Paper:
Gorodnichenko, Y. & Talavera, O. (2017). Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration. *American Economic Review, 107*(1), 249-282.
http://dx.doi.org/10.1257/aer.20141127

Price setting in online markets: Basic facts, international comparisons, and cross-border integration*

BY YURIY GORODNICHENKO AND OLEKSANDR TALAVERA

Abstract

We document basic facts about prices in online markets in the U.S. and Canada, which is a rapidly growing segment of the retail sector. Relative to prices in regular stores, prices in online markets are more flexible and exhibit stronger pass-through (60-75 percent) and faster convergence (half-life less than 2 months) in response to movements of the nominal exchange rate. Multiple margins of adjustment are active in the process of responding to nominal exchange rate shocks. Properties of goods, sellers and markets are systematically related to pass-through and the speed of price adjustment for international price differentials.

JEL: E3, F3, F40, F41
Keywords: Online markets, prices, pass-through, border effects, law of one price.

* Yuriy Gorodnichenko: Department of Econimics, University of California – Berkeley, Berkeley, CA 94720-3880, USA (ygorodni@econ.berkeley.edu). Talavera: Management School, University of Sheffield, Sheffield, S10 1FL, United Kingdom (o.talavera@shef.ac.uk).We thank Viacheslav (Slavik) Sheremirov for excellent research assistance. We are deeply grateful to Alberto Cavallo for his feedback and help with the data. We also thank Olivier Coibion, Robert Deckle, Johannes Wieland, and participants of numerous conferences and seminars for comments. Gorodnichenko also thanks the NSF (1151168) and the Sloan Foundation for financial support.

E-commerce is a rapidly increasing segment of the retail market. The U.S. Census Bureau estimated that total e-commerce sales for 2013 were $263.3 billion, which is approximately 5.6 percent of total retail sales in the U.S. economy.[1] This represents an increase of 16.9 percent from 2012, while total retail sales increased by 4.2 percent in 2013; this pattern is consistent with historical trends: online sales have grown much faster (10 or more percent) than sales of brick-and-mortar stores. Forrester Research, an independent technology and market research company, predicts that by 2016, online sales will account for more than 9 percent of total retail sales.[2] While e-commerce is young, its digital presence is a major force revolutionizing retail as we know it: according to Deloitte (2015), the internet is projected to influence 64 percent of in-store retail sales by the end of 2015. To the extent that market valuation reflects prospects of companies, stock market participants believe that Amazon.com has a brighter future than Walmart (even though Amazon.com has only a quarter of Walmart's revenue) and that the future of retail is in online markets.

However, despite a significant and rapidly expanding share of e-commerce, the properties of online prices are still relatively understudied, even though these prices can shed new light on a number of key puzzles. Indeed, online markets have unique characteristics. For example, the physical cost of changing prices is negligible for internet stores, and therefore internet prices can fluctuate every instant (e.g., minute, day, week) in response to shifting demand and supply conditions. Searching for best online prices for very narrowly defined goods is particularly cheap and simple as consumers do not need to travel anywhere, buyers can establish the distribution of prices with just a few clicks, and pressure for price convergence is especially strong with ubiquitous price comparison websites

[1] For the same period, U.S. manufacturers reported e-commerce shipments of $3.3 trillion, which is 57 percent of all manufacturing shipments. See U.S. Census Bureau (2015).

[2] These patterns are very similar in other developed countries. For example, according to the Centre for Retail Research, online retail sales in Europe jumped 20 percent this year, far outstripping the 1.4 percent growth in store-based sales. Furthermore, the share of online sales in total sales is larger in Europe than in the USA. For instance, the share is 9.5 percent in the U.K.

(PCWs). More generally, the geographical location of consumers and stores is largely irrelevant in e-commerce, and therefore administrative borders and similar frictions are likely to play a much more limited role.

These special properties of online markets can help understand why pass-through of exchange rate fluctuations and reversion to the law of one price are generally weak in international data and thus constitute one of the central puzzles in international economics (Obstfeld and Rogoff 2000). In a highly integrated market with low frictions of price adjustment, easy search and price comparisons, and limited influence of geographical barriers, one can rule out some popular explanations of the puzzle and narrow down a set of plausible theories. Using internet prices in the U.S. and Canada for a broad array of products, we try to exploit these insights and provide new evidence on the nature and sources of frictions in price adjustment and departures from the law of one price.

To document and study the properties of online prices, we have constructed a unique dataset of price quotes. Specifically, we gathered prices and other relevant information from a leading PCW for a duration of 5 years. The data include each good's unique identifier (similar to barcodes in the scanner price data), each good's description, prices for each seller, each seller's unique identifier, the number of seller reviews, the ranking of seller quality, reviews of goods, etc. The dataset covers a broad range of goods that are sold online, including software, electronics, tools, computer parts, and photo equipment. We have collected information for more than 115,000 goods and nearly 20 million price quotes.

There are several advantages of using our data. First, the time span (almost 5 years) is considerably longer than the time span usually available for researchers studying online prices (typically a year or less). This dimension is important when we study dynamic properties of prices, such as duration of price spells, speed of price convergence, and pass-through. Second, the coverage of goods is much broader than in previous analyses of online prices, which typically have focused on books and CDs. The latter types of goods are easy to compare across sellers or countries, but they also have a number of unusual properties that make generalizations difficult. Our dataset is heavily populated by durable goods that tend to be under-represented in typical scanner price data and that are much more likely to be traded and moved across distant locations. Third, we collected prices for identical goods in the U.S. and Canada so that comparison of prices is direct and simple. Thus, we can avoid a number of pitfalls associated with comparing price indexes or goods that are only broadly similar. Fourth, our data include information on important attributes such as the reputation of sellers and goods as revealed by ratings of sellers and products. We can use these attributes to explore the predictors of pass-through and speed of price adjustment for online prices. In contrast, previous research on basic properties of prices had only very limited (if any) information about characteristics of goods for which prices were available. Fifth, our data include many sellers—most stores in our sample sell goods only online and do not have conventional, brick-and-mortar retail outlets (e.g., Amazon.com)—rather than one retail chain; therefore, we can assess the relative importance of different sources of price variation. This multi-seller dimension is important because branches of a single seller are less likely to engage in competition between each other than with branches of different sellers. Finally, the high frequency of our data allows us to time reactions of prices to other high frequency events such as changes in the exchange rate or natural experiments, thus making identification more clear-cut.

Using this dataset, we report properties of various pricing moments (e.g., the frequency and size of price changes) in e-commerce and thus complement earlier studies (e.g., Nakamura and Steinsson 2008) that present the same information for regular, brick-and-mortar stores. We find that the size of price changes in online stores (approx. 4 percent) is less than half the size of price changes in regular stores (approx. 10 percent). We also find that price changes occur much more frequently in online stores (approx. once every 3 weeks or less) than in regular stores (once every 4-5 months or more). This evidence is consistent with the view that online prices are much more flexible than prices in regular stores. However, the fact that we still observe some rigidity in online prices suggests that the costs of changing prices are more complex than just physical menu costs and instead are likely to involve costs of gathering and processing information as well as potentially coordinating price changes with customers, suppliers, or other sellers. We also document that price dispersion is substantial and persistent, even for very narrowly defined goods. For example, the average standard deviation of log prices in a given week for a precisely defined good at the bar-code level is between 0.13 and 0.16.

Once these basic facts are established, we study the sensitivity of online prices to fluctuations of the nominal exchange rate. Since adjustment of online prices is unlikely to have any physical costs, and with easy shipping the physical location of the seller is much less important, pass-through could be quick and nearly complete, while it can be slow and partial in the prices of regular stores because of the frictions associated with trade flows and mobility of buyers. We find that, on average, pass-through in online markets is incomplete but large and amounts to approximately 60-75 percent, which is greater than the 20-40 percent pass-through documented for regular markets. The speed of price adjustment to equilibrium levels is substantially faster in online markets (half-life is about 2-2.5 months) than in regular markets (half-life varies from 3 quarters to a few years).

There is significant heterogeneity in pass-through and the speed of price adjustment across goods. Using the richness of our data, we show that for goods with certain characteristics, pass-through can be close to 100 percent. We also document that the size of pass-through and the speed of price adjustment are systematically associated with the degree of price stickiness, turnover of sellers, returns to search, synchronization of price changes, reputation of sellers, and the degree of competition. These results help reconcile the heterogeneity of estimated pass-throughs and the speeds of adjustment across studies and provide new facts for theoretical models to match.

This paper is related to several strands of research. The first strand is focused on assessing whether the law of one price (or its milder versions such as the purchasing power parity (PPP) hypothesis) holds and how quickly deviations from the law of one price are eliminated. The early generation of this literature could use only price indexes collected at the country or regional level, which led to a number of practical and conceptual issues with the interpretation of the results. Rogoff (1996) summarizes this literature as documenting that PPP is likely to hold in the long run, but it takes a long time for prices to converge to the PPP (i.e., the half-life is routinely estimated to be over a year and in most cases multiple years). This literature also found that deviations from PPP can be quite large and heterogeneous across countries and time (e.g., Takhtamanova 2010, Campa and Goldberg 2005, Barhoumi 2005) which can be only partially explained by sticky prices and exchange rate regimes, constituting the PPP puzzle.

Data limitations of the first strand motivated the second generation of studies, which focused on using micro-level price data to measure pass-through and the speed of price adjustment for goods defined more precisely. Imbs et al. (2005, 2010), Crucini and Shintani (2008), Broda and Weinstein (2008), and others showed that pass-through and the speed of price adjustment are higher when prices for narrowly-defined goods are considered: the half-life of price adjustment falls to about a year. These papers demonstrate that the PPP puzzle observed in price indexes can be explained at least to some extent by aggregation biases. We contribute to this literature by examining the behavior of prices at the level of precisely defined goods sold by multiple stores in different countries in a market with arguably low frictions.

Easier access to micro-level price data also allows the exploration of the predictors of pass-through and the speed of price adjustment. For example, Menon (1996), Kardasz and Stollery (2001), Gaulier, Lahreche-Revil, and Mejean (2006), Bachis and Piga (2011), Goldberg and Hellerstein (2013), and Mayoral and Gardea (2011) relate market structure, market power (including adjustment of mark-ups), tariffs, presence of multinationals, and importance of non-traded inputs for price stickiness of final goods and the size of pass-through.

We contribute to this literature by exploring the predictors of pass-through and the speed of price adjustment for online markets.

The third strand of research is focused on documenting price rigidities at the micro-level, which can be used later to calibrate macroeconomic models (see, e.g., Nakamura and Steinsson (2008). Studies in this literature concentrate almost exclusively on prices collected in regular, brick-and-mortar stores. In contrast, we focus on online prices, which describe a rapidly growing part of the retail sector. Online prices will play an increasingly important role in the future; therefore, macroeconomists should incorporate properties of a broader set of goods including goods sold online when they characterize micro-foundations of their macroeconomic models. To this end, we complement Cavallo (2015) by covering a different set of goods (i.e., most durables in our data and mostly grocery items in his).

The fourth strand of research documents basic facts about properties of online prices. In a study representative of this literature, Brynjolfsson and Smith (2000) compare online and conventional-store prices for books and CDs. They find that online prices are 9-16 percent lower than prices in regular stores, and the changes in prices are much smaller for online prices, yet quotes of internet prices are quite dispersed, even for precisely defined goods. Much of the subsequent literature has tried to, mostly theoretically, explain the dramatic dispersion of prices in online markets (e.g., Baye and Morgan 2001, 2004, 2009, Morgan, Orzen, and Sefton 2006) by information frictions (e.g., bounded rationality), sellers' ability to discriminate consumers (e.g., based on what sellers know about customers; see Deck and Wilson (2006)), and differences in advertisement (e.g., investment in building brand reputation). We complement this literature by covering a broad set of goods and provide evidence that considerable price dispersion in online markets is a typical characteristic.

The most relevant studies to our paper are Lünnemann and Wintr (2011), Boivin, Clark, and Vincent (2012), and Cavallo, Neiman, and Rigobon (2014). Lünnemann and Wintr (2011) document stickiness of online prices in the U.S. and large European markets (Germany, France, Italy, and the U.K.). They find that internet prices are more flexible than their offline counterparts with half of the spells ending within a month. While Lünnemann and Wintr (2011) have online price data for multiple countries, they do not study the behavior of international price differentials. In contrast, Boivin, Clark, and Vincent (2012) focus on the dynamics of online price differences across three online book sellers in Canada and the U.S.: Amazon.com (and Amazon.ca), BN.com (Barnes & Noble website), and Chapters.ca. They find that price differentials (or relative quantities) for books *do not* react to fluctuations in the relative price of foreign competitors following exchange rate movement; this is consistent with extensive market segmentation and pervasive violations of the law of one price. Similar to Boivin, Clark, and Vincent (2012), Cavallo, Neiman, and Rigobon (2014) collect online prices for four large retailers (Apple, H&M, Zara, and IKEA) in a number of countries and document that the violations of the law of one price—for example, they compare prices for a given IKEA product in IKEA websites in Germany and Sweden—appear only for countries outside currency unions and arise at the time goods are introduced rather than at later stages of product life. We merge these lines by exploring a larger, complementary set of goods (including coverage of generic and branded products) using longer time series and price quotes from multiple sellers, exploiting significant movements in the nominal exchange rate, and investigating predictors of observed pass-through and the speed of price adjustment.

The rest of the paper is structured as follows. In the next section, we describe the dataset and how it was collected. In Section 3, we document the basic properties of online prices. In Section 4, we do extensive international price comparisons and estimate pass-through and the speed of price adjustment for online prices. In addition, we explore the predictors and various margins of price adjustment in response to changes in the nominal exchange rate. In Section 5, we discuss our results and make concluding remarks.

I. Data Description

A. Data collection

This study uses data collected from a PCW that provides price quotes for two countries: USA (.com domain) and Canada (.ca domain).[3] Styles of pages with price quotes are similar across countries, which simplifies data extraction and identification of exactly identical products listed by Canadian and U.S. sellers. Identifiers for goods listed on the website are similar to barcodes used in the analysis of scanner price data. For example, manufacturing product number (MPN) 0S03110 uniquely identifies Hitachi Touro Mobile Pro Portable External 750 GB 2.5” Hard Drive. Figure 1 shows screenshots of typical web pages from PCWs.

Although the price comparison platform we use has similar websites in other countries, we limit the set of countries to the U.S. and Canada for several reasons. First, the link between the U.S. and Canadian websites greatly simplifies linking goods across countries. Second, trade flows are more likely to be affected by trans-ocean shipping costs, language differences, etc. if we compare prices in, for instance, Japan and the U.S. Finally, we want to study countries with strong trade ties. The U.S.-Canada pair is ideal in this respect as flows of goods and services between these two countries are strong even for online markets. For example, Statistics Canada (2013) reports that 63 percent of Canadian online shoppers placed an order with a U.S. online store in 2012. This is comparable to the 82 percent share of Canadian online shoppers who placed an order with a Canadian online store.

In contrast to a few previous studies that investigate properties of online prices and typically have up to one year of data (e.g., Lünnemann and Wintr 2011), our data cover nearly five years. The data collection was launched on November 16, 2008 and continued until September 2013. Importantly, this

[3] The U.S. part of the website was among the top 10 Web portals based on total unique visitors in January 2010. Comscore, January 2010. The website reported in 2012 that tens of millions of people visited it every month.

timeframe includes a period of significant appreciation of the Canadian dollar against the U.S. dollar from 1.30 in the end of 2008 to 0.95 in the middle of 2011 (see Figure 2). A longer time series combined with significant changes in the exchange rate will help us to obtain precise estimates.

Every Saturday at midnight, a Tcl/python script was triggered to collect webpages with price information. The script has several stages. First, it collects information on the universe of goods available for a given type of goods on the comparison website. For each good, there exists a link to a unique webpage with price quotes. The script constructs a dictionary of goods and associated links. Second, the script follows the links and downloads web pages with price quotes. It usually takes about 24 to 48 hours to download a complete set of pages for all goods in targeted categories. Third, after the web pages are downloaded, the Python part of the script extracts a good's description, unique manufacturing product number (MPN), prices for each seller, and sellers' unique ids from every webpage. Our price quotes are net prices (i.e., prices *before* taxes and shipping/handling costs). Figure 3 shows an example of price quotes extracted from the web pages for a good popular in the U.S. and Canada. Whenever possible, we also collected gross prices (i.e., net prices plus taxes and shipping/handling costs) where the destination was an address in Berkeley, CA. Gross prices are available for about one half of net price quotes.

In the end, we obtained information for more than 115,000 goods and nearly 20 million good-seller-week-country quotes. Our price data cover 55 types of goods in four main categories: computers (20 types, e.g., laptops), electronics (13 types, e.g., GPS), software (12 types, e.g., computer games), and cameras (10 types, e.g., digital cameras). Table 1 presents the list of categories and types of goods in our sample.[4] The majority of stores only operate online (Table 2), but

[4] The price comparison website used in this study has been introducing more detailed categories over time. To ensure consistency in our data, we use the classification of goods available at the time when we started to collect our data. Our choice of product coverage was motivated by several considerations. First, we wanted to cover goods where having sellers in the U.S. and Canada was common. For some categories such as clothes, furniture, etc., it is a tangible restriction because many of these goods are local (e.g., flip-flops for Californians) and are branded or sold exclusively in one

there is also a significant presence of stores selling both online and offline. While we have a wide distribution of store sizes, the top 5 percent of sellers by size account for approximately 90 percent of price quotes in our data. Appendix D provides additional details on the properties of the data set. The selection of goods, length of the time sample, and variation in exchange rates in our time sample provide us with a number of advantages relative to what researchers used in previous studies.[5]

First, our dataset covers a relatively diverse set of goods, while the vast majority of papers on online prices almost exclusively focus on books or CDs for which it was relatively easy to ensure that the same good is compared across sellers. Prices of these goods have, however, a number of unusual properties, such as very long spells of constant prices. Furthermore, the market for books and CDs is dominated by a handful of major sellers, such as Amazon.com and Barnes&Noble. Thus, it may be hard to generalize results beyond books and CDs. The diversity of goods in our sample will be essential when we study predictors of the size of exchange rate pass-through and the speed of price adjustment.

Second, a great deal of research on the law of one price has used data on goods for which transaction costs for cross-border purchases are likely to outweigh even large departures from the law of one price. For example, consumers are unlikely to directly take advantage of arbitrage opportunities in grocery products, which are typically available in scanner price data or cost-of-living surveys (e.g., Economist Intelligence Unit). In contrast, we focus on

---

country. Second, we had to select categories where goods have an identifier akin to the universal product code (UPC) because we need to link goods over time and across countries. For some categories (e.g., furniture, toys, food), this restriction was a barrier in earlier years because the coding was missing or not sufficiently detailed to ensure that the identifier is unique. Third, we didn't want to cover books, CDs/DVDs because these goods are unusual in many respects.

[5] We have no information on the quantities of goods bought at quoted prices, and some price quotes may be irrelevant for consumers. However, in another dataset with online quotes and clicks associated with these quotes, Gorodnichenko, Sheremirov and Talavera (2014) found that pricing moments are qualitatively similar for equally weighted price quotes and for price quotes weighted by clicks. Because click-weighted moments point to more price flexibility, one may interpret our results as a lower bound on how quickly prices adjust to movements in the exchange rate.

goods for which transaction costs are small and consumers are essentially free to exploit even small arbitrage opportunities. Indeed, goods in our sample are durable, standardized, and easy to ship. Most goods in our sample are produced outside the U.S. or Canada, and marginal cost shocks can be effectively differenced out when we take the ratio of Canadian and U.S. prices. These qualities are also likely to limit the importance on non-tradables, which often account for a significant share of the cost of selling goods in regular stores.

Third, goods in our data are precisely defined; therefore, one can be more certain that he or she compares prices of the same good when he or she contemplates a purchase. For example, we treat red and blue iPods that otherwise share exactly the same technical characteristics as separate goods. This contrasts with previous research using price indexes or prices for broadly defined goods (e.g., toothpaste).

Fourth, our dataset collects price quotes from multiple sellers while previous research (e.g., Gopinath et al. 2011, Cavallo, Neiman, and Rigobon 2014) typically used micro-level price data from one seller (e.g., because scanner price data are supplied by one retail chain). This aspect is potentially important because branches of the same seller in different countries (e.g., Amazon.com and Amazon.ca) are less likely to compete with each other than outlets of different sellers (e.g., Amazon.com and Rakuten.com). Our data covers a broad spectrum of sellers, such as large general stores (Amazon, Newegg), large specialized or branded stores (B&H or Dell), and niche stores (Memory4less).

Finally, data are collected at weekly frequency; this allows us to study responses of prices at relatively high frequency and makes identification cleaner.

At the same time, one should bear in mind limitations of our data. First, the composition of goods in our sample is skewed towards electronics. While this makes our analysis potentially specific to the electronics market, this market is sufficiently large to be interesting in itself. According to the estimates of the U.S. Census Bureau[6], 30 percent of revenue in e-commerce retail in 2008-2009

[6] http://www.census.gov/econ/estats/2013/all2013tables.html, Historical Table 5.

was generated by categories we cover (i.e., computer hardware, computer software, electronics and appliances, office equipment and supplies). The share declined to 20 percent in 2013 as other categories of goods penetrated e-commerce, but goods in our sample continue to be a major market in internet retail. Furthermore, Gorodnichenko, Sheremirov, and Talaver (2014) document that properties of online prices relative to offline prices are similar for electronics and other product categories; thus, one may expect our results to generalize.

Second, price quotes listed on the PCW may be not representative of prices offered by online stores. Indeed, competition on PCWs is fierce, and PCWs often charge per click or per listing. As a result, stores may choose to post only their best prices on PCWs. Such behavior can affect some moments of the data (e.g., cross-sectional price dispersion). While this pattern is certainly a valid concern if one is interested in the distribution of *all* price quotes, the issue is likely to be insignificant if one is interested in the behavior of price quotes at which consumers make purchases. There is considerable evidence (e.g., Baye et al. 2009, Chevalier and Kashyap 2011, Gorodnichenko, Sheremirov, and Talaver 2014) documenting that transaction prices are heavily concentrated in the competitive (bottom) part of the price distribution so that prices listed on PCW are likely close to transaction prices. As a result, our data are suitable for analyzing international price comparisons but may provide a potentially distorted picture of the micro-level properties of *all* online prices.

Third, most of the sellers in our sample are online-only (see **Error! Reference source not found.**); thus, we do not capture the full spectrum of pricing behavior in the internet retail. However, there are advantages of focusing on this type of sellers. For example, sellers with online and offline presence (e.g., Walmart) have to coordinate their online and offline prices to ensure that consumers do not exploit pricing differentials across the retail modes. Because offline prices are rather sticky, they can delay adjustment of online prices. In contrast, online-only stores do not face such a drag and can react to shocks and competitors' prices faster. Thus, an emphasis on online-only stores may offer a better environment to test the predictions of the law of one price in a friction-free setting.

B. Data filters

Because price data are extraordinarily heterogeneous in our sample, we apply a series of filters to minimize the effects of missing values, extreme observations, etc. Specifically, we drop the top and bottom 1 percent of prices within each category-country. For time series analyses focused on dynamic responses, we keep only goods with at least twenty observations. We remove price quotes for used/refurbished goods, which effectively means excluding many "marketplace" sellers, such as eBay. Finally, because we are interested in international price comparisons, we constrain the sample only to goods that were sold by both U.S. and Canadian online sellers.

This last filter may be fairly restrictive: goods sold in multiple countries typically constitute only a small fraction of goods sold locally. For example, Gopinath et al. (2011) use price data from a large grocery chain prominently present in the U.S. and Canada. Given the universe of approximately 120,000 UPCs sold by the chain, they can match only 3.3 percent of UPCs across the U.S.-Canada border (approx. 4,000 goods). Broda and Weinstein (2008) document a similar effect using a much larger universe of UPCs: only 7.5 percent of the goods are available in both the U.S. and Canada. Fortunately, the overlap in our data is high: the match rate is more than 50 percent.

These filters reduce the number of goods in our sample from 115,000 to about 24,000. We verified that selection into the estimation sample is likely to be random as various pricing moments are approximately the same in the full and estimation samples. For example, the distribution of price levels for the estimation sample is close to the distribution for the full sample. Likewise, the key moments are very similar for the full and estimation samples (see Appendix D).

C. Data quality

PCWs are convenient and popular aggregators of price information. A major study by the European Commission (2013) reports that 74 percent of *all* shoppers in the E.U. use internet comparison tools (PCW is the most popular one: 73 percent of comparison tool users) to compare prices (69 percent of users) and find the cheapest price (68 percent of users). Electric/electronic appliances is the product category with the most intensive use of price comparison tools (63 percent of users). 48 percent of users check a PCW before making an online purchase, and 35 percent of users report that the use of a comparison tool results in a purchase. E-commerce merchants use PCWs to attract new customers and increase sales.

PCWs routinely allow automatic export of product feeds so that whenever an online seller changes a good's price, the new price is immediately reflected on PCWs. Online sellers are also interested in keeping their prices as current as possible because they often pay for clicks on PCWs, and if a price is outdated or a good is out of stock, online sellers waste money.[7] However, there could be systematic discrepancies between prices reported on PCWs and prices listed on the websites of sellers because, for example, online sellers may engage in "bait and switch" strategies. To assess the quantitative importance of this concern, we use several approaches.

First, we compare prices from both sources (that is, from the PCW and from a seller listed on the PCW) for a random sample of 100 goods.[8] Specifically, a script clicks on a link for each seller listed on our PCW and collects price information from the seller's webpage (if necessary, this information is checked manually). We find (Figure 4) that while there are some discrepancies, price quotes (Panel A) are remarkably consistent across sources.

[7] For example, our price comparison website charges between \$0.35 and \$1.15 per click depending on the product category (the website does not charge per listing during the sample period).

[8] We are extremely grateful to Alberto Cavallo for generating price data from websites of online sellers and sharing these data with us.

When we aggregate price quotes across sellers and focus on the average price for a given good (Panel B), the difference between the sources is small. The differences are somewhat larger when we consider dispersion of prices across sellers measured in terms of standard deviation (Panel D) and interquartile range (Panel C). However, even for price dispersion, the PCW provides quite accurate information. If we regress a moment based on prices from sellers' websites on the corresponding moment based on prices from the PCW, we get an estimated slope close to one and an estimated intercept close to zero with R^2 approaching to one. We cannot reject equality of moments across the sources of price information. In a similar spirit, when we compare price quotes for Apple products listed on our PCW and on Apple store website (price quotes for the latter are provided by Cavallo, Neiman, and Rigobon (2014)), we find a high correlation (ρ =0.98) of price quotes across sources (see Appendix E).

Second, we compare the dynamics of prices in our data with the dynamics of prices collected by the Bureau of Labor Statistics (BLS). Specifically, we restrict our sample to product categories that can be matched to disaggregated price indices constructed by the BLS. For example, we can compare the dynamics of "RA01 Televisions" price index constructed by the BLS with the dynamics of an equally weighted price index based on PCW quotes in the Plasma/LCD TV category. Using six matches to the BLS data, we find that the dynamics of prices in our data and the BLS data are similar (see Appendix D for more details).

Third, one may be concerned that PCWs may post outdated price quotes. While it is difficult to establish the lag in price updates, we use a natural experiment to assess the quantitative importance of this potential problem. Specifically, in Appendix A, we explore how price quotes on our PCW responded to the 2011 Thailand floods that had a major impact on the global production of hard drives. We document that prices for hard drives reacted within a week with the peak response within a month. We also observe the significant exit of sellers from the PCW, which is consistent with depleted inventories. These results suggest that price quotes are updated quickly, which is consistent with the assessment in European Commission (2013). We conclude that the quality of price data from the PCW is reasonably high.

II. Basic facts about price setting in online markets

Panels A and B of Table 3 show descriptive statistics for our data.[9] Let i, t, s, c index goods, time (weeks), sellers, and countries, respectively. The average log price $\log P_{itsc}$ in our sample is 5 (or approx. \$150). This magnitude is significantly larger than the level of prices considered in previous studies (e.g., with scanner price data or online prices of books and CDs), where goods routinely have prices below \$10. It is also not unusual in our sample to observe prices of goods above \$600 (approx. 75th percentile) or \$1400 dollars (approx. 90th percentile). Since we focus on how quickly cross-border arbitrage opportunities dissipate, the level of prices is important as search effort is likely to be larger for big-price-tag items. The level of prices is approximately the same in the U.S. and Canada.

Goods routinely have multiple sellers in our data. The average number of sellers is approximately 2.4 in Canada and 3.4 in the U.S. This is consistent with the notion that the U.S. market is larger than the Canadian market, but the difference is not as striking as one observes in the numbers of regular, brick-and-mortar stores in two countries. In part, this difference is smaller because online markets tend to be more concentrated. The stability of sellers—we define stability as the ratio of the number of stores selling a good in a given week to the number of stores ever selling this good in the month which covers the given week—is similar in Canada (0.90) than in the U.S. (0.89).

Similar to previous studies of online prices (e.g., Brynjolfsson and Smith 2000, Baye et al. 2006), we observe dramatic cross-sectional dispersion of prices which is calculated as

$$\sigma_{itc} \equiv \left\{\frac{1}{\#(\mathcal{S}_{itc})}\sum_{s\in\mathcal{S}_{itc}}\left(\log P_{itsc} - \frac{1}{\#(\mathcal{S}_{itc})}\sum_{s\in\mathcal{S}_{itc}}\log P_{itsc}\right)^2\right\}^{0.5},$$

[9] We present selected statistics by category of goods in Appendix G.

where $\mathcal{S}_{itc}$ is the set of stores that sell good *i* in week *t* in country *c*. On average, across goods and time periods, the standard deviation of log prices within a country is 0.13-0.16, which is significant but smaller than one can observe for the dispersion of prices across regular stores.[10,11] Given that the levels of prices are large in our sample, these price differentials correspond to significant dollar amounts. In some cases, the differences between cheapest and most expensive prices are in multiple hundreds of dollars, which could be surprising given easy search for the best prices in online markets. However, we do observe that the size of price differentials is negatively correlated with the level of prices. That is, more expensive goods tend to have smaller (log) price dispersion. We also find that the cross-sectional dispersion of prices in any given market is fairly persistent. The serial correlation of the log or level of σ_{itc} is routinely above 0.85.

The frequency of price changes is high: 20 to 30 percent of prices change in a given week, implying that the average duration of price spells is just a few weeks.[12] Price increases and decreases are equally likely in our data. The average price change is slightly negative, which captures the fact that goods in our sample are subject to technical improvements over time; thus, prices of existing goods tend to depreciate with the age of goods. Temporary price cuts ("sales") are relatively infrequent (approx. 2-3 percent of goods are on sale in a given week) and small (the average size is 5 to 10 percent). In contrast, prices in scanner price data (e.g., Kehoe and Midrigan 2015), in government surveys of prices (e.g., Nakamura and Steinsson 2008), or in online prices for books (e.g., Boivin, Clark, and Vincent 2012) have a much lower frequency of price

[10] For example, Coibion, Gorodnichenko and Hong (2015) report that the standard deviation in the log price for a given unique product code (UPC), a given market (metro area), and a given week is 28 percent on average across periods, markets, and UPCs. Sheremirov (2015) documents similar evidence.

[11] Rating of sellers is a strong predictor of price deviations for a given good; thus, some price dispersion is due to compensating differentials for seller reputation. However, the dispersion remains high even after controlling for store rankings.

[12] We define a price change as a movement in prices larger than one percent in absolute value. We discard very small price changes (less than one percent in absolute value) as these changes are likely to arise from measurement errors (e.g., Eichenbaum et al. 2014).

changes, a larger size of price changes, and more prevalent and deeper sales. At the same time, our moments are consistent with Lünnemann and Wintr (2011), who analyze a similar set of goods but have data only for one year. Higher frequency and smaller sizes of price changes for online prices are consistent with "menu" costs being smaller for online sellers than for regular stores.

As a final measure of price stickiness, we consider synchronization of price changes across sellers. Specifically, we define synchronization in a given week for a given good as the fraction of price quotes with a price change conditional on at least one price change and having at least two sellers at this point in time:

$$Synchronization_{itc} = \frac{\sum_{s\in\mathcal{S}_{itc}} \mathbf{1}\{P_{itsc} \neq P_{i,t-1,sc}\} - 1}{\sum_{s\in\mathcal{S}_{itc}} \mathbf{1}\{P_{itsc} \neq \text{missing } \cap P_{i,t-1,sc} \neq \text{missing}\} - 1},$$

where we code $Synchronization_{itc}$ as missing if $\sum_{s\in\mathcal{S}_{itc}} \mathbf{1}\{P_{itsc} \neq P_{i,t-1,sc}\} < 1$. The average synchronization is 19 percent in the U.S. and 23 percent in Canada. These magnitudes are very similar to the unconditional frequencies of price changes and hence point to little synchronization of price changes across sellers.

While our results point to greater flexibility of online prices, one may be concerned that this outcome is determined by differences in the composition of goods sold online and in regular stores. To address this concern, we compare moments for narrowly defined categories of goods for price data from our PCW, from a major online shopping platform (Gorodnichenko, Sheremirov, and Talavera 2014), and from conventional stores (Nakamura and Steinsson 2008). Consistent with our earlier results, we find (Table 4) that relative to prices in conventional stores, online prices have a higher frequency and smaller size of price changes as well as less frequent and smaller sales. Prices from our PCW have properties (frequency, size, and synchronization of price changes and cross-sectional dispersion of prices) similar to the properties of prices directly provided by a major PCW/shopping platform. Thus, differences in the composition of goods are not a likely explanation for differences in pricing moments in online and offline retail.

III. International price differentials

A. Descriptive statistics

We focus on two popular measures of international price differentials: the relative exchange rate $\log\left(P_{it}^{CA}/P_{it}^{US}\right)$ and the real exchange rate $\log\left({EX_t}^{-1} \times P_{it}^{CA}/P_{it}^{US}\right)$, where *i* and *t* index goods and time, respectively, P_{it}^{CA} (P_{it}^{US}) is a price measure for a given good in Canada (U.S.), and EX is the CAD/USD nominal exchange rate. Since for any given period/good/country there are multiple sellers and hence multiple prices, we consider several measures of prices at the country level: mean price across sellers; median price across sellers; and minimum price across sellers.[13] Each of these measures has pros and cons. For example, while the mean price was often used in previous research, median prices are less sensitive to extreme price quotes. In light of Baye et al. (2009), Chevalier and Kashyap (2011), and Gorodnichenko, Sheremirov, and Talavera (2014), one may conjecture that minimum prices are closer to transaction prices and thus are more likely to capture prices relevant for consumers.

Irrespective of which measure of prices we use, international price differentials are moderately large (Panel C, Table 3). The mean of $\log\left(P_{it}^{CA}/P_{it}^{US}\right)$ and $\log\left({EX_t}^{-1}P_{it}^{CA}/P_{it}^{US}\right)$ is about 5 to 12 percent. Some of the price dispersion across countries can be explained by differences in taxes. For example, the value added tax (federal and provincial) in Canada is about 13 percent, and there is big variation in sales taxes across U.S. states.[14] However, differences in taxes are unlikely to be the whole story. First, there is dramatic variation in price differentials (columns (4) and (5) in Table 2): the 25th percentile of the mean

[13] We also considered mean price weighted by the reputation of sellers, where reputation is measured as the number of stars, from 1 to 5, that consumers assign to sellers. Results for star-weighted moments are similar to the results reported in the paper. We also constrained our sample to include sellers with 4+ star reviews. We found similar results.

[14] Although we use an address in Berkeley, CA, online sellers with no physical presence in California do not have to collect sales tax (close to 10 percent) on behalf of the state of California. As a result, Californian consumers often pay no sales tax on their online purchases.

price differential is negative, while the 75th percentile is between 15 and 25 percent. The AR(1) coefficient for either exchange rate is between 0.80 and 0.92 (at weekly frequency), depending on whether we control for good/type fixed effects so that the implied half-life is 3 to 6 weeks, which is much shorter than half-lives estimated on prices collected in regular stores. If price differentials were mainly determined by taxes, one would expect to see little if any variation in price differentials across goods or over time. Second, for a subsample of goods that we have information for gross prices that include taxes and shipping costs, we observe similar international price differentials (Appendix Table F1).[15]

The standard deviation of price differentials across countries—which ranges from 0.22 to 0.27 see column (2)—is much larger than the standard deviation of price differentials within countries, which is between 0.09 and 0.11. This finding is qualitatively consistent with results reported in the earlier literature comparing price differentials within and across countries (e.g., Engel and Rogers 1996, Gorodnichenko and Tesar 2009). However, moments for the real and relative exchange rates are broadly similar so that fluctuations in the nominal exchange rate are unlikely to be the main factor in cross-border price differentials.

In summary, properties of online price differentials are qualitatively similar to properties of prices in regular markets, but the magnitude and persistence of price differentials are smaller relative to counterparts reported in previous studies for brick-and-mortar stores. Thus, this first pass at the data suggests that frictions are much smaller in online markets, but non-negligible cross-sectional dispersion of prices and some persistence of price differentials are consistent with some border frictions in online markets. In the following

[15] The price comparison web page was redesigned for various goods in various times, and in many versions of the webpages, we could specify the location of the buyer and thus obtain net and gross prices. We used the address of the Department of Economics at UC Berkeley as the shipping destination. Gross prices are available for about half of quotes for which we have net prices.

sections, we will examine predictors of these persistent and volatile cross-border price differentials in online markets.

B. Pass-through and the speed of price adjustment

To characterize the dynamics of cross-border price differentials, economists commonly use two metrics: pass-through (i.e., how movements in the nominal exchange rate are translated into movements of prices of goods) and the speed of price adjustment to equilibrium levels. While there is a variety of versions of these two metrics, we employ two basic econometric specifications to construct these metrics:

Pass-through α:

(1) $$\log\left(\frac{P_{it}^{CA}}{P_{it}^{US}}\right) = \alpha EX_t + Controls + error_{it},$$

Speed of price adjustment β:

(2) $$d\log\left(\frac{P_{it}^{CA}}{P_{it}^{US}}\right) = \beta\left(\log\left(\frac{P_{i,t-1}^{CA}}{P_{i,t-1}^{US}}\right) - \alpha EX_{t-1}\right) + \phi_1 d\log\left(\frac{P_{i,t-1}^{CA}}{P_{i,t-1}^{US}}\right) + \lambda_1 dEX_{t-1} + \boldsymbol{Controls} + error_{it} \ ,$$

where $\boldsymbol{Controls}$ is a set of control variables, and $dx_t \equiv x_t - x_{t-1}$ is the first difference operator.[16] Specification (1) estimates the long-run pass-through and is a generic specification estimated in the literature (see Goldberg and Knetter (1997) for a survey). The law of one price predicts that α should be equal to one and, hence, values of α closer to one correspond to smaller departures from the law of one price. Specification (2) is set in the error-correction/cointegration form where β quantifies how quickly the deviation from equilibrium is eliminated. More negative values of β mean faster adjustment. In specification (2), equilibrium relationship between relative and the exchange rate (coefficient α) is determined according to specification (1). Thus, while the equilibrium relationship nests the law of one price, it also allows deviations from the law of

[16] We use BIC to select the number of lags for $d\log\left(P_{i,t-1}^{CA}/P_{i,t-1}^{US}\right)$ and dEX_{t-1}.

one price (i.e., α can be less than one).[17] In our preferred specification, $Controls$ include good fixed effects.

A key assumption behind specifications (1) and (2) is that price differentials have a common stochastic trend, which is captured by the nominal exchange rate. Because the error term is almost certainly correlated across goods, and hence standard panel-data unit root tests are not suitable, we use the Bai and Ng (2004) approach to extract a common component from price differentials and test it for a unit root and for cointegration with the nominal exchange rate. The results of these tests (Appendix B) indicate that there is indeed a common stochastic trend cointegrated with the nominal exchange rate. Hence, specifications (1) and (2) are valid.

Table 4 reports estimated specifications (1) and (2) on pooled data. To account for the fact that error terms in specifications (1) and (2) can be correlated across time, goods, and countries as well as the fact that EX_t is common across goods and countries, we use the Driscoll and Kraay (1998) standard errors. Note that for specification (2) we have fewer observations because we restrict the sample only to goods with at least twenty observations.

The estimated exchange rate pass-through (Panel A) is about 60 to 75 percent, which is considerably larger than 20 to 40 percent pass-through typically reported in previous studies based on prices collected from regular stores (Menon 1996, Kardasz and Stollery 2001, Goldberg and Verboven 2001, Barhoumi 2005, Campa and Goldberg 2005, Gaulier, Lahreche-Revil, and Mejean 2006, Takhtamanova 2010, Gopinath and Rigobon 2008, Cao, Dong, and Tomlin 2012). This increased pass-through is consistent with salient features of online markets: i) prices are more flexible, ii) competition is fierce,

[17] Since we use an estimated α in equation (2), one may be concerned about the consistency of estimated β as well as using standard inference for estimated β. These concerns are unlikely to be quantitatively important for several reasons. First, exchange rates are fairly persistent and approach a unit root so that an estimate of α in specification (1) can be super-consistent. Second, the error terms in specifications (1) and (2) are essentially uncorrelated; thus, adjustment for the generated regressors is minimal. Hence, we can first estimate specification (1) and then use $\hat{\alpha}$ to construct the deviation from equilibrium relationship in specification (2).

iii) consumers can easily buys goods from the U.S. or Canada, iv) distribution/non-tradable costs are small, and v) most goods are produced overseas so that the costs are similar across countries.

Estimated β's (Panel B) suggest a fast correction of prices toward a long-run equilibrium. If we abstract from the short-run dynamics (i.e., ϕ and λ) in specification (2), 7 percent or more of the gap from the long-run relationship is closed in a week (correspondingly about 25 percent of the gap is closed in a month and 60 percent in a quarter), which implies the half-life of 2-2.5 months or less. This speed of adjustment is considerably faster than the speed estimated on price indexes (e.g., Rogoff (1996) estimates a half-life of 3 to 5 years) or scanner price data, where prices of exact same goods sold in regular stores are compared across countries (e.g., Broda and Weinstein (2008) estimate a half-life of 2.9 quarters). This speed of price adjustment, however, would probably not surprise observers of the online markets. For example, Baye et al. (2007) emphasize that i) online customers compare prices within goods, not within stores; ii) the number of sellers and prices changes frequently; and iii) firms need to constantly monitor prices of their rivals. All of these factors are likely to accelerate price adjustment.

One may be concerned that high pass-through and the speed of price adjustment are potentially determined by idiosyncratic, transitory shocks such as sales and measurement errors in our data. To address this concern, we perform several checks. First, we run a series of calibrated Monte Carlo experiments to show that it would take implausibly large measurement errors to drive our results (see Appendix C). Second, we aggregate data to monthly frequency to reduce the importance of transitory shocks in the data. Pass-through and the speed of price adjustment estimated at a monthly frequency (**Error! Reference source not found.**) are similar to the estimates at a weekly frequency. Third, we re-estimate specificions (1) and (2) on regular prices (i.e., excluding sales) and find similar results (**Error! Reference source not found.**).[18] One should also note that we use

[18] We use ∧- and ∨-shaped filters as in Nakamura and Steinsson (2008) to identify sales.

prices averaged across sellers so that adverse effects of idiosyncratic shocks on estimated pass-through and the speed of price adjustment are likely attenuated. Thus, we conclude that idiosyncratic, transitory shocks are unlikely to drive our estimates.

The speed of adjustment in our data is much higher than the speed estimated by Boivin, Clark, and Vincent (2012) for online prices of books or by Cavallo, Neiman, and Rigobon (2014) for online prices of Apple products. The discrepancy in the results for books is likely to reflect the specifics of book markets, which tend to have much stickier prices and higher market power of sellers. While Apple goods are seemingly similar to goods in our sample, there are important differences. Most importantly, Apple has considerable market power and can limit price competition across sellers and its own Apple store. As a result, Apple products have stickier prices, fewer and smaller sales, lower cross-sectional price dispersion as well as lower pass-through and slower speed of price adjustment (see Appendix E). More generally, one may expect that sellers present in both online markets (e.g., Amazon.com and Amazon.ca) can price discriminate consumers in Canada and the U.S. and reduce competition between their branches in different countries. This behavior should reduce pass-through and the speed of price adjustment. Results in Panel C of Table 5 are consistent with this intuition and may explain why previous studies (e.g., Gopinath et al. 2011, Cavallo, Neiman, and Rigobon 2014) using price comparisons across branches of the same seller in different countries tend to find low pass-through.

C. Predictors of pass-through and the speed of price adjustment

While in the previous section we focus on pooled estimates of pass-through and the speed of price adjustment to present simple summary statistics, there is dramatic heterogeneity of these characteristics across goods (Table 5) when we estimate α and β at the level of individual goods. A key question is as follows: what factors are systematically related to the size of pass-through and the speed of price adjustment? Usually, it is hard to answer this question because the data are available only at the aggregate level or little is known about the properties of goods and, as a result, previous research (e.g., Yang 1997, Campa and Goldberg 2005) focused on macroeconomic determinants (e.g., exchange rate regime, level of inflation) of pass-through. Fortunately, our dataset contains information about a number of potentially important determinants at the micro level.

To be clear, we have observational data, and, therefore, our results should not be interpreted as causal; they document correlations. However, these correlations are informative about equilibrium relationships in the data, and, therefore, they provide important inputs for theoretical efforts aimed at rationalizing the behavior of international price differentials. In what follows, we discuss several groups of factors that are arguably related to the behavior of international price differentials and then explore if estimated correlations are consistent with theoretical predictions.

First, Head, Kumar, and Lapham(2010), Richards, Gómez, and Lee (2014), and others argue that the degree of pass-through is negatively related to search costs. The return to search effort should be higher for expensive goods. For example, consumers are more likely to search for better deals on computers and plasma TVs than on toothpaste or beer. A higher search intensity should put a larger pressure on price convergence across sellers and countries. Thus, one may expect that more expensive goods should exhibit a larger pass-through and faster speed of price adjustment. Our dataset has a wide distribution of goods in terms of their prices, and we can exploit this variation to examine and quantify this channel. Specifically, we use log median prices to proxy for returns on search.

Second, a number of studies (Rogoff 1996, Apslund and Friberg 2001, Bergin and Feenstra 2001, Imbs et al. 2005, Mayoral and Gadea 2011, Devereux and Yetman 2010, Takhtamanova 2010) suggest that price stickiness can be an important force in determining how deviations from the law of one price are eliminated. With flexible prices, adjustment can be deep and quick. In contrast, sticky prices can delay price adjustment and make it incomplete. We can measure the degree of price stickiness using the mean frequency of price changes for a given good in our sample. More frequent price changes should be associated with larger pass-through and faster price adjustment. In addition, we use prevalence of convenient prices (e.g., prices like $199, $99, $39.99) and frequency of sales to capture price rigidity more completely. Intuitively, convenient prices create barriers to price adjustment because pricing points ending in, e.g., 9, tend to be far apart; hence, firms may choose to stick to a convenient price even in spite of relatively large shocks. Knotek (2011) documents that high incidence of convenient price is indeed associated with increased price rigidity. On the other hand, sales may be interpreted as a form of price flexibility used by a firm to respond to shocks when the firm cannot change its regular price (Kehoe and Midrigan 2015).

Third, the degree of synchronization in price changes can be important because pass-through and the speed of price adjustment could be affected not only by the degree of price stickiness at the level of individual sellers but also to what extent price setting is staggered (see Neiman 2010). Indeed, in many macroeconomic models, one needs staggered price setting in addition to strategic complementarity to generate gradual adjustment of prices. As argued by Bhaskar (2002) and others, if prices are set simultaneously (i.e., synchronization is high), the reaction of prices to shocks is stronger, and departures from equilibrium levels are quickly eliminated.

Fourth, Feenstra, Gagnon, and Knetter (1996), Atkeson and Burstein (2008), and many others emphasize that market power can affect the magnitude of pass-through. While the theory often stresses market share, we do not have information on sales of individual stores, and we will instead use a proxy for the degree of market power. Specifically, the number of sellers should be indicative of the degree of competition. With more sellers, one should expect a larger pass-through and speed of adjustment.

Fifth, Gust, Leduc, and Vigfusson (2010) argue that firm entry can increase exchange rate pass-through. Indeed, an easier entry into selling a good is likely to make competition stronger (e.g., hit-and-run strategy) and, as a result, make pass-through larger and price adjustment faster. A stronger turnover of sellers is likely to be indicative of how easy it is to start selling a given good. We proxy for the turnover using our stability measure (i.e., a more stable set of sellers means a lower turnover), and, hence, we should expect a negative correlation between stability and pass-through and between stability and the speed of price adjustment.

Finally, reputation of sellers can influence pass-through and speed of price adjustment. Specifically, consumers are more likely to take advantage of price differentials if sellers of a given good have a high reputation because price differentials then likely present a genuine opportunity to have a good deal rather than capture a compensating differential for lack of reputation or heterogeneity in some other dimension (see Imbs et al. 2010 for a discussion). This logic suggests that pass-through and speed should be high if sellers have a high reputation.

To test these predictions, we estimate specifications (1) and (2) for each good separately and then regress estimated $\hat{\alpha}$ and $\hat{\beta}$ on the factors we describe above:

(3)
$$\begin{aligned} Outcome_i = {} & \gamma_1 \log(\bar{P}_i) + \gamma_2[\log(\bar{P}_i)\,]^2 + \gamma_3 Frequency_i + \\ & \gamma_4 \log(Sellers_i) + \gamma_5[\log(Sellers_i)]^2 \\ & +\gamma_6 StabilitySellers_i + \gamma_7 Synchronization_i \\ & +\gamma_8 Reputation_i + \gamma_9 Sales_i + \gamma_{10} Convenient_i + T_i + C_i + error_i, \end{aligned}$$

where *i* indexes goods, $Outcome_i = \{\hat{\alpha}_i, \hat{\beta}_i\}$, $\bar{P}_i$ is the median price of good *i* in the U.S., $Frequency_i$ is the average frequency of price changes in Canada and the U.S., $Sellers_i$ is the number of sellers in the U.S. and Canada, $StabilitySellers_i$ is the average stability of sellers in the U.S. and Canada, $Synchronization_i$ is the average synchronization rate of price changes in the U.S. and Canada, $Reputation_i$ is the average star rating of U.S. and Canadian sellers, $Sales_i$ is the average frequency of sales in the U.S. and Canada, $Convenient_i$ is the average share of convenient prices in the U.S. and Canada,[19]

[19] We define convenient prices as prices that end with 9 in the $1-$100 range (e.g., $39, $59.99, $79.50) or that end with 99, 98, 97, 96, or 95 in the $100+ range (e.g, $199, $399.99,

T_i is a set of fixed effects for periods over which $\hat{\alpha}_i$ and $\hat{\beta}_i$ are estimated, and C_i is a set of fixed effects for categories of goods. Each variable on the right-hand side is calculated as the time series average. Table 6 reports estimated coefficients for specification (3) by least squares for various measures of prices.

We have conjectured a positive relationship between the size of pass-through and returns on search proxied by the price of a good. The estimates suggest a non-linear relationship. For goods with prices less than approximately $150—which is close to the median price of goods in our data—the relationship is positive, but it turns into a negative one for more expensive goods. This inverted-U relationship suggests that pass-through and search have an interplay that is more complex than often assumed. Indeed, pass-through and search are determined simultaneously in equilibrium, and firms can respond to endogenous search effort by pricing goods in such a way that returns to search are reduced for expensive goods where search is likely to be most intensive and hence the elasticity of demand can be particularly high. For example, a manufacturer can require online stores to sell its good at a price set by the manufacturer rather than by retailers, thus limiting price dispersion and competition between stores. In addition, manufactures could be more likely to sell high-price goods (e.g., laptops) directly to customers than low-price goods (e.g., cables), and they may be interested in preserving sales through their websites again by limiting price dispersion. While we are not able to test hypotheses of this type with our data, there is anecdotal evidence consistent with this explanation.[20]

Interestingly, we also find an inverted-U relationship between a good's price and the speed of price adjustment, where the speed is the slowest for goods priced around $150, which is approximately the price where the estimated pass-through is the highest. Note that $\hat{\alpha}_i$ and $\hat{\beta}_i$ are essentially uncorrelated, and,

$999.50). Note that in defining convenient prices, we ignore cents and focus only on dollar amounts. As a result, prices like $30.99 are not considered convenient.

[20] For example, Apple products sold in a broad array of online stores show little, if any, price dispersion across online stores because Apple apparently coordinates prices across sellers (see an article on zdnet.net).

therefore, it is unlikely that this pattern arises mechanically from the way we estimate these parameters. It is more likely that this pattern reflects incentives to adjust prices. Intuitively, if pass-through is close to 100 percent, returns to arbitrage are second-order as the profit function is approximately flat. As a result, the speed of price adjustment is slow. In contrast, when pass-through is low, returns to arbitrage are high (the slope of the profit function is steep), and, thus, the speed is fast.

There is also a non-linear relationship between the number of sellers and pricing dynamics. Specifically, raising the number of sellers from two sellers (the minimum number) to 4-5 sellers (approximately, the average number of sellers) is associated with increased pass-through. Further increases in the number of sellers are associated with decreasing pass-through. The speed of price adjustment is not significantly correlated with the number of sellers.

There is a strong positive relationship between the size of the estimated pass-through and frequency of price changes. Specifically, a one standard deviation increase in the frequency of price changes (approx. 0.17) is associated with a 34 percentage point increase in pass-through. High frequency of price changes is also strongly associated with faster price adjustment. Estimates for other proxies of price stickiness (prevalence of convenient prices) and price flexibility (frequency of sales) paint a similar picture. Overall, consistent with theoretical predictions, goods with stickier prices have a lower speed of price adjustment.

Greater synchronization of price changes is associated with lower pass-through. At the same time, we find weak evidence that synchronization decelerates price adjustment. These results suggest that synchronization likely captures market power, enabling coordination of price changes and limiting the ability of online sellers to eliminate arbitrage opportunities.

The stability of sellers is significantly negatively correlated with the speed of price adjustment: a lower turnover of sellers (higher stability) reduces the speed (i.e., $\hat{\beta}$ becomes larger and closer to zero). This finding is consistent with the view that easy entry into a market and limited time-horizons for sellers, which limits the scope for collusion, are likely to eliminate arbitrage opportunities and mis-pricing of goods faster. The quantitative effect of seller stability is large. A one standard deviation increase in stability (approximately 0.05) is associated with a 0.05 reduction in the speed. At the same time, we do not find a significant relationship between pass-through and stability.

In summary, although we cannot establish causal links in our data, estimated correlations shed useful light on the relative roles of potential forces that determine pass-through and the speed of price adjustment. Future work that makes identifying assumptions (i.e., structural approach) or employs (quasi-) experimental design may quantify causal chains in the data. Our results summarizing patterns in the data supply moments to be matched in this future work.

D. Margins of price adjustment

While the previous section documents that pass-through and the speed of price adjustment are high in online markets, one can learn more about these two objects by exploring what margins of price adjustment are used in response to movements in the nominal exchange rate. We use our specification (1) to construct a measure of the deviation from equilibrium EC:

$$(4) \qquad \widehat{EC}_{it} = \log\left(\frac{P_{it}^{CA}}{P_{it}^{US}}\right) - \hat{\alpha} EX_t.$$

where, as before, i and t index goods and time (weeks), respectively, P is a measure of a price (e.g., median price, mean price, minimum price), and EX is the nominal exchange rate. Note that α is estimated for each price measure separately.

We measure the intensive margin of price adjustment as the average price change (conditional on price change) across sellers of good i in country c and week t:

(5) $$\overline{dP}_{ict} = \frac{\sum_{s=1}^{S_{itc}} \log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) \times \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| > 0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| > 0.01\right\}}.$$

We also calculate the mean size of price increases and price decreases separately:

(5') $$\overline{dP}_{ict}^{decrease} = \frac{\sum_{s=1}^{S_{itc}} \log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) \times \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) < -0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) < -0.01\right\}},$$

(5'') $$\overline{dP}_{ict}^{increase} = \frac{\sum_{s=1}^{S_{itc}} \log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) \times \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) > 0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) > 0.01\right\}}.$$

The extensive margin of price adjustment—again with the distinction for any price change, price increase, and price decreases—is measured as

(6) $$\Pr(dP \neq 0)_{ict} = \frac{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| > 0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| \text{ is not missing}\right\}}$$

(6') $$\Pr(dP > 0)_{ict} = \frac{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) > 0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| \text{ is not missing}\right\}}$$

(6'') $$\Pr(dP < 0)_{ict} = \frac{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right) < -0.01\right\}}{\sum_{s=1}^{S_{itc}} \mathbf{1}\left\{\left|\log\left(\frac{P_{isct}}{P_{isc,t-1}}\right)\right| \text{ is not missing}\right\}}$$

and is thus a fraction of sellers that change their prices in the set of sellers that have listed good i in weeks t and $t-1$.

Finally, stores with the best prices may run out of inventories faster than other stores; thus, cheap stores can be more likely to exit the market until they replenish their inventories. We calculate the probability of exit as follows:

(7) $$\Pr(exit)_{ict} = \frac{\sum_{s=1}^{S_{ic,t-1}} \mathbf{1}\{P_{isct} \text{ is missing} \cap P_{isc,t-1} \text{ is not missing}\}}{\sum_{s=1}^{S_{ic,t-1}} \mathbf{1}\{P_{isc,t-1} \text{ is not missing}\}}.$$

Using these measures, we estimate the following generic specification with a pricing moment given in (5)-(7) as the dependent variable:

(8) $$Moment_{ict} = \gamma_c + \psi_c \widehat{EC}_{i,t-1} + \kappa_{c1} EX_{t-1} + \kappa_{c2} Moment_{ic,t-1} + \lambda_{ic} + error_{ict}.$$

Note that specification (8) is estimated for each country separately as the direction of the change in the pricing moment can depend on whether equilibrium error EC is positive or negative; thus, estimated coefficients may move in opposite directions for Canada and the U.S. For example, if $EC > 0$ (goods in Canada are relatively expensive), one may expect prices in Canada to decrease (i.e., $\overline{dP}_{i,CA,t} < 0$) and prices in the U.S. to increase (i.e., $\overline{dP}_{i,US,t} > 0$) and hence $\psi_{CA} < 0$ and $\psi_{US} > 0$.

Table 7 presents estimates of ψ_c, which is the key parameter in specification (8), for various pricing moments and measures of prices. For the response of the mean price change $\overline{dP}_{ict}$, we consistently find (row 1) that if prices in Canada are 10 percentage points above equilibrium level, prices in Canada fall by 0.8 to 1.3 percentage points on impact, while prices in the U.S. increase by 0.4 to 0.7 percentage point on impact. Consistent with our previous findings, these results suggest fast adjustment of prices to equilibrium levels. This pattern also applies to both price increases (row 2) and price decreases (row 3). For example, if we focus on the mean prices in the U.S. and Canada, a positive equilibrium error EC (i.e., prices are more expensive in Canada), price increases in Canada become smaller, while price decreases become larger (more negative). Likewise, a positive equilibrium error EC tends to lead to larger price increases and smaller (i.e., less negative) prices decreases in the U.S. Hence, we do not observe strong asymmetric effects in the size of price adjustment as prices appear to be equally flexible in terms of increases and decreases. The magnitude of the response is generally larger for Canada than for the U.S., which is consistent with the view that price adjustment is likely to be larger in smaller markets.

The frequency of price adjustment for all price changes (row 4) does not exhibit a robust relationship to equilibrium errors. However, this lack of correlation reflects that movements in frequencies of price increases and frequencies of price decreases roughly offset each other. Once we focus on the frequency of price increases (row 5) and the frequency of price decreases (row 6) separately, the data indicates a strong link between the frequencies of price adjustment and equilibrium errors. Consider the frequency of price increases when we use mean prices. A positive 10 percentage point equilibrium error EC reduces the frequency of prices increases in Canada by 0.85 percentage points and increases the frequency of price increase in the U.S. by 0.29 percentage points. This finding is in line with the price adjustments along the intensive margin where positive EC leads to smaller price increases in Canada and larger price increases in the U.S. The effect is in the opposite direction for the frequency of price decreases: a positive 10 percentage point equilibrium error EC increases the frequency of prices decreases in Canada by 0.76 percentage points and decreases the frequency of price decrease in the U.S. by 0.20 percentage points. One can immediately see that the movements of the frequency of price increases and the frequency of price decreases have similar magnitudes, and thus the effect on the frequency of all price changes becomes weak. Similar to the results for the intensive margin, the response of the extensive margin is stronger for Canada than for the U.S.

The exit of goods with cheap prices is not strongly correlated with equilibrium errors. We only find one case with minimum prices with significant statistical evidence that a positive equilibrium error makes exit of stores less likely in Canada and more likely in the U.S. While one should expect this pattern, we conjecture that we do not find the same patterns for other price measures because the consumer pressure arising from price differentials is likely to be the highest for stores offering lowest prices. Indeed, price sensitive consumers are likely to buy at the cheapest prices and thus are more likely to respond to arbitrage opportunities when relative prices shift. At the same time, given fairly large dispersion of prices within countries, mean or median prices at the level of countries may be too coarse to detect changes in demand arising from shifts in relative prices.

To further explore margins of price adjustment, Figure 5 plots the time series of mean price changes (i.e., all price changes, price increases, and price decreases in Panels A, B, and C) when we aggregate across goods (with equal weights) to the country level. We also report the estimated slope from regressing each series on the nominal exchange rate. In general, price increases (decreases) in Canada are negatively (positively) correlated with the nominal exchange rate (CAD/USD), and the pattern of correlations is reversed for the U.S. One can also observe that the correlation between the size of price decreases in the U.S. and in Canada is negative.

In a similar manner, we aggregate frequencies of price adjustment across goods to the country level (Panels D, E, and F). These aggregate frequencies for the U.S. and especially for Canada tend to be positively correlated with the nominal exchange rate. However, a decomposition of price changes into price increases (Panel E) and price decreases (Panel F) suggests that the correlation with the nominal exchange rate is the strongest for price increases in Canada and equally weak for price increases and price decreases in the U.S.

The frequency of price increases and decreases in Canada was the highest in late 2008 and early 2009 when the Canadian dollar was strongly appreciating. The fact that the frequency of price changes rose for both price increases and price decreases highlights that the exchange rate movements induced firms to review their prices with possible adjustment in either direction rather than move all Canadian prices in one direction. In other words, firms appeared to be re-optimizing their prices rather than mechanically adjusting their prices by changes in the exchange rate. Obviously, these price adjustments happened during the Great Recession, so perhaps this "churning" of price changes reflects increased intensity of price adjustment in recessions rather than responsiveness of prices to exchange rate fluctuations. However, we observe only a moderate to weak increase in the frequency of price adjustment for U.S. retailers; therefore, it is hard to see the contribution of the Great Recession to increased frequency of price adjustment in Canada.

To explore this issue further, we regress the frequency of price increases and the frequency of price decreases on the CAD/USD exchange rate over the period that excludes the Great Recession; that is, we use data after June 2009. We find that the frequency of price decreases in Canada is not statistically or economically sensitive to the exchange rate, while the frequency of price increases continues to stay highly significant in statistical and economic terms. At the same time, the frequency of price decreases in the U.S. is positively related to the CAD/USD exchange rate (although the sensitivity is smaller than that for Canada), while the frequency of price increases in the U.S. does not exhibit a significant correlation with the exchange rate. This pattern of responses is consistent with the predictions of economic theory on how firms should adjust their prices, and it therefore corroborates our findings in Table 7.

The exit frequency (Figure 6) is positively correlated with the nominal exchange rate for both the U.S. and Canada, but, similar to other margins, the exit margin in Canada is more sensitive to fluctuations in the nominal exchange rate. Some of the positive correlation is determined by the coincidence of high turnover of sellers and goods (i.e., high exit frequency) and depreciation of the Canadian dollar in the Great Recession. If we exclude the Great Recession, the exit frequency in the U.S. shows no sensitivity to the exchange rate, while the exit frequency in Canada is even more strongly positively related to the CAD/USD exchange rate. It appears that when the Canadian dollar depreciates, the U.S. consumers take advantage of cheap Canadian prices and deplete inventories of Canadian stores, while the pool of Canadian customers is unable to exercise the same pressure on U.S. stores when the Canadian dollar appreciates.

IV. Concluding remarks

While the law of one price is an appealing concept, the vast majority of previous research has emphasized various frictions that prevent the law from holding over relative long periods. These frictions can take a variety of forms, but the most popular barriers leading to violations of the law are search costs, costs of nominal price adjustment, and transportation/distribution costs. Assessing the contribution of these frictions has been remarkably difficult as these frictions are ubiquitous in standard markets with brick-and-mortar stores.

Online markets have unusual characteristics, such as low search costs, irrelevance of physical locations of buyers and sellers, and negligible physical costs of price changes; thus, studying price setting in online markets offers a unique opportunity to rule out the prominent frictions and explore whether the law of one price holds in this close-to-ideal setting.

We construct a new, massive dataset of online price quotes in the U.S. and Canada. This dataset has a number of desirable features, such as long time series, large cross sections, and multiple sellers. We document that, relative to prices in regular stores, prices in online markets are more flexible as well as exhibit stronger pass-through and faster convergence in response to movements of the nominal exchange rate. Multiple margins of adjustment (frequency of price changes, direction of price changes, size of price changes, exit of sellers) are active in the process of responding to nominal exchange rate shocks. Furthermore, we use the richness of our dataset to show that the sensitivity of prices to changes in the nominal exchange rate is systematically correlated with the characteristics of goods and markets (e.g., the degree of competition). To the extent future retail will shift to the internet, one can therefore expect that cross-country price differentials are going to be smaller and less persistent, bringing the law of one price closer to reality.

Scraping online prices is a cheap and fast approach to collecting price quotes at high frequencies; therefore, it is attractive to statistical agencies. While these data open new, unprecedented research opportunities (e.g., the Billion Prices Project run by Alberto Cavallo and Roberto Rigobon), economists should also appreciate limitations of many currently available datasets, including the dataset used in this paper. Perhaps the most important one is the lack of information about volumes of purchases associated with price quotes. Using the number of clicks may provide a simple proxy for quantities of goods sold in online stores, but the quality of this and similar proxies should be verified with alternative information. As information technology progresses and internet retailers become more willing to share transaction data, one may expect major improvements in the quality of data so that one can answer questions that seem currently insurmountable. For example, these new data can help us to understand how stores selling goods online and offline (e.g., Walmart) set prices and conduct sales in these interconnected markets. One may also be able to trace consumers' history of searches to transactions and, hence, have a better understanding of how searching operates and how it is related to price dispersion and adjustment.

References

Asplund, Marcus and Richard Friberg. 2001. "The Law of One Price in Scandinavian Duty-Free Stores." *American Economic Review* 91(4): 1072-1083.

Atkeson, Andrew and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review* 98(5): 1998-2031.

Bachis, Enrico and Claudio Piga. 2011. "Low-cost airlines and online price dispersion." *International Journal of Industrial Organization* 29(6): 655-667.

Bai, Jushan and Serena Ng. 2004. "A PANIC Attack on Unit Roots and Cointegration." *Econometrica* 72(4): 1127-1177.

Barhoumi, Karim. 2005. "Long Run Exchange Rate Pass-Through Into Import Prices In Developing Countries: An Homogeneous or Heterogeneous Phenomenon?" *Economics Bulletin* 6(14): 1-12.

Baye, Michael R. and John Morgan. 2004. "Price Dispersion in the Lab and on the Internet: Theory and Evidence." *RAND Journal of Economics* 35(3): 449-466.

Baye, Michael R. and John Morgan. 2009. "Brand and Price Advertising in Online Markets." *Management Science* 55(7): 1139-1151.

Baye, Michael R. and John Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review* 91(3): 454-474.

Baye, Michael R., J. Rupert J. Gatti, Paul Kattuman, and John Morgan. 2006. "Did the Euro Foster Online Price Competition? Evidence from an International Price Comparison Site." *Economic Inquiry* 44(2): 265-279.

Baye, Michael R., J. Rupert J. Gatti, Paul Kattuman, and John Morgan. 2007. "A Dashboard for Online Pricing." *California Management Review* 50(1): 202-216.

Baye, Michael R., J. Rupert J. Gatti, Paul Kattuman, and John Morgan. 2009. "Clicks, Discontinuities, and Firm Demand Online." *Journal of Economics and Management Strategy* 18(4): 935-975.

Bergin, Paul R. and Robert C. Feenstra. 2001. "Pricing-to-market, Staggered Contracts, and Real Exchange Rate Persistence." *Journal of International Economics* 54(2): 333-359.

Bhaskar, V. 2002. "On Endogenously Staggered Prices." *Review of Economic Studies* 69(1): 97-116.

Boivin, Jean, Robert Clark, and Nicolas Vincent. 2012. "Virtual Borders." *Journal of International Economics* 86(2): 327-335.

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

Brynjolfsson, Erik and Michael D. Smith. 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science* 46(4): 563-585.

Campa, Jose Manuel and Linda S. Goldberg. 2005. "Exchange Rate Pass-Through into Import Prices." *Review of Economics and Statistics* 87(4): 679-690.

Cao, Shutao, Wei Dong, and Ben Tomlin. 2012. "The Sensitivity of Producer Prices to Exchange Rates: Insights from Micro Data." Bank of Canada Working Papers 12-20.

Cavallo, Alberto. 2015. "Scraped Data and Sticky Prices." NBER Working Paper 21490.

Cavallo, Alberto, Brent Neiman, and Roberto Rigobon. 2014. "Currency Unions, Product Introductions, and the Real Exchange Rate." *Quarterly Journal of Economics* 129(2): 529-595.

Chevalier, Judith A., and Anil K. Kashyap. 2011. "Best Prices," NBER Working Paper 16680.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong. 2015. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review* 105(3): 993-1029.

Crucini, Mario J., and Mototsugu Shintani. 2008. "Persistence in Law of One Price Deviations: Evidence from Micro-Data." *Journal of Monetary Economics* 55(3): 629-644.

Deck, Cary A., and Bart J. Wilson. 2006. "Tracking Customer Search to Price Discriminate." *Economic Inquiry* 44(2): 280-295.

Deloitte. 2015. *Navigating the new Digital Divide: Capitalizing on Digital Influence in Retail*. Available at https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consumer-business/us-cb-navigating-the-new-digital-divide-v2-051315.pdf.

Devereux, Michael B., and James Yetman. 2010. "Price Adjustment and Exchange Rate Pass-Through." *Journal of International Money and Finance* 29(1): 181-200.

Driscoll, John C., and Aart C. Kraay. 1998. "Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data." *Review of Economics and Statistics* 80(4): 549-560.

Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith, 2014. "How Frequent Are Small Price Changes?" *American Economic Journal: Macroeconomics* 6(2): 137-155.

Engel, Charles, and John H. Rogers. 1996. "How Wide Is the Border?" *American Economic Review* 86(5): 1112-1125.

European Commission. 2013. *Study on the coverage, functioning and consumer use of comparison tools and third-party verification schemes for such tools*.

Available at http://ec.europa.eu/consumers/consumer_evidence/market_studies/docs/final_report_study_on_comparison_tools.pdf

Feenstra, Robert C., Joseph E. Gagnon, and Michael M. Knetter. 1996. "Market Share and Exchange Rate Pass-Through in World Automobile Trade." *Journal of International Economics* 40(1-2): 187-207.

Gaulier, Guillaume, Amina Lahreche-Revil, and Isabelle Mejean. 2006. "Structural Determinants of the Exchange-Rate Pass-Through." CEPII research center Working Paper 2006-03.

Goldberg, Pinelopi K., and Rebecca Hellerstein. 2013. "A Structural Approach to Identifying the Sources of Local-Currency Price Stability." *Review of Economic Studies* 80(1): 175-210.

Goldberg, Pinelopi K., and Michael M. Knetter. 1997. "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature* 35(3): 1243-1272.

Goldberg, Pinelopi K., and Frank Verboven. 2001. "The Evolution of Price Dispersion in the European Car Market." *Review of Economic Studies* 68(4): 811-848.

Gopinath, Gita, and Roberto Rigobon. 2008. "Sticky Borders." *Quarterly Journal of Economics* 123(2): 531-575.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li. 2011. "International Prices, Cost and Markup Differences." *American Economic Review* 101(6): 2450-2486.

Gorodnichenko, Yuriy, and Linda L. Tesar. 2009. "Border Effect or Country Effect? Seattle May Not Be So Far from Vancouver After All." *American Economic Journal: Macroeconomics* 1(1): 219-241.

Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera. 2014. "Price Setting in Online Markets: Evidence from A Major Shopping Platform." NBER Working Paper 20819.

Gust, Christopher, Sylvain Leduc, and Robert Vigfusson. 2010. "Trade Integration, Competition, and the Decline in Exchange-Rate Pass-Through." *Journal of Monetary Economics* 57(3): 309-324.

Head, Allen, Alok Kumar, and Beverly Lapham. 2010. "Market Power, Price Adjustment, And Inflation." *International Economic Review* 51(1): 73-98.

Imbs, Jean, Haroon Mumtaz, Morten O. Ravn, and Hélène Rey. 2005. "PPP Strikes Back: Aggregation and the Real Exchange Rate." *Quarterly Journal of Economics* 120(1): 1-43.

Imbs, Jean, Haroon Mumtaz, Morten O. Ravn, and Hélène Rey. 2010. "One TV, One Price?" *Scandinavian Journal of Economics* 112(4): 753-781.

Kardasz, Stanley W. and Kenneth R. Stollery. 2001. "Exchange Rate Pass-Through and Its Determinants in Canadian Manufacturing Industries." *Canadian Journal of Economics* 34(3): 719-738.

Kehoe, Patrick J. and Virgiliu Midrigan. 2015. "Prices Are Sticky After All." *Journal of Monetary Economics* 75(1): 35–53.

Knotek, Edward. 2011. "Convenient Prices and Price Rigidity: Cross-Sectional Evidence." *Review of Economics and Statistics* 93(3): 1076-1086.

Lünnemann, Patrick and Ladislav Wintr. 2011. "Price Stickiness in the US and Europe Revisited: Evidence from Internet Prices." *Oxford Bulletin of Economics and Statistics* 73(5): 593-621.

Mayoral, Laura and Maria D. Gadea. 2011. "Aggregate Real Exchange Rate Persistence through the Lens of Sectoral Data." *Journal of Monetary Economics* 58(3): 290-304.

Menon, Jayant. 1996. "The Degree and Determinants of Exchange Rate Pass-Through: Market Structure, Non-Tariff Barriers and Multinational Corporations." *Economic Journal* 106(435): 434-444.

Morgan, John, Henrik Orzen, and Martin Sefton. 2006. "An Experimental Study of Price Dispersion." *Games and Economic Behavior* 54(1): 134-158.

Nakamura, Emi, and Jon Steinsson. 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123(4): 1415-1464.

Neiman, Brent. 2010. "Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices." *Journal of Monetary Economics* 57(3): 295-308.

Obstfeld, Maurice and Kenneth Rogoff. 2000. "The Six Major Puzzles in International Macroeconomics. Is There a Common Cause?" in *NBER Macroeconomics Annual* 2000, B.S. Bernanke and K. Rogoff, eds., Cambridge, MA: MIT Press.

Richards, Timothy J., Miguel I. Gómez, and Jun Lee. 2014. "Pass-Through and Consumer Search: An Empirical Analysis." *American Journal of Agricultural Economics* 96(4): 1049-1069.

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

Sheremirov, Viacheslav. 2015. "Price Dispersion and Inflation: New Facts and Theoretical Implications." Federal Reserve Bank of Boston Working Paper No. 15-10.

Statistics Canada. 2013. *Individual Internet use and e-commerce, 2012*. Available at http://www.statcan.gc.ca/daily-quotidien/131028/dq131028a-eng.pdf.

Takhtamanova, Yelena F. 2010. "Understanding Changes in Exchange Rate Pass-Through." *Journal of Macroeconomics* 32(4): 1118-1130.

United States Census Bureau, 2015. *E-Stats 2013: Measuring the Electronic Economy*. Available at http://www.census.gov/econ/estats/e13-estats.pdf.

Yang, Jiawen. 1997. "Exchange Rate Pass-Through In U.S. Manufacturing Industries." *Review of Economics and Statistics* 79(1): 95-104.

Figure 1. Screenshots of typical web pages from price comparison websites.

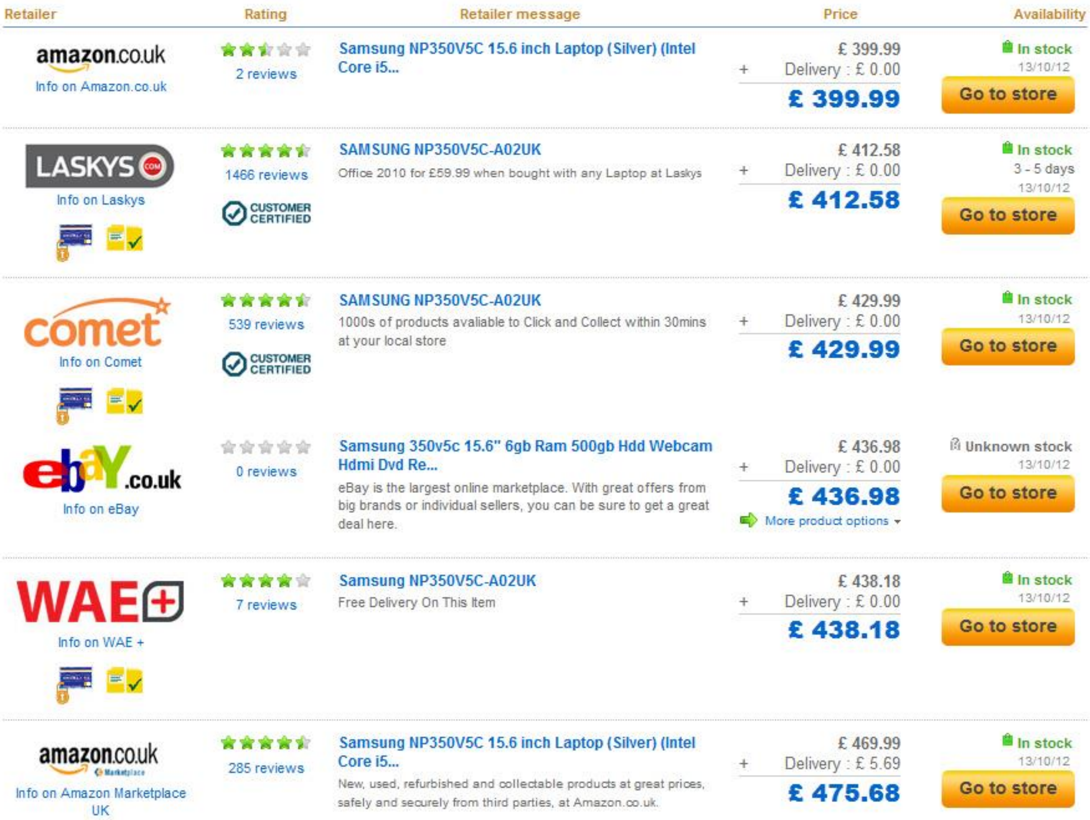

Retailer	Rating	Retailer message	Price	Availability
amazon.co.uk Info on Amazon.co.uk	2 reviews	Samsung NP350V5C 15.6 inch Laptop (Silver) (Intel Core i5...	£ 399.99 + Delivery : £ 0.00 £ 399.99	In stock 13/10/12 Go to store
LASKYS Info on Laskys	1466 reviews CUSTOMER CERTIFIED	SAMSUNG NP350V5C-A02UK Office 2010 for £59.99 when bought with any Laptop at Laskys	£ 412.58 + Delivery : £ 0.00 £ 412.58	In stock 3 - 5 days 13/10/12 Go to store
comet Info on Comet	539 reviews CUSTOMER CERTIFIED	SAMSUNG NP350V5C-A02UK 1000s of products avaliable to Click and Collect within 30mins at your local store	£ 429.99 + Delivery : £ 0.00 £ 429.99	In stock 13/10/12 Go to store
ebay.co.uk Info on eBay	0 reviews	Samsung 350v5c 15.6" 6gb Ram 500gb Hdd Webcam Hdmi Dvd Re... eBay is the largest online marketplace. With great offers from big brands or individual sellers, you can be sure to get a great deal here.	£ 436.98 + Delivery : £ 0.00 £ 436.98 More product options	Unknown stock 13/10/12 Go to store
WAE+ Info on WAE +	7 reviews	Samsung NP350V5C-A02UK Free Delivery On This Item	£ 438.18 + Delivery : £ 0.00 £ 438.18	In stock 13/10/12 Go to store
amazon.co.uk Marketplace Info on Amazon Marketplace UK	285 reviews	Samsung NP350V5C 15.6 inch Laptop (Silver) (Intel Core i5... New, used, refurbished and collectable products at great prices, safely and securely from third parties, at Amazon.co.uk.	£ 469.99 + Delivery : £ 5.69 £ 475.68	In stock 13/10/12 Go to store

Figure 2. Time series of CAD/USD exchange rate.

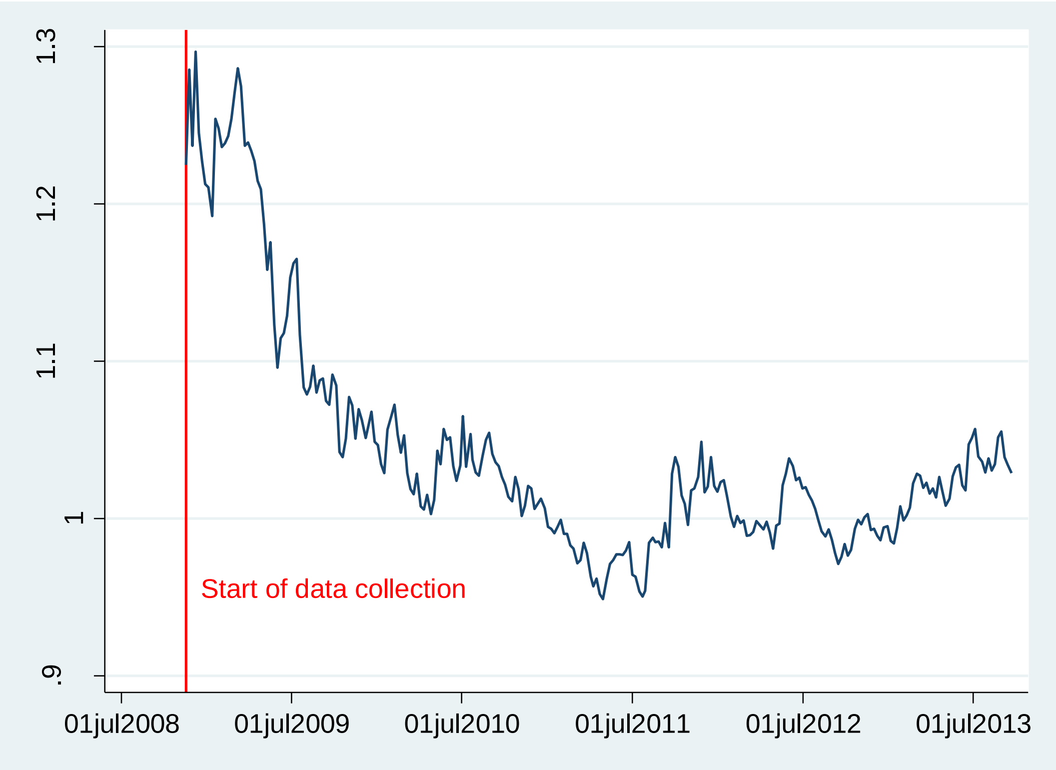

Notes: Source: Board of Governors.

Figure 3. Price quotes.

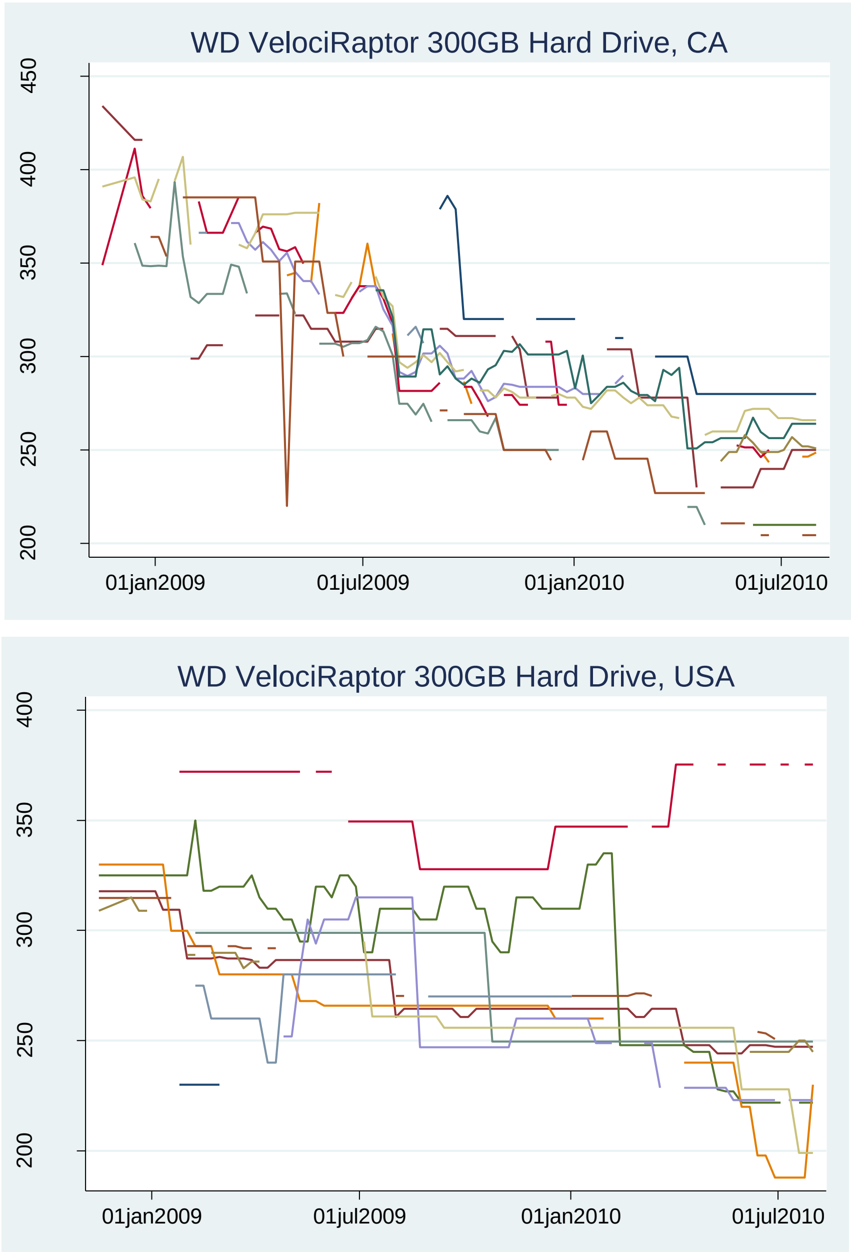

Notes: Each line shows a path of price quotes for a given online seller of the WD VelociRaptor 300Gb hard drive. The left panel is for Canadian sellers. The right panel is for U.S. sellers.

Figure 4. Price quotes listed on the price comparison website and seller websit

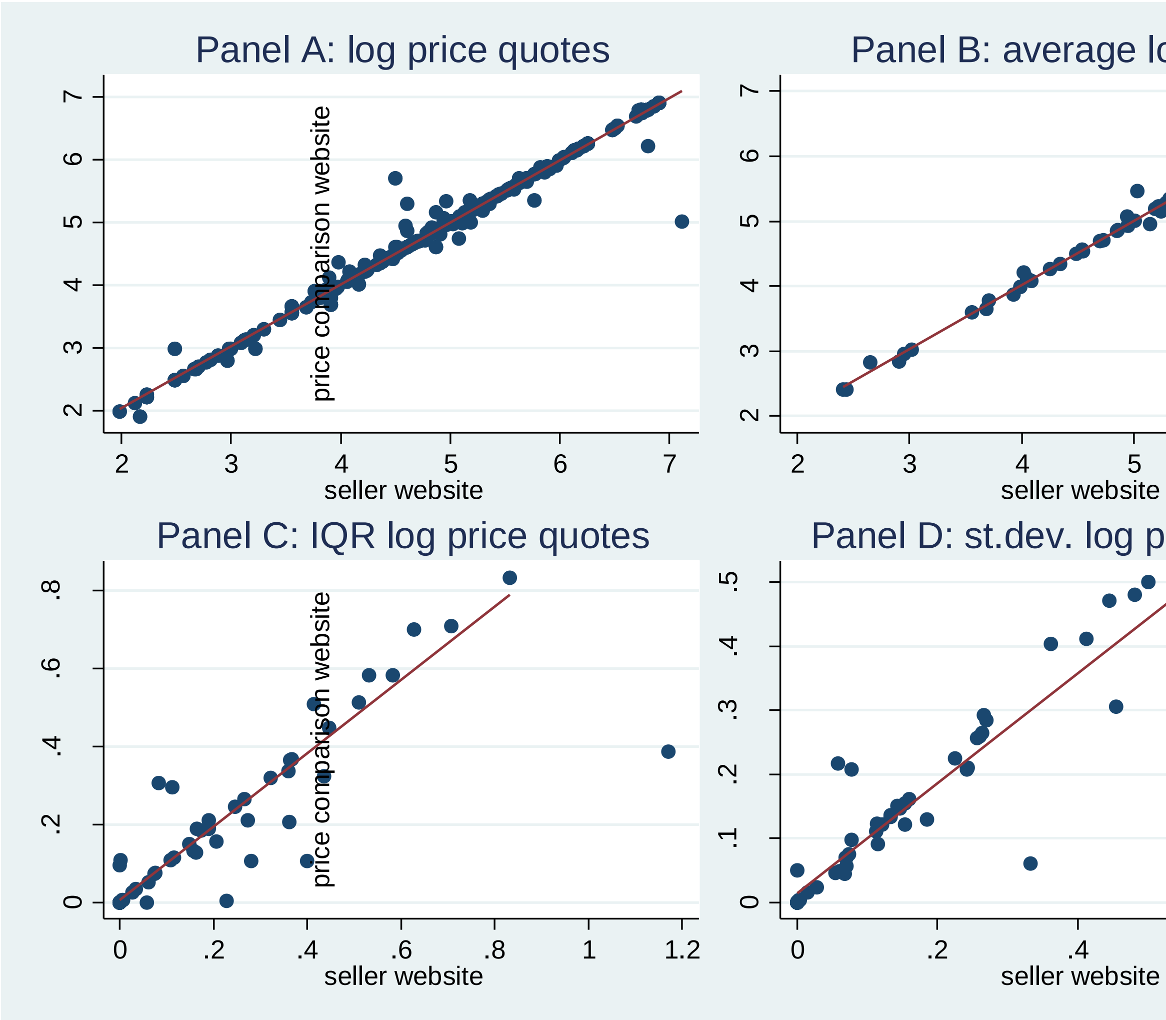

Notes: Panel A shows price quotes listed on the price comparison website and seller websites for each good, that is, each point is a good-
price quote is calculated for each source of price information, that is, each point shows an average log price for a good. Panel C shows the
sellers for each good in both sources of price information. Panel D shows the standard deviation of log prices across sellers for each good

Figure 5. Intensive and extensive margins of price adjustment.

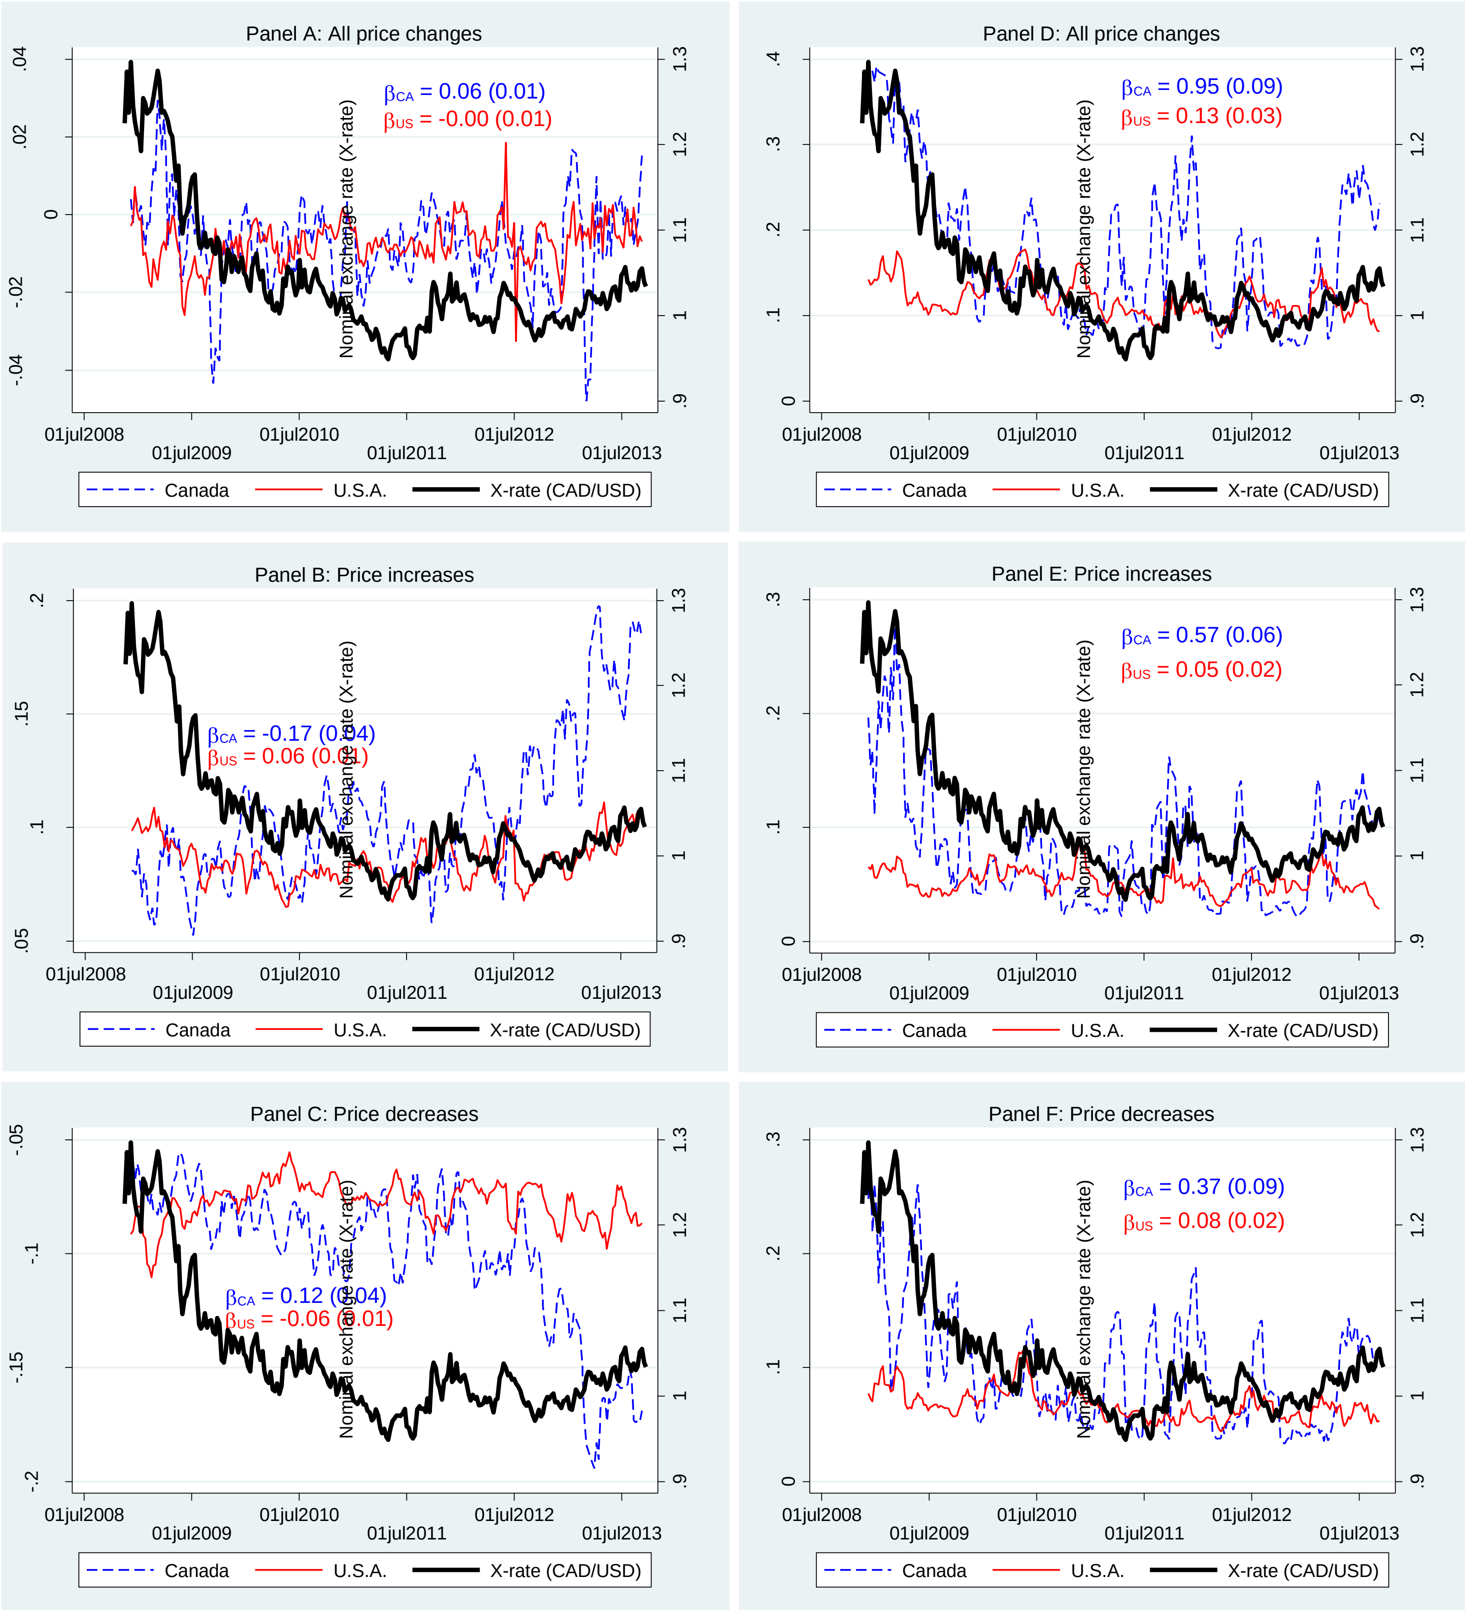

Notes: β_{CA} and β_{US} show estimated slopes of regressing a given variable for Canada and the U.S. on the nominal CAD/USD exchange rate. Newey-West standard errors are in parentheses. See section 4.D for further details.

Figure 6. Exit margin of price adjustment.

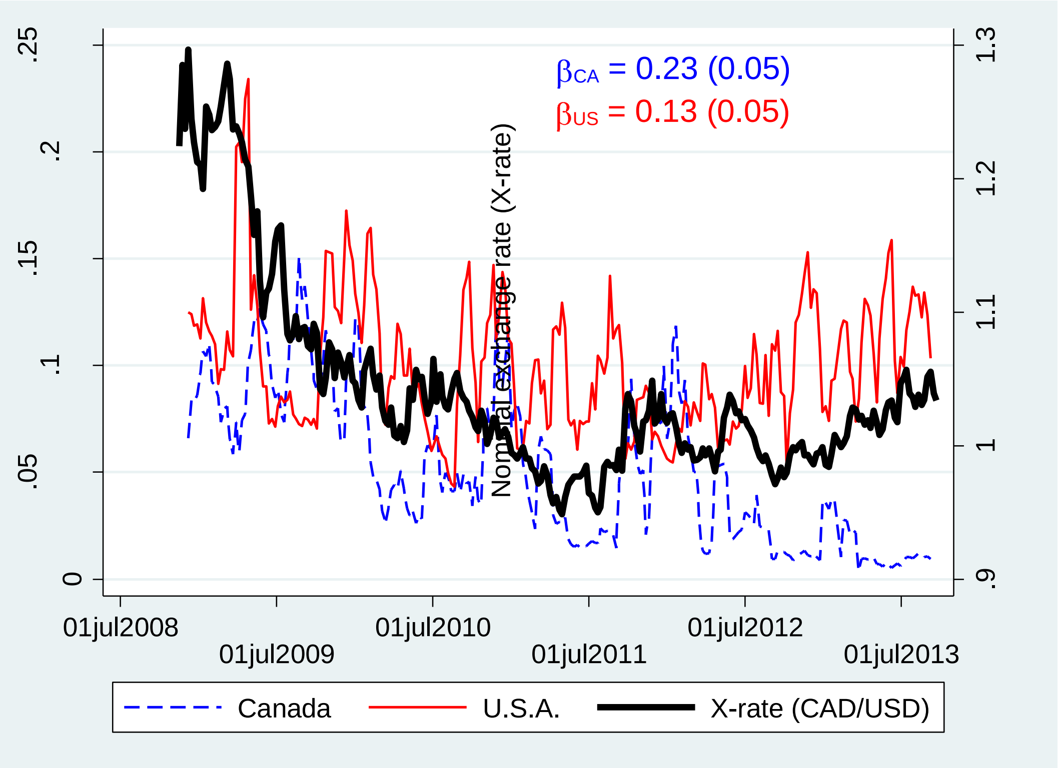

Notes: β_{CA} and β_{US} show estimated slopes of regressing a given variable for Canada and the U.S. on the nominal exchange rate. Newey-West standard errors are in parentheses. See section 4.D for further details.

Table 1. Description of categories.

Category	Type	Quotes	G
Cameras (10 categories)	35mm SLR lens Accessories, Bags and Cases, Binoculars, Camcorders, Camcorder Batteries, Camcorder Accessories, Dedicated Flashes, Digital Cameras, SLR Lenses, Tripods	1,398,396 (543,587)	1 (
Computers (20 categories)	Cases, Desktops, Flash Memory, Flat Panel LCD monitors, Hard Drives, Hubs, Keyboards, Laptop, Laptop Memory, Microphones and Headsets, Modems, Motherboards, Network Adapters, Power Supply, Processors, Scanners, Speakers, Storage Media, UPSS, Webcams	11,260,217 (8,368,381)	5 (1
Electronics (13 categories)	Audio Cables, AV Accessories, Calculators, Cash Registers, GPS, Headphones, MP3 players, Portable Device Accessories, Projectors, Projection Screens, Plasma/LCD TV, TV Accessories, Video Cables	4,313,179 (2,704,025)	3 (8
Software (12 categories)	Anti-Virus, Audio/Video Utilities, Computer Games, Engineering and Design, Databases, Financial and Legal Software, Graphics and Publishing, Office Suites, Programming, Security, System Utilities, Windows Operating Systems	1,628,044 (726,704)	1 (

Notes: The last four columns report the number of unique price quotes, goods, and sellers as well as the median number of goods p
corresponding statistics for the sample of goods used in Table 5.

Table 2. Composition of sellers in the sample.

Seller type	Canada	USA	Pooled
Offline-online	11.53	3.21	7.00
Online-only	78.05	76.21	77.05
Marketplace	-	1.52	0.83
Not classified	10.42	19.06	15.13
Total	100.00	100.00	100.00

Notes: “Offline-online” sellers include stores that sell goods online and that have conventional, brick-and-mortar retail outlets (e.g., Walmart). “Online-only” sellers cover stores that sell goods online and that do not have conventional, brick-and-mortar retail outlets (e.g., Amazon.com). “Marketplace” sellers are multi-vendor online shops (e.g., eBay.com). For “not classified” stores, we could not establish if a seller has a conventional retail outlet.

Table 3. Descriptive statistics.

	Mean (1)	St.Dev (2)	Median (3)	P25 (4)	P75 (5)
Panel A: Canada					
Cross-sectional distribution of prices					
St.dev. log(Price)	0.128	0.090	0.111	0.066	0.160
IQR log(Price)	0.111	0.083	0.091	0.051	0.158
Median log(Price)	5.403	1.407	5.292	4.448	6.602
Frequency of price changes	0.367	0.169	0.367	0.246	0.462
Size of price changes					
Median dlog(Price)	-0.006	0.019	-0.003	-0.007	-0.002
Median abs(dlog(Price))	0.029	0.044	0.017	0.008	0.031
Sales					
Mean size	0.067	0.101	0.028	0.018	0.071
Frequency	0.027	0.032	0.023	0.000	0.039
Synchronization of price changes	0.231	0.210	0.182	0.037	0.374
Properties of sellers					
Number of sellers	2.426	1.209	1.871	1.585	3.127
Stability	0.899	0.065	0.907	0.850	0.952
Freq. of convenient prices	0.196	0.187	0.137	0.061	0.262
Panel B: USA					
Cross-sectional distribution of prices					
St.dev. log(Price)	0.159	0.113	0.140	0.077	0.220
IQR log(Price)	0.173	0.139	0.142	0.075	0.250
Median log(Price)	5.328	1.415	5.191	4.365	6.541
Frequency of price changes	0.197	0.155	0.191	0.055	0.300
Size of price changes					
Median dlog(Price)	-0.006	0.033	-0.004	-0.011	0.000
Median abs(dlog(Price))	0.042	0.052	0.030	0.017	0.049
Sales					
Mean size	0.071	0.087	0.046	0.026	0.082
Frequency	0.022	0.031	0.010	0.000	0.035
Synchronization of price changes	0.187	0.124	0.176	0.101	0.258
Properties of sellers					
Number of sellers	3.370	1.920	2.870	1.868	4.306
Stability	0.887	0.052	0.887	0.856	0.926
Freq. of convenient prices	0.194	0.203	0.141	0.034	0.280
Panel C: International price differentials					
Mean prices					
Relative exchange rate	0.074	0.225	0.050	-0.035	0.183
Real exchange rate	0.051	0.218	0.034	-0.048	0.142
Median prices					
Relative exchange rate	0.081	0.227	0.056	-0.028	0.189
Real exchange rate	0.058	0.221	0.038	-0.039	0.148
Minimum prices					
Relative exchange rate	0.123	0.272	0.085	-0.007	0.234
Real exchange rate	0.100	0.268	0.069	-0.025	0.196

Notes: P25 and P75 in columns (4) and (5) show 25th and 75th percentile of the statistics indicated in the first column. Relative exchange rate is calculated as $\log(P_{it}^{CA}/P_{it}^{US})$ where *i* and *t* index goods and weeks, respectively, P^{CA} is the price in Canada, and P^{US} is the price in the U.S. The real exchange rate is calculated as $\log\left(EX_t{}^{-1} \times P_{it}^{CA}/P_{it}^{US}\right)$ where EX_t is the nominal CAD/USD exchange rate. See text for further details.

Table 4. Comparison of pricing moments

		Price comparison website	Leading shopping platform: no weights	Leading shopping platform: click weighted	Conventional stores
		(1)	(2)	(3)	(4)
Frequency of posted price changes, per week					
EE011	Personal Computers and Per. Equipment	27.15	16.25	21.94	7.74
EE021	Computer Software	20.32	13.33	24.17	2.60
EE042	Calculators and Adding Machines	10.10	9.81	14.74	1.95
RA011	Televisions	28.80	25.76	23.10	7.02
RA051	Radio and Tape Recorders/Players	14.94	11.35	20.37	5.22
RD012	Still Camera	24.90	11.37	33.28	4.47
Mean $\lvert \Delta \log P \rvert$, percent					
EE011	Personal Computers and Per. Equipment	4.77	11.50	11.57	11.26
EE021	Computer Software	8.00	11.41	11.47	22.65
EE042	Calculators and Adding Machines	11.10	19.67	17.64	19.94
RA011	Televisions	5.00	7.36	8.20	9.71
RA051	Radio and Tape Recorders/Players	8.94	16.72	17.00	12.60
RD012	Still Camera	7.32	13.33	13.37	10.54
Frequency of sales, per week					
EE011	Personal Computers and Per. Equipment	2.80	1.21	1.95	5.87
EE021	Computer Software	2.91	0.66	1.71	6.12
EE042	Calculators and Adding Machines	2.90	0.81	0.98	6.02
RA011	Televisions	2.80	1.51	2.19	12.30
RA051	Radio and Tape Recorders/Players	3.53	1.08	1.84	14.12
RD012	Still Camera	3.86	0.99	2.76	9.73
Mean abs. size of sales, percent					
EE011	Personal Computers and Per. Equipment	5.67	10.23	9.75	9.32
EE021	Computer Software	8.40	7.59	9.65	18.21
EE042	Calculators and Adding Machines	6.40	-	-	14.93
RA011	Televisions	6.70	11.94	13.74	6.61
RA051	Radio and Tape Recorders/Players	9.52	15.12	12.38	9.71
RD012	Still Camera	8.49	10.70	11.74	7.78
Cross-sectional dispersion, $st.dev. \log P$, percent					
EE011	Personal Computers and Per. Equipment	10.63	20.80	14.40	-
EE021	Computer Software	20.03	14.80	13.70	-
EE042	Calculators and Adding Machines	16.70	18.70	22.70	-
RA011	Televisions	8.80	14.10	11.60	-
RA051	Radio and Tape Recorders/Players	17.84	18.80	16.90	-
RD012	Still Camera	8.94	14.70	12.80	-
Within-good price synchronization					
EE011	Personal Computers and Per. Equipment	20.18	15.09	17.69	-
EE021	Computer Software	15.98	8.48	15.41	-
EE042	Calculators and Adding Machines	5.40	12.49	16.13	-
RA011	Televisions	17.40	18.19	20.15	-
RA051	Radio and Tape Recorders/Players	12.02	9.53	17.50	-
RD012	Still Camera	20.08	11.53	23.27	-

Notes. The table compares the frequency and absolute size of price changes and sales, cross-sectional dispersion and price within-good price synchronization for selected narrow categories in online data used in this paper, data used in Gorodnichenko, Sheremirov and Talavera (2014), and data for conventional stores (column 4) are from Nakamura and Steinsson (2008). All data are for the U.S. Only matched categories are shown.

Table 5. Pass-through and the speed of price adjustment.

	No Fixed effects	Type Fixed effects	Good Fixed effects	N
	(1)	(2)	(3)	(4)
Panel A: Pass-through				
Mean Price	0.765	0.727	0.670	1,739,845
	(0.100)	(0.091)	(0.086)	
Median Price	0.747	0.710	0.666	1,739,384
	(0.101)	(0.092)	(0.089)	
Minimum Price	0.706	0.695	0.620	1,738,222
	(0.071)	(0.061)	(0.045)	
Panel B: Speed of Adjustment				
Mean Price	-0.062	-0.067	-0.154	1,400,705
	(0.004)	(0.004)	(0.007)	
Median Price	-0.070	-0.075	-0.168	1,399,840
	(0.004)	(0.004)	(0.007)	
Minimum Price	-0.069	-0.075	-0.162	1,399,055
	(0.004)	(0.004)	(0.007)	
Panel C: Intra-seller prices				
Pass-through	0.553	0.240	0.206	84,143
	(0.069)	(0.060)	(0.060)	
Speed of Adjustment	0.005	-0.055	-0.100	63,496
	(0.017)	(0.013)	(0.027)	

Notes: Panel A presents estimates of α in specification (1). Panel B presents estimates of β in specification (2). Panel C reports estimates of α (the first row) and β (the second row) for the sample of price quotes by the same seller in the U.S. and Canada. Driscoll and Kraay (1998) standard errors are in parentheses.

Table 6. Determinants of pass-through and the speed of price adjustment.

	Pass-Through, $\hat{\alpha}$			Speed of Adjustment, $\hat{\beta}$		
	Mean price	Median price	Minimum price	Mean price	Median price	Minimum price
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Median Price)	0.227	0.338	0.566	0.051	0.048	0.022
	(0.088)	(0.087)	(0.122)	(0.008)	(0.009)	(0.009)
Log(Median Price)^2	-0.024	-0.033	-0.053	-0.004	-0.004	-0.002
	(0.008)	(0.008)	(0.011)	(0.001)	(0.001)	(0.001)
Freq. of price change	1.947	1.964	2.062	-0.126	-0.132	-0.143
	(0.194)	(0.183)	(0.224)	(0.017)	(0.017)	(0.025)
Log(Sellers)	1.287	1.262	1.498	-0.025	-0.016	0.000
	(0.282)	(0.299)	(0.279)	(0.030)	(0.033)	(0.037)
Log(Sellers)^2	-0.421	-0.404	-0.486	0.010	0.006	-0.000
	(0.084)	(0.091)	(0.087)	(0.008)	(0.009)	(0.010)
Stability of Sellers	0.296	0.548	-0.969	0.871	0.966	1.014
	(0.658)	(0.586)	(0.643)	(0.074)	(0.082)	(0.082)
Synchronization	-0.342	-0.366	-0.356	0.035	0.013	-0.017
	(0.157)	(0.152)	(0.160)	(0.017)	(0.016)	(0.015)
Average Reputation	-0.120	-0.127	0.011	-0.015	-0.018	-0.025
	(0.057)	(0.055)	(0.064)	(0.005)	(0.006)	(0.007)
Freq. of Sales	1.040	1.157	0.635	-0.402	-0.388	-0.400
	(0.756)	(0.798)	(0.616)	(0.054)	(0.056)	(0.065)
Freq. of Convenient Prices	0.111	0.178	0.028	0.024	0.030	-0.018
	(0.101)	(0.097)	(0.161)	(0.011)	(0.014)	(0.014)
Observations	21,734	21,667	21,750	22,068	22,118	22,072
R^2	0.15	0.15	0.25	0.16	0.16	0.18
Descriptive statistics for dependent variables						
Mean	0.636	0.639	0.904	-0.347	-0.365	-0.491
St.Dev.	1.908	1.951	2.380	0.342	0.347	0.856
Median	0.616	0.608	0.860	-0.223	-0.244	-0.231
P25	-0.091	-0.101	-0.039	-0.472	-0.495	-0.467
P75	1.407	1.406	1.881	-0.106	-0.118	-0.105

Notes: Columns (1)-(3) and (4)-(6) report estimated specification (3) for pass-through and the speed of price adjustment, respectively. Category fixed effects C_i and time fixed effects T_i are included but not reported. The regressions are run on samples where top and bottom 1 percent of estimated $\hat{\alpha}$ and $\hat{\beta}$ are winsorized. Standard errors are clustered by good type. The last two rows show 25th and 75th percentiles. The number of goods is 24,129.

Table 7. Margins of price adjustment.

	Mean price		Median price		Minimum Price	
	CA	US	CA	US	CA	US
	(1)	(2)	(3)	(4)	(5)	(6)
Mean price change						
Any, $\overline{dP}_{ict}$	-0.128	0.066	-0.109	0.059	-0.081	0.039
	(0.014)	(0.006)	(0.013)	(0.006)	(0.008)	(0.003)
Increase, $\overline{dP}_{ict}^{increase}$	-0.046	0.031	-0.031	0.019	-0.037	0.052
	(0.011)	(0.008)	(0.010)	(0.006)	(0.005)	(0.003)
Decrease, $\overline{dP}_{ict}^{decrease}$	-0.088	0.051	-0.073	0.047	-0.055	0.002
	(0.011)	(0.006)	(0.009)	(0.005)	(0.008)	(0.002)
Probability of price adjustment						
Any, $\Pr(dP \neq 0)$	-0.008	0.009	-0.006	0.005	-0.019	0.010
	(0.015)	(0.006)	(0.015)	(0.005)	(0.013)	(0.003)
Increase, $\Pr(dP > 0)$	-0.085	0.029	-0.079	0.027	-0.061	0.023
	(0.010)	(0.005)	(0.009)	(0.005)	(0.007)	(0.003)
Decrease, $\Pr(dP < 0)$	0.076	-0.020	0.072	-0.022	0.042	-0.013
	(0.011)	(0.004)	(0.011)	(0.004)	(0.010)	(0.002)
Probability of exit						
$\Pr(exit)$	-0.015	-0.001	-0.015	0.004	-0.045	0.034
	(0.009)	(0.007)	(0.008)	(0.007)	(0.005)	(0.005)

Notes: The table reports estimated ψ in specification (8). Good fixed effects are included but not reported. Newey-West standard errors are in parentheses.