consumers could easily find stores willing to sell products online. This difference enabled them to expand without developing warehouses and extensive logistics networks. Thus, while only 40 percent of Amazon's transaction value comes from third-party sellers, virtually all of Rakuten's transactions come from this source. Since its entry into Japan, Rakuten experienced explosive growth. By April of 2000, when it 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.[1] Already at this point, the Rakuten website was getting 95 million hits per month—almost one hit for every man,

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

woman, and child in Japan. 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. By 2010, Rakuten had a 30 percent market share in the Japanese e-commerce market, eclipsing global giants like Amazon and Yahoo.[2] Rakuten transaction values continued to skyrocket reaching 12.7 trillion yen by 2017: equal to about 2.3 percent of Japanese GDP.[3].

Since we can observe both which goods tend to be sold intensively online and can observe how pricing behavior for goods changed before and after the entry of Rakuten, we can adopt differences-in-differences strategy to identify how pricing behavior changed for goods suitable for e-retail relative to goods not suitable for e-retail. We first characterize the set of goods suitable for e-retail by calculating the difference between the share of consumer expenditure for a good sold on the Rakuten platform with that in the Japanese Family Income and Expenditure Survey (FIES), which is use to compute the expenditure weights for the Japanese consumer price index. While we find that virtually all goods are sold online, there is enormous variation in e-retail intensity. Not surprisingly, highly perishable, non-standardized goods are not sold intensively online. Neither are highly time-sensitive goods and goods requiring consumer identification (e.g., medicines) or goods with high weight-to-value ratios. To a first approximation, goods sold online closely resemble the set of goods that historically dominated the catalog business: books, clothing, footwear, hardware, gardening supplies, and (later) electronics. [4] Moreover, we find that regional variation in e-commerce sales intensity is entirely driven by the share of college-educated people, with no role for urban-rural or young-old divides once education once one controls for education.

We use this feature of the data to first document evidence that the entry of Rakuten is associated with price drops in physical retailers selling goods that are available online. This result is complementary to that of Goolsbee and Klenow [2018] who found that goods

[2]Rakuten Annual Report (2010).
[3]Rakuten Annual Report (2017).
[4]see https://dingley.com/five-pivotal-moments-in-catalog-history/

available online in the U.S. have a lower rate of inflation than goods not available online. Our result differs in two important dimensions. First, we show that this price impact is apparent in the pricing decisions of physical retailers. Second, we show that the association between goods that can be sold online and their relative rates of price changes did not exist prior to the entry of large e-commerce firms. Thus, we provide evidence that advent of e-commerce changed relative pricing behavior by lowering the relative prices of goods suitable for e-retail by reducing relative retail markups, improving logistics, or generating some other efficiency gain. If we ignore general equilibrium effects by assuming that e-commerce did not affect prices in sectors untouched by e-retail, this appears to have lowered the Japanese price level by about one percentage point between 1997 and 2016.

We next consider the impact of e-commerce on welfare through enhancing consumer access to new varieties. Computing these gains for the class of models introduced by Krugman (1980), Melitz (2003), and Eaton and Kortum (2002) can be done quite easily with our data by simply computing the share of online purchases. This share rose rapidly to levels quite similar to those observed in the U.S. Based on these sales shares, we obtain an estimate of the welfare gain from e-commerce of 0.7 percent in 2016.

Finally, we examine how the difference in convergence rates across cities between goods available online and those not available online changed after the entry of Rakuten into the market. 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 convergence or half-lives of price differentials of nine years. In contrast, we find that half-lives for price differentials across Japanese cities are only 4.8 years for both goods that are suitable and unsuitable for online sales prior to Rakuten's entrance into the market. The difference probably arises from the fact that Japanese CPI data is based on the sampling of identical or extremely similar goods across cities, which makes it superior for this type of analysis. Rakuten's entry is associated with a 25 percent drop in this half-life for goods traded intensively online, which we assume is due to the difficulty of maintaining intercity

price differences in a world with e-retail.

With these estimates in hand, we next compute the impact of e-retail on Japanese welfare through improved price arbitrage using the model developed in Jensen [2007]. An interesting feature of this model is that not only is it the case that consumers in low-priced cities lose while consumers in high-priced cities gain, the price movements imply that on net consumers lose. This result doesn't mean that arbitrage isn't beneficial. The net welfare gain occurs because producers in the low-priced cities can now sell in the high-priced cities. While this does entail a loss for producers in the high-priced cities, the increase in surplus for producers in the low-priced cities exceeds the losses for producers in the high-priced cities and the net loss of consumer surplus, resulting in higher real income in the economy. Thus, the internet acts very much like a reduction in trade costs in that there are lots of winners and losers. Relative producer surplus falls and consumer surplus rises in the sectors with comparative disadvantage (locations with relatively high-priced goods), but relative producer surplus rises and consumer surplus falls in the sectors with comparative advantage (locations with relatively low-priced goods). However, just like a reduction in trade costs, the net gains are positive.

In order to estimate these effects, we compute counterfactual price paths for the Japanese economy. The first is based on our estimates of how intercity price differentials should dissipate in the presence of e-retail, and the second is based on how intercity price differences would dissipate in the absence of e-retail. The faster rate of price convergence in the presence of e-retail implies that welfare gains will be higher after Rakuten entered the market. When we quantify this for Japan, we estimate that the entry of Rakuten raised the welfare of Japanese consumers by about 0.06 percent per year.

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 number of papers have 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] examined the impact of the introduction of mobile phones on fish or agricultural markets in India, Niger, and the Philippines, and Steinwender [2018] examined 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.

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 government surveys of e-commerce sales to compute e-commerce sales shares (an important input into welfare calculations) rather than using private Alibaba or credit-card data to assess e-commerce expenditure shares, which are a key input into welfare calculations. These private measures of e-commerce intensity deviate substantially from official measures. For example, the U.S. Census Bureau reports that in 2014, e-commerce firms accounted for about 6 percent of total retail sales, and the Japanese Ministry of Economy Trade and Industry reports that in the same year, Japanese e-commerce firms had a 4.4 percent market share.[5] By contrast, the e-commerce sales intensity in credit card or Ali Research data is around 20 percent for the U.S. and 10.6 percent for China. Part of the discrepancy comes from the fact that credit cards are more likely to be used for e-commerce transactions and

[5]https://census.gov/retail/mrts/www/data/pdf/ec_current.pdf

balance-sheet sales numbers for e-commerce firms often include revenues from other sources—finance, telecommunications, etc.—that do not necessarily involve e-commerce. For example, of Rakuten's net sales, 42 percent comes from e-commerce, 7 percent from travel, and of the remainder most comes from credit cards, banking, securities, and telecommunications.

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).

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. Goolsbee and Klenow [2018] show that price trends of goods sold online exhibited lower inflation rates in the US between 2014 and 2017 than in the CPI. While we also show that goods available online exhibit lower inflation rates than goods not available online, there are two important differences between our study and theirs. First, we can show that this pattern was not present prior to the introduction of e-commerce. This fact makes it clear that the differential patterns in inflation rates is not a characteristic of the goods themselves but rather it was an effect that appeared for goods sold online only after the introduction of the internet. Second, we show that goods sold intensively online exhibit lower rates of price increase in physical stores. This is consistent with the idea that online price trends are mirrored in the prices of physical retailers.

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. Bergin et al. [2017] employ a similar triple difference strategy to show that rates of price convergence across European countries increased after joining the euro area. Our contribution lies in adapting some of the techniques developed in these earlier papers for understanding how e-commerce affects regional price convergence.

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-retail suitability. We present our main estimates for the impact of Rakuten on price convergence and welfare in Section 4, and Section 5 concludes.

2 Theory

Estimating the impact of e-commerce on average prices and in New Trade Theory is very straightforward following Arkolakis, Costinot, and Rodriguez-Clare (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 4.2. In Section 2.1, we model the impact that e-commerce has had on interregional price differential, and we show how the decline in these differentials raises welfare in Section 2.2.

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

We begin by defining some notation. Let $p_{ict} \equiv \ln P_{ict}$ be the log price of item i in city c in time t. Define the Δ^k operator as $\Delta^k p_{ict} \equiv p_{ict} - p_{ic,t-k}$; thus, if we set $k = 1$, we can examine annual changes, but we can also examine longer differences by setting k equal to a whole number larger than one. Let $x_i \in [0, 1]$ be the "e-retail suitability" of the good, where

zero indicates it is not suitable for e-retail sales 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-retail sales are positive in period t and zero otherwise. We assume that the change in the price of any item in a city c can be written as a standard purchasing price parity specification in which we introduce a modification that allows the rate of price converge for goods available online to change, i.e.,

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

where α_{it} is a item-time fixed effect; β_{ct} is a city-time fixed effect; γ is a parameter that captures the rate of intercity price convergence for goods not available online; δ_1 is a parameter that captures the rate of price convergence for goods available online prior to the entry of Rakuten; δ_2 captures the increase in rate of price convergence for online goods after the entry of Rakuten; and ϵ_{ict} is an iid error term. We think of this error as price shocks arising from period t local supply-and-demand conditions for an item in a city that are not shared by all items in the city and are uncorrelated with past prices.

In this specification, a critical parameter is the rate of convergence given by $(\gamma + \delta_1 x_i + \delta_2 D_t x_i)$, which we expect to be between -1 and 0. A value of -1 means that equation (1) collapses to $p_{ict} = \alpha_{it} + \beta_{ct} + \epsilon_{ict}$, and therefore the price of any item can be decomposed into its national price (α_{it}), a common local market premium (β_{ct}), and an iid error term that is not persistent. In this case, any idiosyncratic price shock to a good in a city (ϵ_{ict}) has no impact on prices in the next period. Hence, price convergence occurs in one period, and prices always equal their conditional mean of $(\alpha_{it} + \beta_{ct})$ plus a random iid shock. At the other extreme, we have the case of where $(\gamma + \delta_1 x_i + \delta_2 D_t x_i) = 0$, which implies that the price of that good i in city c follows a random walk with a drift term given by $(\alpha_{it} + \beta_{ct})$. In intermediate cases where $(\gamma + \delta_1 x_i + \delta_2 D_t x_i) \in (-1, 0)$, price differences across cities can persist for more than k years.

In our setup, we can write the approximate half-life[6] of any price deviation from the

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

steady-state price (measured in intervals of length k) as[7]

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

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 δ_t into rates of convergence. Thus, the impact of e-retail on the rate of convergence for any good i can be written as:

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

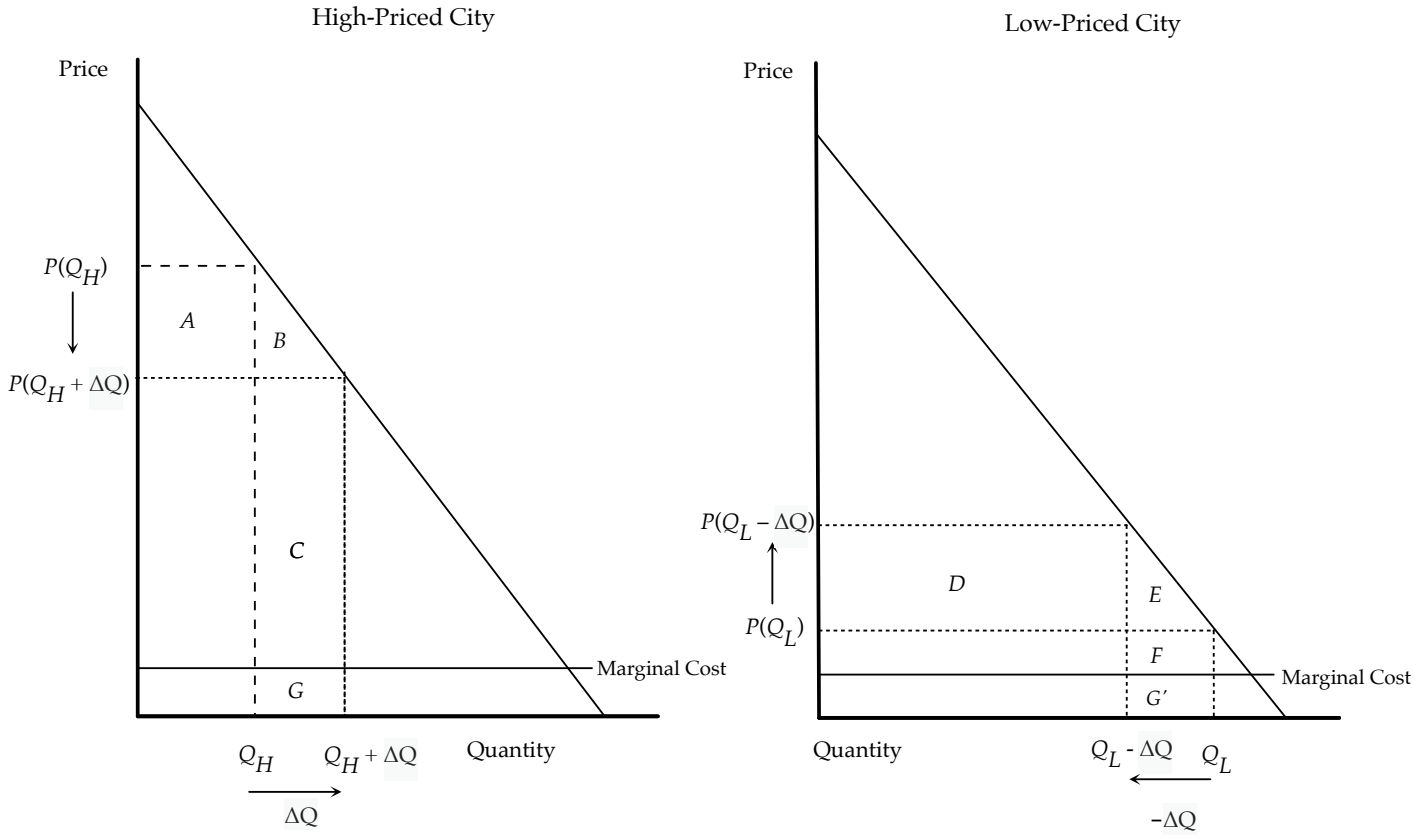
2.2 Welfare

We can use these estimates to inform us about the welfare gains from e-retailing 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). This framework can easily be applied to the e-retailing context since firms like Rakuten provide a platform that enables consumers in any city to purchase goods from a large number of retailers spread across Japan. If e-retailing enables retailers in the low-priced region to make ΔQ units of sales to the high-priced region, 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

[7]The steady state price is given by the price at which setting $\Delta p_{ict} = 0$. This price equals $p_{ic} = -(\alpha_{it}+\beta_{ct})/(\gamma+\delta_1 x_i+\delta_2 D_t x_i)$.

the heights of their parallel sides.

Figure 1: Welfare Gains from Arbitrage in the Jensen Model



Jensen 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 $(B + C + G) - \left(E + F + G'\right) = (B + C) - (E + F)$.[8]

[8]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.

In order to compute the welfare gain, we need to compute the price change associated with the arbitrage opportunity associated with e-retail. We begin by writing the change in welfare due to the price change of good i in city c over k periods as

$$\Delta W_{ict} = \frac{1}{2}\left(2P_{ic,t-k} + \Delta P_{ict}\right)\Delta Q_{ic,t} - m\Delta Q_{ic,t}, \tag{4}$$

where m is the marginal cost of producing the good. Without loss of generality we can decompose prices and quantities into two components: a national component that captures national movements in the price of good i (ΔP_{it}), a city-specific component that captures relative movements in prices in that city (ΔP^R_{ict}):

$$\Delta P_{ict} = \Delta P_{it} + \Delta P^R_{ict}. \tag{5}$$

Let the total quantity demanded be denoted by $Q_{it} \equiv \sum_c Q_{ict}$. We can now decompose quantity movements (ΔQ_{ict}) into a national component ($\Delta Q^N_{ict} \equiv \left(Q_{ic,t-k}/Q_{i,t-k}\right)\Delta Q_{it}$), which tells us how consumption in the city would have moved if it followed the national trend, and a city-specific ($\Delta Q^R_{ict} \equiv \Delta Q_{ic,t} - \Delta Q^N_{it}$) component. This lets us rewrite equation (4) as

$$\Delta W_{ict} = \frac{1}{2}\left(2P_{ic,t-1} + \Delta P_{it} + \Delta P^R_{ict}\right)\left(\Delta Q^N_{ict} + \Delta Q^R_{ict}\right) - m\Delta Q_{ic,t-1} \tag{6}$$

Rearranging terms produces

$$\begin{aligned}
\Delta W_{ict} &= \frac{\left(2P_{ic,t-1} + \Delta P_{it} + \Delta P^R_{ict}\right)\Delta Q^N_{ict}}{2} + \frac{\left(2P_{ic,t-k} + \Delta P_{it} + \Delta P^R_{ict}\right)\Delta Q^R_{ict}}{2} - m\Delta Q_{ict} \\
&= \frac{\left(2P_{ic,t-1} + \Delta P_{it}\right)\Delta Q^N_{ict}}{2} + \frac{\left(2P_{ic,t-k} + \Delta P^R_{ict}\right)\Delta Q^R_{ict}}{2} \\
&\quad + \frac{\Delta P^R_{ict}\Delta Q^N_{ict}}{2} + \frac{\Delta P_{it}\Delta Q^R_{ict}}{2} - m\Delta Q_{ict} \\
&= \left(\Delta W^N_{it} - m\Delta Q^N_{ict}\right) + \left(\Delta W^R_{ict} - m\Delta Q^R_{ict}\right) + \frac{\Delta P^R_{ict}\Delta Q^N_{ict}}{2} + \frac{\Delta P_{it}\Delta Q^R_{ict}}{2}
\end{aligned} \tag{7}$$

In other words, the change in welfare in a city can be decomposed into a term that depends on how the national change in prices and quantities affected welfare, a second term that depends

on how price movements in that city relative to the national average affected welfare, and two second-order terms that capture the fact that a relative price decline matters more for welfare if it occurs for a good that is on average growing in consumption and one that captures the fact that a national drop in prices raises welfare more if local demand is also rising. Our focus will be on the aggregate gains arising from the second term ($\sum_c \Delta W^R_{ict}$=$\sum_c \Delta W^R_{ict} - m \sum_c \Delta Q^R_{ict}$), which captures the first-order impact of relative price movements on welfare.

Defining the terms this way lets us write $\Delta Q^R_{ict}/Q_{ict} = \Delta Q_{ict}/Q_{ict} - \Delta Q_{it}/Q_{it}$, which is more convenient to write as a log approximation given by $\Delta q^R_{ict} \equiv \Delta q_{ict} - \Delta q_{it}$. If we assume that there is no regional variation in demand elasticities ($\eta_c = \eta \ \forall c$), we then have

$$\Delta q^R_{ict} \equiv \Delta q_{ict} - \Delta q_{it} = -\eta \left(\Delta p_{ict} - \Delta p_{it} \right). \tag{8}$$

If we multiply this expression by $Q_{ic,t-5}$ and sum across all cities we obtain an expression for the aggregate change in in quantity due to relative price movements across the cities:

$$\sum_c \Delta Q^R_{ict} \equiv \sum_c Q_{ic,t-1} \Delta q^R_{ict} = -\eta \left(\sum_c Q_{ic,t-1} \Delta p_{ict} - \sum_c Q_{ic,t-1} \Delta p_{it} \right) = 0, \tag{9}$$

where the last equality follows from our assumption (which is the same as that of Jensen [2007]) that relative price changes arising from new arbitrage opportunities do not affect aggregate demand for the good. This expression lets us solve for the expression for national prices:

$$\Delta p_{it} = \frac{\sum_c Q_{ic,t-1} \Delta p_{ict}}{\sum_c Q_{ic,t-1}} \tag{10}$$

In other words, the log-change in national price index is just a quantity-weighted average of the log price change in each city. We now can conduct our counterfactual welfare analysis. Based on equation (1), we can write the price change due to arbitrage as:

$$\widehat{\Delta p_{ict}} = \hat{\alpha}_{it} + \hat{\beta}_{ct} + \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D_t x_i \right) p_{ic,t-1}. \tag{11}$$

Therefore, the counterfactual aggregate price change for any good is

$$\widehat{\Delta p}_{it} = \hat{\alpha}_{it} + \sum_c \frac{Q_{ic,t-1}\hat{\beta}_{ct}}{\sum_c Q_{ic,t-1}} + \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D_t x_i\right) \sum_c \frac{Q_{ic,t-1}p_{ic,t-1}}{\sum_c Q_{ic,t-1}}, \tag{12}$$

The evolution of prices in each city relative to the national price increase can be expressed as the difference between equations (11) and (12):

$$\widehat{\Delta p}^R_{ict}(D_t) = \left[\hat{\beta}_{ct} - \sum_c \frac{Q_{ic,t-1}\hat{\beta}_{ct}}{\sum_c Q_{ic,t-1}}\right] + \left(\hat{\gamma} + \hat{\delta}_1 x_i + \hat{\delta}_2 D_t x_i\right) \left[p_{ic,t-k} - \sum_c \frac{Q_{ic,t-1}p_{ic,t-1}}{\sum_c Q_{ic,t-1}}\right]. \tag{13}$$

We now can write the impact of a price change on welfare in a city as

$$\Delta W^R_{ict}(D_t) = \frac{-\eta}{2}\left[2P_{ic,t-k} + P_{ic,t-k}\widehat{\Delta p}^R_{ict}(D_t)\right] Q_{ic,t-k}\widehat{\Delta p}^R_{ict}(D_t) \tag{14}$$

The national gain in welfare can then be written as:

$$\widehat{\Delta W}^R_t(D_t) \equiv \sum_c \sum_i \widehat{\Delta W}^R_{ict}(D_t). \tag{15}$$

The real gain due to enhanced arbitrage from e-retail can then be written as:

$$\widehat{\Delta W}^E_t \equiv \left[\widehat{\Delta W}^R_t(D_t = 1) - \widehat{\Delta W}^R_t(D_t = 0)\right]/P_t, \tag{16}$$

where P_t is the price level in year t. Note that this gain only applies to any particular period. If we want to compute the aggregate gain over more than one period, we can calculate the cumulative value of these gains between t and T:

$$\widehat{\Delta W}^E_T = \sum_{s=t}^{T} \widehat{\Delta W}^R_s. \tag{17}$$

3 Data

A major challenge in prior research is the fact that e-retailers are typically loathe to release sales data for products. Thus, prior work has not been able to observe the value of e-retail sales. Fortunately, Rakuten, the largest e-retail company in Japan, provided us with 2010 internet sales data (aggregated across buyers and merchants) for each of their narrowest product categories covering the universe of transactions on their platform. In that year, Rakuten had a 30 percent market share of all Japanese e-commerce, so we believe that its transactions are likely to be representative of the overall internet sales intensity of goods.[9]

When uploading the product information on to the Rakuten platform, sellers designate whether their product falls into one of approximately 40,000 product categories or "genres". We then matched these genres to the expenditure categories in the 2010 Japanese Family Income and Expenditure Survey (FIES) and the categories that appear in 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, and of 364 expenditure items that we use to characterize all Japanese expenditures.

We construct e-retail intensity of an expenditure category by comparing average household's total expenditure on that category with average household's online expenditure on it. We measure total expenditure e_i on category i by using national average expenditures per household in 2010 taken from FIES. We denote expenditures on the Rakuten platform in category i by s_i. We compute s_i by summing gross merchant sales across all genres contained in the set of genres in category i (Ω^i), i.e., $s_i \equiv \sum_{j \in \Omega^i} s_j$. We then define e-retail intensity x_i of category i by taking the ratio of the online to total expenditures, normalized by the maximum value of this ratio:

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

In order to see how e-retail intensity varies across products, we aggregated the FIES codes

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

into some broader categories in Table 1 so that we could display the data in a compact form. The rows are ordered by a category's share of Japanese expenditures. The first column of Table 1 reports the percentage of expenditures in category ℓ as reported in the FIES ($E_\ell \equiv \sum_{i \in \Omega^\ell} e_i / \sum_j e_j \times 100$), where Ω^ℓ is the set of items in some more aggregated category ℓ. In the second column, we report the percentage of Rakuten sales 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 sales on the Rakuten platform). The final column reports the "internet intensity", which we define to be the ratio of the two previous columns divided it by the maximum value of S_ℓ / E_ℓ (i.e., $x_i \equiv S_i / E_i / [\max_j \{S_j / E_j\}]$). Thus, our measure of internet intensity takes on a value of zero if there are no transactions involving an expenditure category and a value of 1 if the expenditures on the Rakuten platform relative to those in the economy is the highest among all categories of goods. Expressing internet intensity this way makes our internet intensity (x_i) invariant to the size of sector i.

Table 1: E-Retail Sales Intensity of Non-Housing Consumer Expenditures

Category	Share of Consumer Expenditures	Share of Rakuten Transactions	$x_i = \frac{S_i/E_i}{\max_j(S_j/E_j)}$
Goods			
Household consumables	4.62	20.69	0.44
Fruits and vegetables	4.16	1.03	0.02
Clothing	3.72	13.02	0.35
Store-bought cooked food	3.26	0.86	0.03
Transportation equiment	2.89	3.83	0.13
Fish and shellfish	2.42	1.09	0.04
Cakes and candies	2.40	1.56	0.06
Cereal	2.39	1.35	0.06
Meat	2.15	0.58	0.03
Household applicances	2.09	6.00	0.28
Recreational goods	2.08	15.71	0.75
Electronics	2.02	8.65	0.42
Non-alcoholic beverages	1.58	2.01	0.12
Alcoholic beverages	1.37	2.89	0.21
Medicine and nutritional supplements	1.37	2.60	0.19
Newspapers and magazines	1.32	0.23	0.02
Dairy products and eggs	1.17	0.14	0.01
Oils, fats and seasonings	1.17	0.45	0.04
Domestic utensils	0.78	3.34	0.42
Furniture and furnishings	0.71	7.18	1.00
Footwear	0.57	4.25	0.73
Books	0.42	1.93	0.45
Total	44.66	99.39	
Services			
Miscellaneous services and remittances	12.94	0.11	0.00
Utilities and telecommunication services	12.37	0.03	0.00
Travel and recreational expenses	6.95	0.46	0.01
Meals outside home	5.79	0.00	0.00
Private transportation insurance and services	4.92	0.00	0.00
School and tutorial fees	3.54	0.00	0.00
Medical services	2.51	0.00	0.00
Public transportation	2.32	0.00	0.00
Repairs and maintenance	2.13	0.01	0.00
Personal care services	1.19	0.00	0.00
Domestic services and laundry	0.67	0.00	0.00
Total	55.34	0.61	

Note: Data are for 2010. Shares are expressed as percentages.

Table 1 makes clear some basic stylized facts of our data. First, because we do not have data for Rakuten's travel business, most of the our coverage of services is negligible.Thus, we will focus on e-commerce's impact on goods prices for all of our main results. Second, within goods categories we see that there are no zeros in the table indicating that at this level of ag-

gregation all categories of goods were available online in Japan in 2010. Third, there is enormous variation in the e-retail 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 transportation equipment, household consumables (which includes non-durable household supplies like paper products and cleaning agents), recreational goods (which includes items like sports equipment and gardening supplies), are sold very intensively online as well. In other words, 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.

In this version of the paper, we only make use of the Rakuten data. However, in future versions, we will also make use of a second dataset: the National Survey of Family Income and Expenditures. The NSFIE is a representative survey of households that reports the purchase channels by product of Japanese households for the same set of products we use in this paper. Starting in 2004, the NSFIE also began a quinquennial recording the expenditure share of each product from online merchants. Since the Rakuten data is based on the point of sale, whereas the NSFIE is based on a survey of household purchases, these two datasets need not yield the same internet intensities of goods. (especially if some of Rakuten sales are to businesses or if there is some misreporting of household purchases). However, in practice the two surveys line up very closely and none of our results are qualitatively affected by how we measure or in what year we measure internet intensity.

There is an important difference between *aggregate* e-commerce sales shares point-of-sale data reported by Japan's Ministry of Economy, Trade and Industry (METI) and the numbers reported in the NSFIE. The difference is likely to due to the way in which the data are collected. METI bases its e-commerce sales shares by conducting a survey of

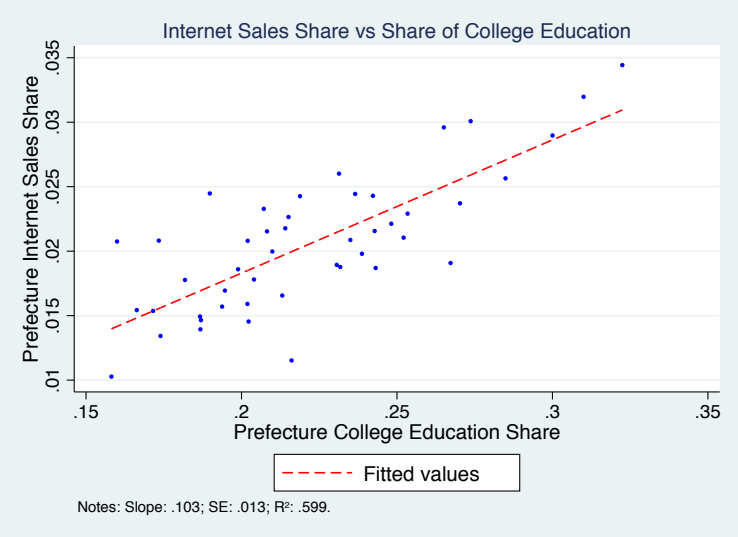
the major e-commerce firms and dividing e-commerce sales by the sales of of all Japanese manufacturers. The NSFIE, however, is based on household surveys and suffers from the problem that a growing share of households have been not answering questions related to where they made their purchases. In particular, in 1999 only 7 percent of expenditures were from undisclosed locations, but by 2014, this number had risen to 13 percent. By contrast, the share of e-commerce sales in the NSFIE is rising over time, but the the level in 2014 was only 2.4 percent of aggregate expenditures, which is about half the METI number and implausibly small given the magnitude of Rakuten sales. After discussing the discrepancies with people knowledgable about Japanese statistics, we came to the conclusion that an important part of the discrepancy is due to households that buy goods online not reporting the sales channel and the purchase channel being left blank.[10] We therefore use the METI data to identify aggregate internet sales shares and the Rakuten or NSFIE data for understanding the *distribution* of internet sales shares.

The NSFIE data also enables us to examine internet purchase intensity by location for the 47 prefectures of Japan. As one can see from Figure 2, there is a very strong association between the share of prefectures expenditures online and the share of a prefecture's population over the age of 15 that went to college. This variable alone explains 60 percent of the variance across prefectures. Interestingly, when we tried to control for other prefectural covariates that we thought might be associated with e-commerce expenditure shares—per capita income, population (urbanization), average age, or other education levels—none of these were even close to statistically or economically significant in any specification in which we also included the share of the population with a college degree. This suggests that e-commerce is largely a technology that serves highly educated people and urban-rural or old-young digital divides are secondary.[11]

[10]We thank Takashi Unayama of Hitotsubashi University for insights into this issue.

[11]Interestingly, Fan et al. [2018] find no link between education and internet sales intensity. One possibility for the difference in the results is that Chinese education levels in their data set are much lower than in our data. 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 and .

Figure 2: Education and E-Commerce Purchases



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 with an aim of 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 bar code, the data leaves limited scope for unobserved quality differences to affect intercity price differentials. For example, Imai et al. [2012] 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 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.[12]

[12]In order to further clean the data, we drop all observations in which the item only appears in one city.

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[13], 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.

Table 2: Summary Statistics for the Sample of Goods

	Mean	Standard Deviation	Min	p10	p50	p90	Max
Δp_{ic1996}	-0.019	0.179	-1.252	-0.244	-0.002	0.170	1.480
x_i	0.046	0.086	0.000	0.000	0.013	0.137	1.000
$p_{ic1991} - \frac{1}{C}\sum_c p_{ic1991}$	0.000	0.151	-0.882	-0.177	0.006	0.173	0.931
$x_i \times p_{ic1991}$	0.368	0.808	0.000	0.002	0.069	1.006	12.154
Observations	25,923						
Δp_{ic2001}	-0.029	0.188	-1.433	-0.254	-0.010	0.158	1.150
x_i	0.047	0.088	0.000	0.000	0.013	0.143	1.000
$p_{ic1996} - \frac{1}{C}\sum_c p_{ic1996}$	0.000	0.162	-0.968	-0.185	0.006	0.178	1.108
$x_i \times p_{ic1996}$	0.378	0.839	0.000	0.002	0.069	1.031	12.211
Observations	27,509						
Δp_{ic2016}	-0.010	0.492	-4.341	-0.428	0.045	0.435	1.886
x_i	0.047	0.091	0.000	0.000	0.013	0.137	1.000
$p_{ic1991} - \frac{1}{C}\sum_c p_{ic1991}$	0.000	0.145	-0.861	-0.168	0.007	0.163	0.900
$x_i \times p_{ic1991}$	0.377	0.833	0.000	0.002	0.075	1.014	12.154
Observations	19,861						

Notes: Prices are in natural log. Δp_{ict} is the five-year log difference in prices; $x_i = \frac{s_i/e_i}{\max_j(s_j/e_j)}$; and the data in the first panel covers the period 1991-1996; the second panel covers 1996-2001; and the last panel covers 1991-2016.

Table 2 reports the sample statistics for our data. As one can see from the table, we typically have close to 30,000 price quotes in any of our five year periods. The first line of the table shows the average rate of inflation across the sample period. As one can see, on average goods prices fell, which reflects the deflation that can be observed in Japan over this

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.

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

time period. The second line reports information on the internet 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-retail intensity of 0.137 over the full sample period, which is more than three times higher than a good with the median intensity. Moreover, at the upper tail of the distribution, we observe goods with an e-retail intensity that is more than 75 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.

One concern about our data is that it is not at the bar-code level. This might create problems if the JSB is sampling goods of very different quality across cities. Hottman et al. [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. We can check whether we have excessive price dispersion (caused by including goods of very different quality within categories) by computing the 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 in the third row of each panel, we find that 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 barcoded goods in the US and Canada suggests that the JSB is probably not including goods that differ substantially in quality when sampling goods in different cities and therefore that quality variation across cities for the same product is unlikely to be a major problem in our data.

4 Estimation

In Section 4.3, we first do some data exploration to first show that price convergence is a central tendency in the data and that there is visual evidence that the internet changed the rate of convergence. We also show that there was no important change in the overall rates of price change for goods available online. 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 4.4. Finally, in Section 4.5, we present our estimates of the welfare gain from e-retail.

4.1 E-Commerce and National Prices

Goolsbee and Klenow [2018] have explored whether we see any differences in rates of inflation for goods available online. Here, we build on this work to show that these differential rates of inflation only appeared after the entry of e-commerce firms. In order to examine this in the data, we regress annual log price changes of goods (Δp_{ict}) on good (i) and city (c) fixed effects along with an indicator variable, D_t, that is one in 1997 (the year Rakuten opened) and zero before as well as the internet sales intensity of the good interacted with this dummy ($x_i D_t$):

$$\Delta p_{ict} = \alpha_i + \beta_c + \phi D_t + \theta x_i D_t + \epsilon_{ict}, \tag{18}$$

The coefficient on D_t (ϕ) tells us whether there was any differential trend in price inflation for goods available online after the entrance of Rakuten and θ, the coefficient on the internet intensity interaction term ($x_i D_t$) tells us about the differential rate of price change for goods traded online after the entry of Rakuten. 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 Rakuten versus the full time period.

Table 3: Relative Price Changes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent Variable	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{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

We present these results in Table 3. Overall, the constant term is negative, which reflects the fact that Japan experienced deflation over much of this time period. In all the regressions with fixed effects, we see a significant negative coefficient on the post-Rakuten internet intensity interaction ($x_i D_t$). Even more interesting is the fact that this effect becomes larger and more statistically significant once we extend the sample to 2016. In terms of economic magnitudes, the results in column 6 imply that a good at the 90^{th} percentile of internet sales intensity had an inflation rates that fell by 0.6 percentage points per year faster after the entry of Rakuten. This is consistent with results in Goolsbee and Klenow [2018] who found that goods traded online have lower inflation rates than goods not available online. However, our results differ in that we show that this difference is observable in the prices in physical stores and that these differential price trends did not exist prior to the entry of e-commerce firms. These results suggest that while there may be some evidence that better logistics, reduced markups, and other factors associated with e-retail might have affected price increases of goods available online.

We can also use these estimates to compute the implied gain from the differential impact on prices estimated in Table 3. If we assume that the entry of e-commerce did not affect the price level in sectors with no e-commerce sales, i.e., we ignore general equilibrium effects, we can obtain an estimate of e-commerce on the price level by computing

$$\Delta p_t^N \equiv \hat{\theta} \sum_i x_i \frac{e_{i,t-1}}{\sum_i e_{i,t-1}} \tag{19}$$

where $\hat{\theta}$ is the estimated coefficient on the internet intensity interaction term in equation (18). Based on the specification in column 7, we estimate that e-commerce lowered the Japanese price level by 0.05 percent per year or about 1 percentage point between 1997 and 2016 due to reductions in markups and efficiency gains. Converting this number to welfare depends on whether one thinks the drop in the Japanese price level was just due to a reduction in retail markups (in which case most of the gain was a transfer from producers to consumers) or a rise in efficiency in which case it was a pure welfare gain.

4.2 Gains in "New Trade Models"

Although Brynjolfsson et al. [2003] computed the variety gains from internet book sales, we know of no study that has done this generally. Fortunately, this is trivial to do with our data. Arkolakis, Costinot, and Rodriguez-Clare (2012) show that in the Krugman (1980), Melitz (2003), and Eaton and Kortum (2002) models, the log change in welfare following a trade liberalization equals $\frac{1}{\epsilon} \ln \lambda$, where $\lambda \in [0, 1)$ is the share of consumer of expenditures on sales from retailers other than e-commerce firms, and ϵ equals the "trade elasticity." The trade elasticity has a different interpretation in different models. For example, in the Krugman (1980) model it equals one minus the elasticity of demand whereas in the Eaton and Kortum (2002) model, it equals the negative of the Fréchet shape parameter. However, in order to make progress, we will use a standard estimate for the trade elasticity of -5.

The computation of the share of household purchases that was made from sources other than e-retailers and be computed straightforwardly. Based on the NSFIE data, we know that the share of household expenditures purchased from all retailers (χ) was 0.62 in 2014, with the remaining expenditures covering utilities, education, and other expenditure items that we will assume do not represent e-commerce. If we couple this with the share of expenditures

in 2014 from all e-retailers, $s = 0.0437$, we we have $\lambda = (1 - s)\,\chi + 1 - \chi = 0.97$. This gives us an estimate of the welfare gain from e-commerce in Japan in 2014 of 0.6 percent. We can also easily do some robustness checks for different values of the trade elasticity. Allowing the trade elasticity to range -3 to -7 results in welfare gains that range from 0.4 percent to 0.9 percent. Similarly, if we had computed the Japanese gains achieved in 2017, we would obtain a number of 0.7 percent. Thus, our point estimate for the welfare gains these e-commerce gains is 0.6, with reasonable estimates ranging from 0.4 percent to 0.9 percent.

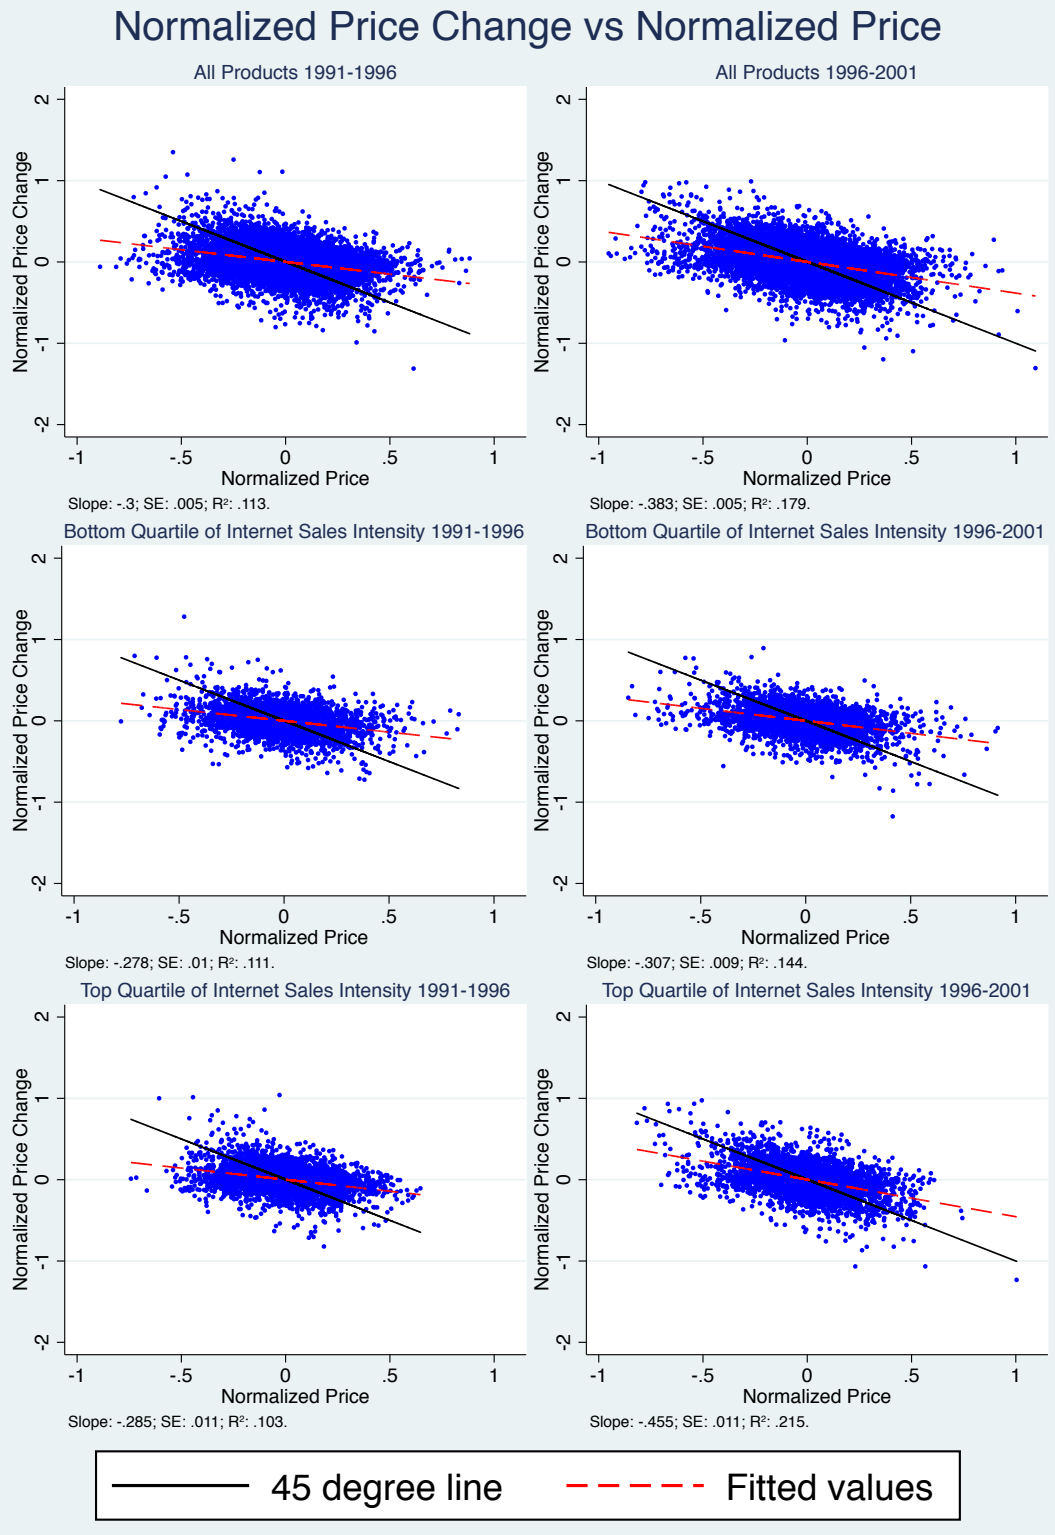
4.3 Gains Due to Price Arbitrage

In order to visualize what is going to drive our results, we first consider two five-year periods. The first five-year period (1991-1996), pre-dates the formation of Rakuten by at least a year, so we can call this period the "pre-Rakuten 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 Δp_{ict} and p_{ict} on product and city fixed effects and construct normalized price changes ($\Delta^5 p_{ict} - \hat{\alpha}_{it} - \hat{\beta}_{ct}$) and normalized price levels ($p_{ic,t-5} - \hat{\alpha}'_{it} - \hat{\beta}'_{ct}$), where $\hat{\alpha}_{it}$ ($\hat{\alpha}'_{it}$) and $\hat{\beta}_{ct}$ ($\hat{\beta}'_{ct}$) are the estimated fixed effects from the regression of Δp_{ict} (p_{ict}) 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 ($p_{ic,t-5} - \hat{\alpha}'_{it} - \hat{\beta}'_{ct}$).

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

Figure 3: Normalized Price Change vs. Normalized Price



tended to have lower rates of inflation than goods with low prices. This mean reversion in the data is the convergence that is likely to be product of price arbitrage. As one can see from these two plots, there is the average rate of price unconditional rate price convergence is about 30 percent and rose to 39 percent in the four-years after the entry of Rakuten. 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. One plausible reason for the weaker evidence of price convergence in the U.S. is that that the American Chamber of Commerce Researchers Association data used in this study 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]).

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 2010 ($x_i < 0.002$) and the set of goods in the highest quartile of the distribution ($x_i > 0.054$). 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 Rakuten platform. The slope of the line for goods not available online in the early period is -0.28, which is almost identical to the slope in the pre-Rakuten period (-0.30) and the slope in the post-Rakuten period for goods not traded much online (-0.29). In other words, the entry of Rakuten did not seem to have had much of an effect on goods not traded on its platform. However, we see a very different pattern for goods with an internet intensity in the upper quartile of the distribution. The slope steepens by 50 percent, rising in magnitude from -0.31 to -0.46. In other words, Figure 3 provides the visual version of the result we will explore econometrically in Section 2.1—following the entry of Rakuten, the rate of price convergence for goods traded intensively rose substantially.

Table 4: Estimates Over Period 1991-2001

	(1)	(2)	(3)	(4)
Dependent Variable	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}
Lagged Price	-0.299***	-0.343***	-0.325***	-0.136***
	(0.006)	(0.006)	(0.004)	(0.002)
Internet Intensity	-0.037		0.124*	0.001
× Lagged Price	(0.070)		(0.068)	(0.023)
Internet Intensity		-0.820***	-1.061***	-0.258***
× Lagged Price × Post Rakuten		(0.068)	(0.085)	(0.028)
t	{1996}	{2001}	{1996,2001}	Annual 1992-2001
Observations	25,923	27,509	51,185	152,958
R^2	0.52	0.53	0.52	0.51

4.4 Estimating Convergence Rates

We now turn to obtaining estimates of the impact of Rakuten on the rate of intercity price convergence. Table 4 presents the results of estimating equation (1) for five- and one-year intervals. In the first column, we present separate regressions for a specification analogous to what we showed in the plots. 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. Comparing the first rows of columns 1 and 2 reveals the convergence rates for goods not suitable for online sales (i.e., those where $x_i = 0$) were almost identical before and after the entry of Rakuten. The coefficient on internet intensity interacted with lagged prices ($x_i p_{ic,t-5}$) in column 1 indicates that the rate of convergence for goods suitable for online sales was not significantly different than the convergence rate of other goods prior to to the entry of Rakuten. However, the negative and significant coefficient on the interaction term ($D_t x_i p_{ic,t-5}$) in the post-Rakuten sample (where we dropped the $x_i p_{ic,t-5}$ term from the specification because we do not have any pre-Rakuten observations) indicates a significantly faster rate for goods available online after Rakuten entered the market. This specification confirms the basic

insight about the increase in the rates of convergence of goods sold intensively online after the entry of Rakuten that we saw in Figure 3.

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 internet intensity coefficient. As one can see from the table, the coefficient is negative and precisely measured. Not surprisingly, the estimated coefficient on $p_{ic,t-k}$, $\hat{\gamma}$, doesn't change much, and we continue to get a negative and significant coefficient on the Rakuten sales intensity interaction term ($\hat{\delta}_2 = -1.06$). Interestingly, the estimate of δ_1, the differential rate of price convergence for e-retail intensive goods flips sign and becomes small and positive, which implies that if we pool the data over the two time periods, goods suitable for online trade had, if anything, a slightly slower rate of price convergence than goods not available online.

One way to assess the economic significance our results is in terms of price-convergence half-lives: i.e., the amount of time it takes for half of any price differential to disappear. Although five-year differences lend themselves easily to the standard differences-in-differences specification, it is not standard to discuss half-lives in terms of five-year periods. The conventional way to estimate half-lives is to use higher frequency data. In column 4 of Table 4, we redo our estimation at an annual frequency. The coefficients fall in magnitude because a given percentage movement in prices over a five-year period implies a smaller movement over any single year. However, we observe the same pattern in the data. The entry of Rakuten is associated with a significant increase in inter-regional rates of price convergence.

We next compute the half-lives for three variants of our specification: one for goods not traded online ($x_i = 0$); one for goods in the 90^{th} percentile of e-retail intensity in the pre-Rakuten period ($x_i = \text{p90}, D_t = 0$); and finally one for one for goods in the 90^{th} percentile of e-retail intensity in the post-Rakuten period ($x_i = \text{p90}, D_t = 1$). For goods not sold online ($x_i = 0$), the coefficient of -0.135 on lagged prices corresponds to a half-life for price

differences of 4.74 years. Similarly, we also see that prior to the entry of Rakuten (i.e., goods for which $x_i = \text{p90}, D_t = 0$) had a similar half-life of 4.75 years. Thus, prior to the entry to Rakuten, we see only a marginal difference in convergence rates for goods available online and off. By contrast, when Cecchetti et al. [2002] estimated this half-life using U.S. data, they found it to be nine years: almost twice as long. Similarly, Parsley and Wei [1996] found a unit root, i.e., an infinite half-life, when running a similar regression with city fixed effects. The differences in these results probably highlights the intercity comparability of the price data in Japan. However, there is a much larger impact of Rakuten on the convergence rate for Japanese goods sold online ($x_i = \text{p90}, D_t = 1$). For these goods, the half-life fell from 4.8 years to 3.6 years after Rakuten's entry: a 25 increase in to convergence rate.

Table 5: Estimates Over Alternative Periods

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{ict}	Δp_{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***
× Lagged Price	(0.073)	(0.078)	(0.072)	(0.073)	(0.025)
Internet Intensity	-0.950***	-1.863***	-1.142***	-1.011***	-0.348***
× Lagged Price × Post Rakuten	(0.091)	(0.091)	(0.085)	(0.078)	(0.026)
t	{1996,2006}	{1996,2011}	{1996,2016}	{1996,2001, 2006,2016}	Annual 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

There are a number of potential problems with the evidence that we have just presented. One obvious concern is that our measure of e-retail suitability dates to 2010 (the earliest year we could obtain), but we don't know these numbers for 2001. There are potentially two concerns arising from this issue. The first is that we have mismeasurement in e-retail suitability (x_i). If this is correct, one should expect that measurement error in e-retail suitability would cause us to underestimate the effects of this variable on price convergence.

A second concern is that perhaps there is something special about the first five years after Rakuten entered the Japanese market. In Table 5, we also examine alternative time periods following Rakuten's entry. In the first three columns, we do a differences in differences based comparing the five years prior to Rakuten's entry (1991-1996) with three alternative non-overlapping periods: 2001-2006, 2006-2011, and 2011-2016. These results confirm that the choice of end period does not matter qualitatively for our basic finding that the rate of price convergence for goods available online rose after the entry of Rakuten. Finally, in the last column of Table 5 we repeat the estimation over the full period (1992-2016) at the annual frequency. The results are similar to the estimates over the shorter time period (1992-2001) reported in Table 4.

The results we have presented so far have only included a sample of goods because prices are easier to measure for goods than services and many services are hard to provide at a distance. However, Rakuten also conducts e-commerce activities in the services sector, so a reasonable question might be whether the results change significantly if we include all goods and services in the analysis. In Table 6, we repeat our analysis using the universe of all goods and services in the Japanese consumer expenditure survey. The results are quite similar to the set of results we obtained only looking at a sample of goods. The coefficient on the post-Rakuten, internet-intensity interaction term does not move significantly, indicating that our results are robust to the choice of sample.

Table 6: Robustness Check Using All Goods and Services

	(1)	(2)	(3)
Dependent Variable	Δp_{ict}	Δp_{ict}	Δp_{ict}
Lagged Price	-0.297***	-0.353***	-0.141***
	(0.004)	(0.003)	(0.001)
Internet Intensity	-0.052	0.145**	0.025
× Lagged Price	(0.065)	(0.070)	(0.024)
Internet Intensity	-1.061***	-1.013***	-0.329***
× Lagged Price × 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

4.5 Welfare Gain

Aggregate consumer gains due to faster price convergence can be calculated from the equation (16) and (17). 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 the welfare gain using an elasticity of six, which corresponds to a trade elasticity of -5. and note that halving this elasticity would halve the welfare gain and doubling it would double the gain. In all cases, we base our estimates of the impact of e-retail on the rate of convergence on Table 5 column 5. All of the data has been converted into 2016 yen. The first columns shows the estimated welfare gains due to price convergence in that year and the second column gives the counterfactual welfare gain that would have occurred if price convergence for goods available online had remained at the pre-Rakuten rate (i.e., $(D_t = 0)$) as given by equation (15). We see that price convergence across regions during this time period led to an average welfare gain of 222.0 billion yen per year (about 2 billion dollars) for residents of our sampled cities. In the second column, we compute the counterfactual gain that would have occurred if the speed of convergence had remained at

the pre-Rakuten rate. This gain is somewhat lower: only 196.6billion yen per year. Thus, the difference between these two columns—25.4 billion yen— (approximately 225 million dollars) constitutes the average annual welfare gain for consumers in our sample of cities.

Table 7: Counterfactual Welfare Gain

Year	$\Delta\hat{W}_t^R(D_t=1)$	$\Delta\hat{W}_t^R(D_t=0)$	$\Delta\hat{W}_t^E$	Total Expenditure	Expenditure on Goods
1997	193.4	171.2	22.2	32,287	20,690
1998	185.5	163.9	21.6	32,076	20,557
1999	200.1	177.0	23.1	35,561	22,231
2000	235.6	208.3	27.3	39,809	24,781
2001	243.4	217.9	25.5	38,222	23,716
2002	211.5	188.6	22.9	38,635	23,928
2003	218.7	193.8	24.9	38,564	24,365
2004	230.7	205.5	25.2	37,897	23,995
2005	222.0	197.0	25.1	42,102	25,538
2006	234.2	207.7	26.6	41,178	25,114
2007	266.9	238.3	28.6	41,483	25,137
2008	250.5	222.0	28.5	40,850	25,073
2009	246.8	218.8	28.0	40,982	25,410
2010	243.8	216.1	27.7	43,400	27,975
2011	235.9	207.6	28.3	42,339	27,429
2012	227.7	201.0	26.7	42,849	27,548
2013	211.0	186.1	24.9	43,501	27,505
2014	197.9	174.0	23.9	42,550	27,174
2015	197.6	173.7	24.0	40,698	26,964
2016	186.9	163.4	23.5	40,210	26,809
Average	222.0	196.6	25.4	39,760	25,097

Note: Unit is in billions of yen.

This number is difficult to interpret because it is only computed for residents in our sampled cities. To obtain a sense of how much this matters for welfare, we should deflate the number by the total amount of expenditures of our sampled households, which is reported in the fifth column of Table 7. We obtain an estimate of the welfare gain which equals 0.06 percent of consumer expenditures. [14]

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

Table 8: Counterfactual Welfare Gain pre-Rakuten Period

Year	$\Delta\hat{W}_t^R(D_t=1)$	$\Delta\hat{W}_t^R(D_t=0)$	$\Delta\hat{W}_t^E$	Total Expenditure	Expenditure on Goods
1992	170.2	149.9	20.2	33,020	21,001
1993	174.8	155.1	19.7	32,613	20,745
1994	161.1	142.9	18.2	32,368	20,582
1995	165.1	145.6	19.4	32,844	21,049
1996	190.8	169.7	21.1	32,848	21,055
Average	172.4	152.7	19.7	32,738	20,886

Note: Unit is in billions of yen.

A concern that one might have about this counterfactual is that the level of price dispersion itself might be a function of the existence of e-retail. One simple way of seeing how important shifts in the distribution of prices are for understanding the welfare gains is to conduct the counterfactual using the pre-Rakuten period (1991-1996). The shocks to prices during this period were independent of e-retail and hence the counterfactual answers the question, "what would the gain to consumers have been given the pre-Rakuten distribution of prices if Rakuten had entered in those years." We report the results from this exercise in Table 8. Our estimates imply that had Rakuten entered the Japanese market five years earlier, the welfare gain would have been about 19.7billion yen per year for consumers in our sample of cities. This estimated gain is comparable to what we observed in the post-Rakuten period (25.4 billion), indicating that it doesn't make a substantial difference whether one computes the welfare gain based on price differential that existed before or after Rakuten entered the market.

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 using a number of popular modeling techniques. First, we consider the possibility that e-commerce

may have lowered prices by driving down the average prices of goods available online. Second, following Brynjolfsson et al. [2003] and Arkolakis, Costinot, and Rodriguez-Clare (2012), we compute the welfare gains due to the ability of e-commerce to enable consumers to purchase goods from other regions. Third, following Jensen [2007], we compute the gains that arise through e-commerce's ability to arbitrage intercity price differences. We find that all three channels produced welfare gains in Japan, but our estimates suggest that the first and second channels are by far the most important, with welfare gains through these channels being eleven to sixteen times larger than through the price arbitrage channel. Overall, we find that increased inter-city arbitrage raised Japanese welfare by 0.06 percent, the gains due to new varieties available through online shopping raised welfare 0.7 percent, and the gains due to overall price reductions for goods available online raised welfare by at most 1 percent.

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.

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.

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.

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, Chihiro Shimizu, and Tsutomu Watanabe. How fast are prices in japan falling? *Discussion papers 12075, Research Institute of Economy, Trade and Industry (RIETI)*, 2012.

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.

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.

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