

Quality and price setting of high-tech goods

Gorodnichenko, Yuriy; Talavera, Oleksandr; Vu, Nam

DOI:

[10.1016/j.econmod.2021.02.010](https://doi.org/10.1016/j.econmod.2021.02.010)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Gorodnichenko, Y, Talavera, O & Vu, N 2021, 'Quality and price setting of high-tech goods', *Economic Modelling*, vol. 98, pp. 69-85. <https://doi.org/10.1016/j.econmod.2021.02.010>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Quality and Price Setting of High-Tech Goods

Yuriy Gorodnichenko^a, Oleksandr Talavera^b, Nam Vu^c

Abstract:

This paper investigates the link between product quality and price setting for central processing units (CPUs). Using thousands of price quotes from a popular price-comparison website, we find that market fundamentals, such as the number of sellers, median price, share of convenient prices and level of seller stability, are important factors for explaining price stickiness and price dispersion. We demonstrate that calculations of price inflation require conditioning not only on CPU quality, but also on market fundamentals to ensure that CPU attributes are priced correctly. Failing to do so can result in an understatement of CPU price deflation in the sample period.

JEL classification: E31, L11, L81, L86

Keywords: price setting, e-commerce, product quality, hedonic pricing, inflation.

^a Department of Economics, University of California-Berkeley, Berkeley, CA 94720. Email: ygorodni@econ.berkeley.edu

^b Department of Economics, University of Birmingham, Edgbaston, Birmingham, United Kingdom B15 2TT. Email: oleksandr.talavera@gmail.com

^c Department of Economics, University of Birmingham, Edgbaston, Birmingham, United Kingdom B15 2TT. Email: hoainam2207@gmail.com

* Standard disclaimer applies. We thank Tho Pham and Daniel Sichel for their valuable support and comments. We are also grateful to participants of Swansea Workshop on Prices and Nowcasting, Scottish Economic Society Annual Conference, Royal Economic Society Annual Conference, ESCoE's Conference on Economic Measurement, Workshop on Empirical Macroeconomics, and MMF Annual Conference for valuable comments and discussions. Any remaining errors are our own.

1. Introduction

The law of one price is the bedrock of many economic models, but the reality often contradicts the law: similar (or even identical) goods can be sold at different prices in highly integrated markets. One may hope that the law of one price may still hold for prices adjusted for quality differences (“price per performance unit”) and other product attributes. This idea is the basis of hedonic pricing used to measure inflation for goods with rapid technical change. Indeed, since Solow (1960), economists and statistical agencies routinely postulate that vintages of capital produced with different technologies should have the same price when these vintages are measured in efficiency units. While intuitive, this approach ignores the fact that price dispersion may arise for other reasons such as market power, price stickiness, and other frictions. Consequently, conventional hedonic pricing may provide a distorted measure of inflation.

To evaluate this hypothesis, we use detailed data on price quotes and performance measures for Central Processing Units (CPUs) to study how quality differences and various frictions influence measured inflation and technical change. We focus on the CPU market for several reasons. First, CPUs have a high rate of quality improvement which helps us to explore the effect of quality changes on prices, as many vintages of CPU technology coexist in a given period. Second, because CPUs are a key element of many electronic devices, quality improvement in CPUs is the major factor which contributes to the increase in the quality of other products and services (desktops, laptops, tablets, mobile phones, software, and cloud computing); see e.g., Jorgenson et al., 2000; Oliner and Sichel, 2000; Brynjolfsson and Hitt, 2003. Finally, CPU quality can be accurately measured using CPU performance scores, which removes the difficulty of quality measurement encountered by previous studies (see e.g., Combris et al., 1997).

A central ingredient of our analysis is our unique dataset, which includes monthly online prices of CPUs in the U.S. The online price quotes are collected from a leading online shopping platform. Our data covers the prices of hundreds of CPUs sold across hundreds of online retailers between April 2009 and December 2012. The substantial size of our data gives us strong statistical power. Finally, for each product, we have CPU performance scores, an integral measure of CPU quality.

Using the dataset, we show that, similar to prices for other products, CPU prices exhibit some stickiness even in a highly competitive e-commerce environment, thus suggesting that various

pricing frictions and market fundamentals—the number of sellers (a proxy for market concentration), median price (a proxy for incentives to search for better prices), share of convenient prices (also called price points such \$9.99 or \$99)—are also important factors for explaining price stickiness and price dispersion. We document a variety of stylized facts about CPU pricing and contrast various pricing moments for CPUs with the corresponding moments for typical goods sold online. Our results suggest that quality differences play an important role in accounting for price differences across different CPUs. When we combine market frictions and quality differences to explain variation in CPU prices, we find that both types of forces are significant predictors of price variation. Using the introduction of new (sometimes “breakthrough”) CPU models, we also examine how prices of other models evolve in response to these technological shocks. We document that the re-pricing of older, less efficient/powerful models is rather modest and sluggish. Finally, we augment standard hedonic-pricing regressions to include various measures of market/product fundamentals to investigate how these fundamentals can affect measured inflation for quality-adjusted products. The estimates suggest that controlling for market/product fundamentals materially influences estimated inflation: prices fall at a faster rate than is implied by quality adjustment alone.

Our work is related to several strands of the literature. The first strand studies the micro foundation of price stickiness (see e.g., Klenow and Malin, 2010 for a survey). Our analysis complements this strand by focusing on goods with rapid quality improvement and by studying the effects of well-measured technological changes on price stickiness and price dispersion.

The second strand of literature focuses on price dispersion at the micro-level. Existing papers often focus on the degree of price dispersion (Lach, 2002; Kaplan et al., 2019; Sheremirov, 2019). Our study contributes to this literature by exploring the impact of product quality on between- and within-product price dispersion, as well as documenting the dynamic properties of price dispersion.

Finally, this study is related to the strand of literature that focuses on constructing hedonic price indices (see e.g., Shiratsuka, 1999; Kryvtsov, 2016; Byrne et al., 2020), especially for computing devices (e.g., Pakes, 2003; Aizcorbe et al., 2020) and computer components (e.g., Byrne et al., 2018). We complement this strand of literature by introducing an improved, direct measure of product quality for a larger set of CPUs to construct the quality-adjusted price indices. More importantly, we quantify the contribution of non-technological factors to correctly measure inflation.

2. Data

2.1. Data collection

We use a comprehensive dataset of CPU price quotes in the U.S. e-commerce market. Each product is uniquely identified by the manufacturer product number (MPN). For instance, MPN “BX80601940” uniquely identifies the “*Intel Core i7-940 2.93GHz Processor*”. Similarly, each online seller (e.g., Amazon.com, CostCentral.com, Memory4Less.com, NextWarehouse.com) is uniquely identified. Our online price-quotes are gathered from a leading price comparison website (PCW) that provides price quotes for the U.S. online market.¹ Specifically, at midnight on the first day of each month, a python script starts to automatically download webpages with price quotes. After that, we extract MPNs, seller IDs and prices for each CPU-Seller. Price quotes in our data are “net”, that is, pre taxes and shipping fees. Relative to previous studies that typically have one year of data or less (e.g., Lünemann and Wint, 2011), we have a longer time series covering 45 months. This feature of our data allows us to better exploit changes in the quality of CPUs over time.

We complement these price quote data with information on the performance scores of CPUs. In this additional dataset, we uniquely identify each CPU model by its official name. Besides performance scores, the data also include the main technical characteristics of the processor such as speed, turbo speed, and the number of cores. The data of CPU performance scores are provided by PassMark Software, a leading authority in software/hardware performance benchmarking and testing. This company is also a Microsoft Partner and Intel Software Partner and owns one of the world’s largest CPU benchmark websites. They not only report the final scores of CPU performances but also provide all of the test results used to compute the performance marks. Furthermore, their testing methods and models to produce the CPU’s performance score are published on their website. Therefore, if they change their performance measure, we can produce new performance scores and update our data to be comparable with new CPU models. Since each CPU model in the latter dataset includes several product versions with different MPNs in the former dataset, we manually matched the performance score of a CPU model with all of its versions. For example, the CPU model “*Intel Core i7-940 2.93GHz*

¹ Gorodnichenko and Talavera (2017) use the same PCW and they provide more details on the prominence and business model of this PCW.

Processor” has three versions with the same performance. Its MPNs are: “AT80601000921AA”, “BX80601940”, and “BXC80601940”.

We use performance scores rather than physical characteristics (e.g., clock speed)—as was done in previous research—because the nature of technical change can evolve over time and some physical characteristics can lose their relevance. For example, Byrne et al. (2018) find that the CPU quality index based on chip characteristics is completely flat over 2010-2013, while the quality index based on performance scores increased sharply. The main reason is that processor producers shifted away from increasing the clock speed due to heat generation. Instead, they improved processor performance by placing multiple cores on a chip.² We complement existing literature by using the performance score as an integral measure of quality for a large set of CPUs.

2.2. Data filters and data quality

We applied a series of filters to ensure the consistency of our analysis. Specifically, we dropped CPUs that do not have performance scores after merging our two data sources. All used or refurbished CPUs were removed from the dataset since their prices are not comparable with prices of new CPUs. In addition, to minimize the effects of extreme values in our data, both the top and the bottom one percent of the prices were dropped. For time-series analyses, CPUs with less than 20 price-quote observations were removed. A CPU is considered as an available product in a month if it is offered by at least three sellers in that month. After applying all filters, our large sample covers monthly prices of 654 CPUs sold across 142 online retailers in the United States from April 2009 to December 2012. We define an observation by its MPN, seller ID and month. Our dataset contains 74,306 product-seller-month price quotes. The CPU market is dominated by Intel and AMD and therefore the price quotes are effectively for these two manufacturers.³

One may be concerned about the quality of price quotes on PCWs, e.g., the price listed on PCWs may be out of date or discrepant from sellers’ actual prices if sellers post lower prices on PCWs than on their websites to attract visits of customers. In fact, online merchants are

² Flamm (2019) provides a review of changes in technology.

³ In the early 2000s, Intel and AMD tended to review their prices quarterly (Goettler and Gordon, 2011). However, at the end of the decade, the frequency of price reductions of Intel products decreased substantially (Byrne et al., 2018).

incentivized to keep updating their latest prices on PCWs as they usually have to pay for clicks on those webpages. Therefore, if their prices are not up to date, they will not gain sales and will waste their money. Similar to our dataset, Gorodnichenko and Talavera (2017) gathered price data from PCWs and find that prices on PCWs and seller websites are highly consistent (correlation 0.98). In addition, they find that price data from PCWs are consistent with the Bureau of Labour Statistics (BLS) data and are updated rapidly in response to shocks. Hence, the price data from the PCWs are of a reasonably high quality.

2.3. Notation and aggregation

We use P_{ist} to denote the price of CPU i sold by seller s at month t . Q_i denotes the performance score of CPU i . We denote the set of all CPUs, sellers and time as $\mathcal{C} = \{1, \dots, C\}$, $\mathcal{S} = \{1, \dots, S\}$, $\mathcal{T} = \{1, \dots, T\}$. For example, C_{st} is the total number of CPUs, which are offered by seller s at time t , while S_{it} represents the total number of sellers that sell CPU i at time t . We will use lower-case letters to denote logs of corresponding variables; e.g., $p_{ist} \equiv \log(P_{ist})$. Bars denote average values; e.g., \bar{p}_{it} is the average log price of CPU i across all sellers at time t .

We aggregate statistics for a given moment m (frequency of price changes, size of price changes, price dispersion, etc.) as follows:

$$\bar{m} = \frac{1}{C} \sum_i \frac{1}{S} \sum_s m_{is}. \quad (1)$$

To highlight the differences in the price per unit of performance, we compute the quality-adjusted price $P_{ist}^{adj} \equiv P_{ist} \times (Q^{(50)}/Q_i)$ where $Q^{(50)}$ is the median performance score between all CPU models.

3. Basic facts

This section describes the evolution of the CPU market over the sample period. We will use institutional details to study how quality improvement affects CPU pricing.

3.1. CPU generations over the last decade

The concept of CPU generations was introduced after Intel released its CPU core i series in *Nehalem* microarchitecture family, which is also known as the first generation. The major difference between CPU generations is the differences in their microarchitecture, which is

reflected by the semiconductor manufacturing process (also called “technology node” or “process technology”). The “process technology” is designated by the process’s minimum feature size that is indicated by the size in nanometres. This size refers to the average half-pitch (half the distance between identical features) of a memory cell (Hoefflinger, 2011). The smaller the process’s minimum feature size, the more powerful and the more efficient in energy consumption the processor. Chip companies were able to shrink the size of their microprocessor and improve the CPU performance, mainly because of the innovation in the semiconductor industry.

Figure 1 shows that Intel was the pioneer of microprocessor technology in our sample period from 2009 to 2012. The first generation of the Intel Core processors (*Nehalem*) was released in November 2008 and uses the 45 nanometres (nm) process, the second-generation processors (*Sandy Bridge*) and the third-generation processors (*Ivy Bridge*) use 32 nm and 22 nm, respectively. Meanwhile, AMD was struggling in the technology race with Intel. AMD released its first-generation CPUs (*Bulldozer* family) using 32 nm technology 9 months after the release of the *Sandy Bridge* processors. The second-generation CPUs of AMD (*Piledriver* family), which were released five months after Intel released *Ivy Bridge* CPUs, were still based on the 32 nm technology.

Table 1 shows that there is variation in performance scores across CPUs within a generation, because manufacturers can tweak various attributes and features depending on market demand. As a result, CPUs in a later generation are not always better than CPUs from earlier generations. However, newer generations do tend to have better performance on average. For example, CPUs in the *Nehalem* generation have, on average, more than 100% improvement in performance scores relative to older CPUs. CPUs in the *Ivy Bridge* generation have, on average, 30% higher scores than CPUs in the *Nehalem* and *Sandy Bridge* generations, but the *Sandy Bridge* generation was roughly comparable to the *Nehalem* family. On average, a new generation improves performance scores on the order of 50%.

Panel A of Figure 2 shows the number of CPU models of each generation over time following a new CPU generation release. The number of available new generation CPU models gradually increases, while old generation CPU models leave the market. It implies that processor producers do not release all models of a generation when they introduce new technology. Instead, they often release a few new models first, then launch more models using that new microarchitecture and stop selling old generation products.

To adopt the technological upgrades in the market from processor manufacturers and consumer demand, sellers respond by updating the list of CPUs that they are offering. Panels B and C of Figure 2 clearly show a rapid increase in the number of sellers that offer new generation processors and in the number of new-generation models offered by a seller, respectively. As a result, the number of observations of new technology CPUs dramatically rose after their releases. These facts suggest that the change in performance scores of CPUs offered by sellers, which is caused by the entry of new CPUs and the exit of old CPUs, should be able to reflect the technological upgrades in the market.

3.2. Price distribution and performance

Table 2 shows the average price of each percentile of the distribution over products (\bar{P}_i), the mean and the standard deviation of the average log price (\bar{p}_i) within the sample. Overall, the median CPU in our data costs \$190.43 and 25% of the CPUs are priced under \$99.95; CPUs that are more expensive than \$334.93 account for the top 25% of highest prices of the sample. When we adjust prices for performance scores, we observe that price dispersion (measured as the standard deviation of log prices) falls by 6 log points from 0.90 to 0.84.

To further explore the relationship between prices and performance, we compute log deviations for prices and performance scores from the median CPU in a given month and present a binscatter plot in Figure 3. Each bin (point in the figure) corresponds to one percent of the sample. The lowess smoother (bandwidth set to 0.1) plots the non-parametric fit. The figure shows that better-performing CPUs command higher prices. However, the relationship is non-linear (concave) as the curve is flatter on the right-hand side of the price distribution, demonstrating that consumers have to spend increasingly more money to achieve a given gain in performance.

To compare the price per performance unit of CPUs in different quality standards or price levels, we divide our sample into four quartile groups of CPUs based on their performance scores. Panel A of Figure 4 shows the distribution of average CPU prices (\bar{P}_i) by quartile. The estimated kernel densities suggest that high-performance CPUs are sold at higher prices and the peaks of estimated densities are shifted to the right as we move from the first performance quartile (lowest performance) to the fourth performance quartile (highest performance). Panel B of Figure 4 shows the corresponding distribution of CPU prices adjusted for quality, i.e., the

average price *per performance unit* of CPUs (\bar{P}_i^{adj}). Note that the distance between distributions is discernibly attenuated and therefore some (but not all) of the observed price dispersion is accounted for by differences in CPU productivity.

Finally, we split our sample by CPU generation (described in the previous section). Panel C of Figure 4 presents the distribution of average prices \bar{P}_i within a CPU generation, while Panel D presents the corresponding results for quality-adjusted prices \bar{P}_i^{adj} . The price per unit of performance for AMD processors was generally lower than the price for Intel processors. This discount could reflect the fact that AMD was often catching up with Intel and thus AMD had to offer discounts to compete with the dominance of Intel in the CPU market (see e.g., Goettler and Gordon, 2011). In addition, we observe that the technical improvement used in later CPU generations benefits customers by reducing the price per unit of performance. However, the difference in the average price per performance unit between a CPU generation and its next-generation is usually small.

3.3. Dynamics of CPU prices

Panel A of Figure 5 shows that the price level in the CPU market (\bar{P}_t) slightly increased as the new-generation CPUs were released at higher prices, but the quality-adjusted prices (\bar{P}_t^{adj}) dropped quickly due to the significant improvement in CPU performance, which is consistent with earlier studies (e.g., Byrne et al., 2017; Byrne et al., 2018).

Panel B of Figure 5 illustrates the dynamics of the median price of each generation. The median prices of AMD generations were often lower than prices for Intel counterparts. We also see that CPU producers usually set higher prices for new generation processors and the median prices tend to gradually decrease over time.⁴ However, the median prices of *Intel Core 2* and *Intel Nehalem* CPUs were consistently higher than the median prices of new generations and remained stable for most of our sample period. This can be explained by the large demand for those two processor generations in our sample period. Since those old CPUs were popular for a long time without major changes to the processor architecture, the old systems, which are not compatible with new processor generations, were also popular. Thus, a large number of customers might choose to purchase an old-generation processor rather than upgrading the

⁴ *Nehalem* prices dropped sharply in early 2010 when Intel introduced a new model in the *Nehalem* family.

whole system, while the supply of these old models is limited as Intel stopped producing them. This could lead to a high price level of *Core 2* and *Nehalem* CPUs.⁵ We observe similar patterns when we examine the dynamics of prices adjusted for quality (Panel C of Figure 5).

4. Pricing facts for CPU e-commerce

In this section, we document the basic facts about price setting for CPUs in the online retail market. Specifically, we focus on the frequency of price changes, size of price changes, and within-CPU price dispersion. To compute these moments, we follow Gorodnichenko, Sheremirov and Talavera (2018) and provide pertinent definitions and notation in Appendix B. We then relate each of these moments to market fundamentals and contrast these moments with the corresponding moments for typical goods sold online, which generally have a lower rate of technical change. Note that one should expect to find smaller pricing frictions in online markets than in conventional “bricks-and-mortar” markets because e-commerce has small nominal price change costs (“menu costs”), small search costs, and small costs of monitoring competitors’ prices (Ellison and Ellison, 2005).⁶ Thus, online price quotes increase our chances of detecting the rapid repricing of CPU models in response to technological advances.

4.1. Frequency and size of price changes

Following previous studies, we determine the frequency of price adjustment as the proportion of non-zero price changes to the total number of price changes observed within our dataset (e.g., Bils and Klenow, 2004; Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008; Eichenbaum et al., 2011). Specifically, we consider a price change that is smaller than 0.1% as a zero-price change, that is, the price is treated as fixed if the deviation from the price in the previous month is very small.

Table 3 reports that the median of the frequency of regular price adjustments is 39.4% per month (that is, 39.4% of prices are changed from one month to the next) and the implied

⁵ CPU producers changed their life-cycle pricing strategy and stopped reducing the price of old models.

⁶ Because we do not have sales flag as in scanner data, we follow previous studies (e.g., Chahrour, 2011) to identify temporary sales with a “sales filter”, which is the V-shape or ^-shape in price changes. Specifically, we consider an increase or decrease in price as temporary price changes if the price returns to its previous price level within one month. The mean frequency of sales across CPUs is 1.88% and the median size of sales across CPUs is 2.28%. Because sales are not prevalent for CPUs, we report statistics only for “regular” prices (that is, prices without sales) to preserve space. Moments for posted prices (that is, prices with sales) are available upon request.

duration of a typical price spell is two months.⁷ These moments are broadly similar to what can be observed for a representative basket of goods sold online. However, in contrast to typical goods in e-commerce, CPUs have a more pronounced bias towards price decreases relative to price increases. Specifically, the frequency of price decreases is 22.4% per month for CPUs and 15.6% per month for a typical product sold online. Meanwhile, the corresponding figures for price increases are 15.2% and 13.9%. This pattern is consistent with a faster technical change in the CPU market which forces CPU sellers to mark down their prices more often.

We observe a similar pattern for the average absolute log price change ($|\Delta p_{ist}|$). The average change for all price adjustments is similar for CPUs (8.1%) and typical goods (10.9%) but there is a clear asymmetry in the size once we condition on the direction of price change. The average size of price decreases for CPUs is 9.2% while the average size of price increases is only 5.6%. In contrast, the size of price changes for typical goods is roughly symmetric: 11.8% for price increases and 10.3% for price decreases.

More generally, these results suggest that, relative to the conventional markets, the frequency of price changes is high and the size of typical price changes is small—facts consistent with easy/cheap nominal price adjustment—however there remains discernible price stickiness which can be rationalized by market/product characteristics. To further explore this conjecture, we regress pricing moments on six variables: (1) number of sellers that sell product i ; (2) quality of product i ; (3) the median price of product i ; (4) share of price points, which is the percentage of price quotes that end at 9, 95 or 99 for product i (e.g., \$199, \$349, \$495); (5) the stability of sellers for product i is the ratio of the number of sellers offering product i in a given month to the number of sellers ever selling this product in the quarter, which covers the given month; and (6) CPU producer, which is a dummy variable that equals 1 if it is an Intel CPU and equals zero otherwise. The first variable proxies for market competition (e.g., Ginsburgh and Michel, 1988; Martin, 1993). The second variable captures the characteristics of goods. The third is a proxy for the returns of buyers' searches (e.g., Head et al., 2010). The fourth reflects the level of inattention to prices when choosing between products (e.g., Knotek, 2011). The fifth variable (stability of sellers) is a proxy for the turnover of sellers (e.g., Gust, Leduc, and Vigfusson, 2010). The last variable aims to capture the special role of Intel in this CPU market.

⁷ If \bar{f}_i is the average frequency of price adjustment for CPU i , the mean implied duration is given by $\bar{d}_i = -[\log(1 - \bar{f}_i)]^{-1}$.

Table 4 reports the results of regression for the frequency and size of price adjustments as outcome variables. All variables are measured at the product level. For instance, the frequency of price changes at the product-level for a specific product is computed as follows: First, the frequency of price adjustments for each seller offering that product is calculated. Then, the data is collapsed to product-level by taking the raw average across sellers to use as a dependent variable and the regression is estimated with no weights.

We find that all explanatory variables have some predictive power. First, a market with more sellers should have a higher (lower) frequency of positive (negative) price changes, and a smaller size of adjustment. Second, the quality of goods in the market is positively associated with the degree of price flexibility (higher frequency and smaller size of price adjustments). This is because a high-quality product often has lower search costs due to higher advertising expenses and higher returns on searches for customers.⁸ We also find that a more powerful processor has price increases less often and price decreases more often. This can be explained by the observed quality premium, which provides more room for sellers to adjust prices. Third, for products with low- and moderate-prices, price changes occur more often and are larger in size when the median price across sellers of a product increases. This result is consistent with Head et al. (2010), documenting that higher returns on searches would put more pressure on the seller to adjust prices. Nevertheless, given the estimated nonlinearity, very expensive CPUs tend to have fewer and smaller price changes. Fourth, a product which has a high proportion of price points has a lower frequency of price changes, particularly positive price changes. This result is consistent with Levy et al. (2011), who find that prices ending with 9 have a lower frequency of price changes. Fifth, an increase in the degree of seller stability, which implies that it is more difficult for new sellers to enter the market, is associated with a decrease in the frequency of price changes, particularly, the frequency of negative price changes. Lastly, Intel processors tend to change prices more frequently than AMD products.

4.2. Price dispersion across sellers

Hedonic pricing of different vintages largely focuses on explaining variations in prices across vintages, rather than within-vintage price dispersion. However, price dispersion across sellers within a product can help us to directly measure potential “mispricing” and frictions in the

⁸ For the negative relationship between customer search and price stickiness level, see i.e., Ellison and Ellison (2009).

market because the quality of the product is held constant. Indeed, if there is dramatic within-product price dispersion, one may be less sanguine about using hedonic regressions to price properties of products since the market does not eliminate arbitrage opportunities.

4.2.1. Intra-month dispersion across sellers

We use several measures of within-CPU price dispersion, which capture different aspects of price variation: the coefficient of variation (CV; the ratio of standard deviation to the mean), standard deviation of the monthly log prices, the log difference of the average and the lowest price (the value of information (VI)), interquartile range (IQR), range (the log difference between the lowest and highest price), and gap (the difference between the two lowest log prices). We calculate the measure of price dispersion between sellers for an identified CPU in a month, then collapse our data to product level by taking the raw average over time.

One might argue that the observed price dispersion is caused by the differences in the shopping experience that customers have with different sellers (see e.g., Stigler, 1961). This difference is not likely to be significant when customers are shopping online as consumers only deal directly with a seller after completing the transaction. To further alleviate this potential problem, we purge “shopping experience” from prices using the following regression:

$$\log p_{ist} = \alpha_i + \gamma_s + \varepsilon_{ist} \quad (2)$$

where α_i and γ_s control for product and seller fixed effects, respectively. The fixed effects can capture the differences in reputation, delivery conditions, and return costs between CPU retailers. Thus, the dispersion of the residuals (ε_{ist}) gives us the price dispersion net of sellers’ heterogeneity in, for example, shipping costs and return policies, which are likely to remain unchanged over a short period of time (see e.g., Nakamura and Steinsson, 2008).

We report price dispersion for P_{ist} and ε_{ist} in Panel A of Table 5. CV (row 1) for P_{ist} is 0.22, thus suggesting considerable heterogeneity in prices, even for a fixed CPU. The standard deviation of log prices (row 2) is also relatively high at 0.23. Seller fixed effects account for around 20% of the variation in actual CPU prices in the U.S. online market (see row (7)) and the residual price dispersion is 0.18. The within-product dispersion for CPUs is generally similar to the within-product dispersion for goods typically sold online (column 2 of the Table).

4.2.2. Price dispersion across sellers over CPU lifetime

The price dispersion across sellers of a product may depend on the stage of the product lifetime. Dispersion is expected to be higher at the release time of a product and then gradually decreases as customers increase their awareness of the pricing strategy of sellers and sellers collect information about their competitors' prices. In a stationary setting (i.e., the set of products is fixed), lifetime price dispersion is not important. However, when there is a turnover of products, one may observe the dispersion of prices in a cross-section simply because we have a mix of "new" products that have large price dispersion and "old" products that have low price dispersion. If price dispersion declines over the lifetime of a product, one may expect that hedonic pricing of attributes works well because the law of one price may hold for established products.

To study this aspect of price dispersion, we calculate the average price dispersion across CPUs for month j after their introduction. Following Gorodnichenko et al. (2018), we identify the time of product introduction by taking the time that the product appeared in the data. CPUs which enter within the first quarter are excluded, as we cannot know whether the CPU was released previously or it returned after being temporarily unavailable. Following this, the measure of price dispersion over a CPU lifetime is computed as follows: We generate the time variable for each processor as the number of months since the month that the processor appeared in the dataset to replace the calendar months. Then we use the cross-sectional price dispersion of each processor and the new time variable to compute the average price dispersion across CPUs for j month after their release month.

We find little evidence of price convergence across sellers offering a CPU over the product lifetime. Figure 6 shows that, since the time that a product is released, price dispersion increases slowly over the first two years from around 0.14 to nearly 0.20. This dynamic is likely to be explained by evolution in the depth of the market, because the number of sellers stays relatively constant over a product's lifetime. Interestingly, price dispersion also increases over the lifetime of goods typically sold online but the profile of the increase is flatter.

4.2.3. Spatial vs. temporal price dispersion

Understanding the nature of price dispersion can show us how persistent mispricing can be, as well as how one can select outlets to correctly price product attributes. The existing literature

on price dispersion posits two main types of price dispersion: spatial and temporal. Spatial dispersion refers to the situation in which some sellers persistently charge higher prices for their product because they have, for example, a better location (see e.g., Baye et al., 2010). In contrast, temporal dispersion refers to the situation in which sellers do not consistently charge lower/higher prices for their product because over time buyers learn sellers' pricing behaviour and identify the seller who offers the best price (see e.g., Varian, 1980; Sheremirov, 2019). Appendix Table A4 provides examples for spatial and temporal price dispersion. The form of price dispersion affects whether or not one can use outlets for sampling prices interchangeably. Specifically, if price dispersion is temporal, the identity of the seller is less important because one can recover the average price in the market by following this seller. In contrast, if price dispersion is spatial, the identity of the seller is important for measuring the level of prices in the market.

Following the method developed by Lach (2002), we assign the price of a product offered by a seller ("price line") to a quartile of the price distribution across all sellers of that product in a given month, then analyze movements across quartiles of that price line over time. For example, the price of retailer s for CPU i in month t is P_{ist} and three cut-off points for product i in month t are $Q_{it}^{(1)}, Q_{it}^{(2)}, Q_{it}^{(3)}$. Then, a seller with the price for product i such that $Q_{it}^{(1)} < P_{ist} < Q_{it}^{(2)}$ is in the second quartile of the cross-sectional distribution in month t , while a seller with the price $P_{ist} > Q_{it}^{(3)}$ is in the fourth quartile (i.e., the price for CPU i is higher than 75% of prices offered by all sellers for CPU i in month t). Following this step, we construct the fraction of time that P_{ist} spends in each quartile and the average fractions across CPUs.

If sellers do not consistently set their price lower/higher than prices of other sellers and instead charge high prices at times and low prices at other times (i.e., temporal price dispersion), the probability of a given price line remaining in each quartile of the price distribution should be similar and near to 25%. If sellers consistently charge lower or higher prices (i.e., spatial dispersion), P_{ist} would spend more time in one of the quartiles. We find evidence consistent with spatial price dispersion in our data.⁹ Rows (12)-(15) of Panel B in Table 5 show that 20.9% of price lines spend more than 95% of the time in one quartile of the cross-sectional distribution. Furthermore, 9.8% of price lines almost always stay in the lowest quartile and 7.3% of price lines almost always stay in the highest quartile. In addition, rows (8)-(11) in the

⁹ Because the results for "raw" prices and residual prices (i.e., ε_{ist} , prices net of seller fixed effects) are similar, we only reported results using "raw" prices.

table show that from 37.2% to 50.1% of price lines spend almost no time in a given quartile. For example, 42.6% of price lines rarely stay in the cheapest quartile and 50.1% of price lines rarely stay in the most expensive quartile.¹⁰ These results are broadly similar to the corresponding moments for a representative basket of goods sold online (column 2).

The last row in Table 5 presents another useful metric: the average standard deviation across price lines of the fractions of time spent in a particular quartile. The average standard deviation equal to 0 implies perfect temporal price dispersion, while the value of $\sqrt{3}/4$ (≈ 0.43) implies perfect spatial price dispersion. In fact, the average standard deviation is 0.29, which is closer to 0.43. For comparison, this metric stands at 0.31 for a typical product sold online.

4.2.4. Predictors of within-product price dispersion

Economists generally rationalize price dispersion with three forces: search costs, price stickiness, and price discrimination.¹¹ To assess the role of these forces to explain within-product price dispersion for CPUs, we estimate the following regression:

$$\bar{\sigma}_i = \beta_0 + \beta_1 \times \log(Q_i) + \gamma \times \bar{X}_i + error_i \quad (3)$$

where $\bar{\sigma}_i$ is the average of monthly within-CPU standard deviation of the log prices for CPU i , Q_i is the performance score of CPU i , \bar{X} is a vector of market/product characteristics of CPU i such as the average number of sellers for CPU i , share of price points (proportion of pricing points ending with 9 or 95-99), stability of sellers' pool, Intel indicator variable, and price stickiness measures (frequency and size of price changes).

Table 6 reports the results for the regression of standard deviation of log price in column (1). We find that CPU performance is negatively correlated with price dispersion. On the other hand, median price, frequency and size of regular price adjustments are positively associated with price dispersion. The results are similar between the regression of posted prices and residual prices after removing seller fixed effects (column 2).

4.3. Taking stock

¹⁰ Furthermore, Appendix Figure A1 plots the distribution of these fractions over observed price lines to provide a clear picture of the existence of the spatial dispersion.

¹¹ See Gorodnichenko, Sheremirov and Talavera (2018) for a discussion.

Our results suggest that price setting in the online market for CPUs is consistent with discernible frictions due to price stickiness and market structure. While there are clear signs of fast technical change in CPU prices (e.g., price decreases are larger and more frequent than price increases), CPU prices tend to have spells as long as spells for other goods sold online, despite the fact that CPUs have faster technical change than many other goods sold online. Within-product price dispersion tends to have properties largely similar to those for a typical product in e-commerce: dispersion is large and “spatial”. Furthermore, price dispersion and stickiness can be explained by market structure (e.g., number of sellers) and product characteristics (e.g., product prices which influence incentives to search for better prices). These facts point to two important observations. First, market/pricing frictions can materially affect price differentials *across* CPUs. Second, the selection of outlets may be important for the quality of hedonic calculations, that is, price quotes from some outlets may result in overstating (or understating) the value of attributes and hence distort estimates of inflation in prices adjusted for quality differences.

5. Technological changes and CPU pricing

Previous sections document that CPU pricing depends not only on CPU performance but also on a variety of market fundamentals.¹² Furthermore, the pricing of CPU performance can vary with market/product characteristics. In other words, we may be unable to use CPU-A’s price to benchmark CPU-B’s price by performance alone because the unit of performance may be priced differently for CPUs A and B due to differences in market power, price stickiness, etc. As a result, when CPU prices are compared, we need to hold constant, not only CPU performance, but also market structure.

To build intuition for this point, let us suppose we need to compare an exiting CPU model (call it CPU A) in period t to an entering CPU model (call it CPU B) in period $t + 1$. The prices for A and B may be different because B is more powerful than A or because the seller (or producer) of A has a different market power relative to the seller of B. To have the correct measure of inflation, we need to take two actions. First, we need to construct the price that the seller of A would have charged if it sold B. Once this price is constructed we can adjust that price for quality differences between A and B. The hedonic regressions adjust for the second step of this

¹² With a slight abuse of terminology, we equate markets with products. In other words, the set of sellers for a given product is treated as a market.

process, but they do not adjust for the first step of the process because they implicitly assume that sellers are interchangeable (e.g., the law of one price holds).

In this section, we investigate how between-CPU price dispersion can be rationalized by differences in CPU performance, as well as market fundamentals. To this end, we first examine the evolution of CPU prices in response to the introduction of new CPU models. We then attempt to predict average CPU prices (across sellers) in a regression context. Finally, we use hedonic regressions to estimate inflation for CPUs, holding performance and market fundamentals constant.

5.1. Pricing moments and technology shocks

As discussed above, there were several breakthroughs in CPU technology: January 2011 when Intel introduced *Sandy Bridge* CPUs (the second generation of Intel Core processors) using 32 nm technology nodes; October 2011 when AMD released *Bulldozer* family 15h (the first generation) using its 32 nm process technology; April 2012 when *Ivy Bridge* processors (the third generation of the Intel Core processors; the 22 nm process) were released, September 2012 when AMD released *Piledriver* family 15h (the second generation, which had some improvements but still used the same module 32-nm design with *Bulldozer*). One may expect major re-pricing of CPUs around the time when new generations become available. Indeed, if a new generation of processors improves performance by 50% on average, prices for existing models could decline by as much as 50%.

Panel A and B of Figure 7 illustrate the monthly frequency and average absolute size of positive and negative price adjustments throughout these shocks. Panel A of the figure uses the full sample, while Panel B uses the data that is restricted to the older generations (namely *AMD K10*, *Intel Core 2*, and *Intel Nehalem*). The release of a new CPU generation generally decreases the frequency and size of positive price changes and increases the frequency and size of negative price changes (Panel A). These effects are more visible for older CPUs (Panel B). However, it is clear that these price changes fall short of a major (50% or so) repricing for existing CPUs. Furthermore, price adjustment is gradual in the data, while theory predicts instantaneous adjustment.

Panel C of Figure 7 paints a similar picture: the introduction of new technologies does not lead to large, rapid changes in the distribution of prices for existing CPUs. For example, the

dispersion of CPU prices in a given month stays roughly constant, irrespective of whether we look at the total price variation (i.e., between-product and within-product variation; black line in the figures) or between-CPU price variation (blue line) or whether we use raw prices (solid line) or quality-adjusted prices (dashed line).

5.2. Predictors of between-product price variation

Hedonic pricing implicitly assumes that the only systematic source of price differences across goods (“vintages”) are attributes that are valuable to consumers (e.g., processor clock speed, energy consumption, etc.). Our analysis of within-product price variation suggests that market/product characteristics can account for some part of price differentials across sellers for a given product, that is, if the quality of the product is held fixed. These findings suggest that pricing of product attributes can also depend on market fundamentals. If this is indeed the case, then hedonic pricing may provide a distorted assessment of what portion of price variation is due to changes in attributes and what portion is purely due to inflation. Intuitively, if attributes are priced differently across two markets (CPUs), we should be cautious when using one of the markets to price goods in the other market. Specifically, to compute inflation, we should hold constant, not only product attributes, but also market fundamentals (that is, the pricing scheme for the attributes).

To explore the quantitative importance of this point, we estimate the following regression:

$$\bar{p}_{it} = \beta_0 + \beta_1 \times \log(Q_i) + \boldsymbol{\gamma} \times \mathbf{X}_{it} + \boldsymbol{\psi} \times \mathbf{X}_{it} \times \log(Q_i) + error_{it} \quad (3)$$

where \bar{p}_{it} is the average price (across sellers) of CPU i in month t , Q_i is the performance score of CPU i , \mathbf{X} is a vector of market/product characteristics such as a CPU’s market concentration, average price (proxy for returns on search), price stickiness, stability of sellers, share of price points and product brand. Coefficient β_1 measures the sensitivity of CPU prices to its performance. Coefficient vector $\boldsymbol{\gamma}$ may be interpreted as measuring markups and other distortions in the market for a given CPU. Coefficient vector $\boldsymbol{\psi}$ captures differences in pricing of the relevant attribute (here the performance scores) across markets (i.e. how a given market characteristic alters the price of the attribute). Standard hedonic regressions impose $\boldsymbol{\psi} = 0$ and often $\boldsymbol{\gamma} = 0$. Note that while the quality (performance) of a CPU is fixed when the CPU is introduced, market fundamentals can evolve as we can have, for example, the entry or exit of sellers.

We find (column (1) of Table 7) that the elasticity of CPU price with respect to its performance score is 0.5 and performance alone can account for approximately 30% of the variation in prices across CPU models. Market characteristics (column (2)) are also statistically significant predictors of CPU prices, but these predictors can account for only 6.3% of price variation and the interpretation of estimated coefficients can be counterintuitive (e.g., as the number of sellers increases (more competition), CPU prices increase), which underscores the predictive rather than causal nature of specification (3). Once we condition on CPU quality and market characteristics (column (3)), the counterintuitive estimates are not an issue. For example, increased competition (more sellers, less stable pool of sellers) is associated with lower prices. Notice that the elasticity of CPU prices with respect to performance increases to 0.6. It can also be seen that Intel appears to charge a 48% premium over AMD. The R^2 of the fitted model increases to 0.39.

Finally, column (4) reports the results when we allow for the interaction of CPU performance and market fundamentals. All interactions are statistically significant and this further improves the fit of the model (R^2 increases to 0.41). The economic intuition underlying the estimated coefficients is also potentially of interest. For example, the positive estimate of the interaction of performance and Intel dummy suggests that Intel can charge higher prices for high-performance CPUs, i.e., the segment of the market where Intel likely faces less competition from AMD. These results suggest that the comparison of prices for different technologies should involve, not only physical properties of the compared technologies, but also the characteristics of the markets in which these technologies are priced.

5.3. Adjustment of prices following technology shocks

The previous section documents that product attributes and market fundamentals are important for pricing CPUs. However, specification (3) does not allow us to assess how quickly CPUs are repriced in response to the entry of new models or the exit of old models. One may expect that an entry for a new, more powerful CPU should lead to price cuts for existing CPU models, but how quickly these cuts happen depends on a number of elements, such as how competitive the markets are and how sticky prices are. To formally explore this insight, we estimate the following specification:

$$m_{it} = \alpha_i + \xi \times X_{it} + \beta_1 \times Upgrade_{it} + \beta_2 \times Downgrade_{it}$$

$$+\psi_1 \times Upgrade_{it} \times \mathbf{X}_{it} + \psi_2 \times Downgrade_{it} \times \mathbf{X}_{it} + error_{it} \quad (4)$$

where m_{it} is a pricing moment for CPU i in month t (frequency of price changes across sellers of CPU i in month t , size of price changes across sellers of CPU i in month t), \mathbf{X} is a vector of market characteristics and α_i are CPU fixed effects. Indicator variable “Upgrade” equals one if the product i jumps to a higher performance quartile at time t compared to time $t - 1$ and equals zero otherwise. For instance, when the number of products entering the market with lower performance scores than product i and/or the number of products exiting the market with better quality than the product i are sufficient, the product i will jump to a higher quality quartile. Indicator variable “Downgrade” equals one if the product i falls to a lower performance quartile at time t compared to the previous period and zero otherwise. We report estimated coefficients in Table 8.¹³

Consistent with time-series evidence presented in section 5.1, we find that upgrading or downgrading a given CPU relative to other CPUs available in the market changes the probability and size of a price change only marginally. For example, coefficients on Upgrade/Downgrade separately and coefficients on interactions with product/market characteristics are generally not significant statistically.

The estimated insensitivity of CPU prices to changes in relative performance position is consistent with several explanations. First, prices may be very sticky and so they exhibit little response to the introduction of new models. Given how quickly online prices change, this hypothesis alone should not be able to account for the observed insensitivity.

Second, new models are priced in such a way that the price of a new CPU does not disrupt the pricing of existing models. For example, if the price per performance unit is fixed, a more powerful model can command a higher absolute price, but it would not mean that existing models are mispriced because they charge the same quality-adjusted price. As discussed above, even when we compare CPU prices per performance unit, there remains considerable price dispersion and performance alone accounts for only 30% of the observed price variation across CPUs. Thus, this hypothesis is also unlikely to be the exclusive explanation.

Third, CPU manufacturers may segment CPU markets. For example, two average CPUs may not be combined to produce one powerful CPU. In a similar spirit, new CPUs may be incompatible with other equipment (e.g., motherboards) thus limiting substitutability of CPUs.

¹³ Results for alternative specifications are reported in Appendix Tables 2 and 3.

AMD and Intel processors are not fully compatible either. Furthermore, chip manufactures may require retail sellers to enforce differential pricings so that AMD/Intel can maximize profits. Such market segmentation can thus limit arbitrage and create potential mispricing (i.e., quality-adjusted prices will not be the same across CPUs). Consistent with this hypothesis, we observe that Intel charges systematically higher prices.

Finally, when a new model (or technology) is released, there could be some uncertainty about the quality of the new release and sellers may be reluctant to re-price “old and tried” models before new releases are thoroughly benchmarked. This demonstrates why benchmarking firms such as Passmark have experienced a high level of success. While we cannot test this hypothesis directly, we conjecture that uncertainty should be dispelled relatively quickly and therefore this hypothesis is unlikely to be a major factor.

In any case, our evidence is consistent with considerable pricing frictions that can limit the consistent pricing of CPU attributes across markets, thus posing a potential threat to hedonic pricing of CPUs.

6. Inflation in CPU prices

Quality improvement has long been a major challenge for measuring inflation, especially for goods with rapid technological advancements. Hedonic regression is a popular approach for ensuring the correct prices for upgrades of product attribute trends (see e.g., Gordon and Griliches, 1997; Shiratsuka, 1999; Bils, 2009). However, this approach ignores the fact that prices are also affected by other factors, such as market and product characteristics.¹⁴

We follow the common approach in the existing literature (see e.g., Byrne et al., 2018; Aizcorbe et al., 2020) and run the hedonic regression for each pair of consecutive years in our sample. More specifically, we include product characteristics and market fundamentals as control variables. Our focus is on the quality-adjusted price movement which is reflected by the coefficient of the year dummy for each time period. We show that hedonic adjustment may be incomplete if pricing of product attributes varies across products. Thus, one should control, not only for product attributes, but also for market fundamentals.

¹⁴ For a related issue in the context of international economics, see e.g., Mallick and Marques (2016) for the impact of income level in importers’ and exporters’ countries as well as product quality on export prices.

To assess how market fundamentals can affect measures of CPU price inflation, we estimate the following specification for each pair of consecutive years (i.e., $\tau = \{2009-2010, 2010-2011, 2011-2012\}$) in our sample:

$$\bar{p}_{it} = \beta_0^{(\tau)} + \beta_1^{(\tau)} \times D_{2t} + \beta_2^{(\tau)} \times \log Q_i + \gamma^{(\tau)} \times \mathbf{X}_{it} + \psi^{(\tau)} \times \log Q_i \times \mathbf{X}_{it} + error_{it} \quad (5)$$

where \bar{p}_{it} is average price of CPU i observed in month t , D_{2t} is a dummy variable that equals 1 if the observation is in the 2nd year of the 2-year period. Vector \mathbf{X} includes market fundamentals such as S_{it} (the median price across sellers), SPP_{it} (number of sellers) and $Stab_{it}$ (seller stability), and $Intel_i$ (an indicator variable equal to 1 if processor i is produced by Intel). Note that we allow coefficients to vary for each two-year sample τ . We run the regression at the product-month level. Table 9 reports $\hat{\beta}_1^{(\tau)}$ for various τ which measure inflation in CPU prices from one year to the next.

When we do not control for product quality or market fundamentals (column 1), we find that inflation was $\approx 25\%$ in 2010 followed by mild deflation in 2011 and 2012. However, rising prices could largely reflect rapid improvement in CPU performance. Indeed, after controlling for performance scores (column 2), the price of CPUs (i.e., quality-adjusted prices) fall by just over 4% in 2010 and by another 11.4% in 2011, but increased by 5.2% in 2012. The cumulative deflation over the sample period is 10.4%. Controlling for changes in the market structure alone (column 3) is not enough to capture quality improvements and implied inflation could be in double digits. Controlling for product quality and market fundamentals, column 4 demonstrates inflation estimates similar to those when we control only for quality (column 2) but the magnitudes and timing of inflation are different. For example, measured inflation in 2010 is 3.7% while the timing of deflation is 10.6% in 2011 and 4.3% in 2012. As a result, the cumulative deflation is only 11.2%. Finally, when we allow for the interaction of market fundamentals and product quality (that is, for differences in how CPU performance scores are priced across markets), we find that cumulative deflation is 16.2% which is considerably larger than the deflation estimated with quality adjustment alone. It should also be noted that the profile of inflation dynamics is steeper than in column (2).

These estimates suggest that conventional hedonic regressions may provide distorted estimates of inflation, because these regressions do not control for changes in market fundamentals or take into account how pricing of product attributes can vary across markets.

7. Conclusion

While the quality of goods is often used to explain price differentials across goods, the measurement, as well as the pricing of quality, is a major challenge. This challenge is particularly apparent for goods which undergo rapid technological advances, which can cloud even high-frequency variation in inflation. Statistical agencies have to wrestle with many thorny issues in this context, but hedonic pricing has proven to be a popular approach for separating inflation from quality improvements. We advance this extremely important tool by examining how market fundamentals and pricing of product attributes (“quality”) can interact and thus potentially affect measured inflation.

Specifically, we couple data on CPU performance (a direct, integral measure of product quality) with a large dataset of CPU price quotes to quantify the role of the quality of goods and market fundamentals in the observed price variation. We document a number of factors consistent with price rigidities and mispricing (i.e., violations of the law of one price for raw prices or quality-adjusted prices) in the retail CPU market when we focus on measures of price stickiness and within-product price dispersion. Furthermore, we can relate the observed frictions (price stickiness, dispersion of prices) to market fundamentals (e.g., the number of sellers which proxies market concentration). This evidence suggests that one may not expect retailers to necessarily provide competitive pricing of CPUs and, in general, the source of price quotes (i.e., the identity of a seller) can be consequential for estimates of inflation.

Building on this insight, we demonstrate that prices for different CPUs depend, not only on CPU performance, but also on market fundamentals. Specifically, pricing of CPU attributes varies across products, thus making the comparison of prices across CPUs more nuanced than commonly postulated by hedonic pricing. We show that failure to account for market fundamentals can result in materially understated deflation of CPU prices in our sample period.

There is much work ahead in this area. For example, future research can gain further insights into the interaction of quality and market fundamentals for other goods by studying a broader spectrum of industries. Longer time series can also provide more variation in order to better isolate the effect of technological advances and inflation on prices.

References

- Aizcorbe, A., Byrne, D. M., & Sichel, D. E. (2020). Getting smart about phones: new price indexes and the allocation of spending between devices and services plans in Personal Consumption Expenditures. In *Measuring Economic Growth and Productivity* (pp. 387–411). Elsevier.
- Baye, M. R., Morgan, J., & Scholten, P. (2010). Information, Search, and Price Dispersion. In Terrence Hendershott (Ed.), *Economics and Information Systems, Handbooks in Information Systems* (pp. 323–376).
- Bils, M. (2009). Do higher prices for new goods reflect quality growth or inflation? *Quarterly Journal of Economics*, 124(2), 637–675.
- Bils, M., & Klenow, P. J. (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5), 947–985.
- Brynjolfsson, E., & Hitt, L. M. (2003). Computing Productivity: Firm-Level Evidence. *Review of Economics and Statistics*, 85(4), 793–808.
- Byrne, D. M., Corrado, C., & Sichel, D. (2020). The Rise of Cloud Computing: Minding your Ps, Qs and Ks. In *Measuring and Accounting for Innovation in the Twenty-First Century*. University of Chicago Press.
- Byrne, D. M., Kovak, B. K., & Michaels, R. (2017). Quality-Adjusted Price Measurement: A New Approach with Evidence from Semiconductors. *The Review of Economics and Statistics*, 99(2), 330–342.
- Byrne, D. M., Oliner, S. D., & Sichel, D. E. (2018). How Fast are Semiconductor Prices Falling? *Review of Income and Wealth*, 64(3), 679–702.
- Chahrour, R. A. (2011). Sales and price spikes in retail scanner data. *Economics Letters*, 110(2), 143–146.
- Combris, P., Lecocq, S., & Visser, M. (1997). Estimation of a Hedonic Price Equation for Bordeaux Wine: Does Quality Matter? *The Economic Journal*, 107(441), 390–402.
- Eichenbaum, M., Jaimovich, N., & Rebelo, S. (2011). Reference prices, costs, and nominal rigidities. *American Economic Review*, 101(1), 234–262.
- Ellison, G. D., & Ellison, S. F. (2009). Search, Obfuscation, and Price Elasticities on the

- Internet. *Econometrica*, 77(2), 427–452.
- Ellison, G., & Ellison, S. F. (2005). Lessons About Markets from the Internet. *Journal of Economic Perspectives*, 19(2), 139–158.
- Flamm, K. (2019). Measuring Moore’s Law: Evidence from Price, Cost, and Quality Indexes. In *Measuring and Accounting for Innovation in the Twenty-First Century*. University of Chicago Press.
- Ginsburgh, V., & Michel, P. (1988). Adjustment Costs, Concentration and Price Behaviour. *The Journal of Industrial Economics*, 36(4), 477–481.
- Goettler, R. L., & Gordon, B. R. (2011). Does AMD spur intel to innovate more? *Journal of Political Economy*, 119(6), 1141–1200.
- Gordon, R. J., & Griliches, Z. (1997). Quality Change and New Products. *American Economic Review*, 87(2), 84–88.
- Gorodnichenko, Y., Sheremirov, V., & Talavera, O. (2018). Price setting in online markets: Does it click? *Journal of the European Economic Association*, 16(6), 1764–1811.
- 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.
- Gust, C., Leduc, S., & Vigfusson, R. (2010). Trade integration, competition, and the decline in exchange-rate pass-through. *Journal of Monetary Economics*, 57(3), 309–324.
- Head, A., Kumar, A., & Lapham, B. (2010). Market power, price adjustment, and inflation. *International Economic Review*, 51(1), 73–98.
- Hoefflinger, B. (2011). ITRS: The International Technology Roadmap for Semiconductors. In *Chips 2020* (pp. 161–174). Springer, Berlin, Heidelberg.
- Jorgenson, D. W., Stiroh, K. J., Gordon, R. J., & Sichel, D. E. (2000). Raising the Speed Limit: U.S. Economic Growth in the Information Age. *Brookings Papers on Economic Activity*, 1, 125–235.
- Kaplan, G., Menzio, G., Rudanko, L., & Trachter, N. (2019). Relative Price Dispersion: Evidence and Theory. *American Economic Journal: Microeconomics*, 11(3), 68–124.
- Klenow, P. J., & Kryvtsov, O. (2008). State-Dependent or Time-Dependent Pricing: Does It

- Matter for Recent U.S. Inflation? *Quarterly Journal of Economics*, 123(3), 863–904.
- Klenow, P. J., & Malin, B. A. (2010). Microeconomic Evidence on Price-Setting. In *Handbook of Monetary Economics* (Vol. 3, Issue C, pp. 231–284).
- Knotek, E. S. (2011). Convenient Prices and Price Rigidity: Cross-Sectional Evidence. *Review of Economics and Statistics*, 93(3), 1076–1086.
- Kryvtsov, O. (2016). Is there a quality bias in the Canadian CPI? Evidence from microdata. *Canadian Journal of Economics*, 49(4).
- Lach, S. (2002). Existence and persistence of price dispersion: An empirical analysis. *Review of Economics and Statistics*, 84(3), 433–444.
- Levy, D., Lee, D., Chen, H. A., Kauffman, R. J., & Bergen, M. (2011). Price Points And Price Rigidity. *Review of Economics and Statistics*, 93(4), 1417–1431.
- Lünnemann, P., & Wint, L. (2011). Price Stickiness in the US and Europe Revisited: Evidence from Internet Prices. *Oxford Bulletin of Economics and Statistics*, 73(5), 593–621.
- Mallick, S., & Marques, H. (2016). Does quality differentiation matter in exporters' pricing behaviour? Comparing China and India. *China Economic Review*, 40, 71–90.
- Martin, C. (1993). Price adjustment and market structure. *Economics Letters*, 41(2), 139–143.
- Nakamura, E., & Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4), 1415–1464.
- Oliner, S. D., & Sichel, D. E. (2000). The resurgence of growth in the late 1990s: Is information technology the story? *Journal of Economic Perspectives*, 14(4), 3–22.
- Pakes, A. (2003). A Reconsideration of Hedonic Price Indexes with an Application to PC's. *American Economic Review*, 93(5), 1578–1596.
- Sheremirov, V. (2019). Price dispersion and inflation: New facts and theoretical implications. *Journal of Monetary Economics*, March 2019.
- Shiratsuka, S. (1999). Measurement errors and quality-adjustment methodology: lessons from the Japanese CPI. *Economic Perspectives*, Q II, 2–13.
- Solow, R. M. (1960). Investment and Technical Progress. In K. J. Arrow, S. Karlin, & P.

Suppes (Eds.), *Mathematical models in the social sciences*, 1959 (pp. 89–104).

California: Stanford University Press.

Stigler, G. J. (1961). The Economics of Information. *Journal of Political Economy*, 69(3), 213–225.

Varian, H. (1980). A Model of Sales. *American Economic Review*, 70(4), 651–659.

Figure 1. CPU generations and release dates.

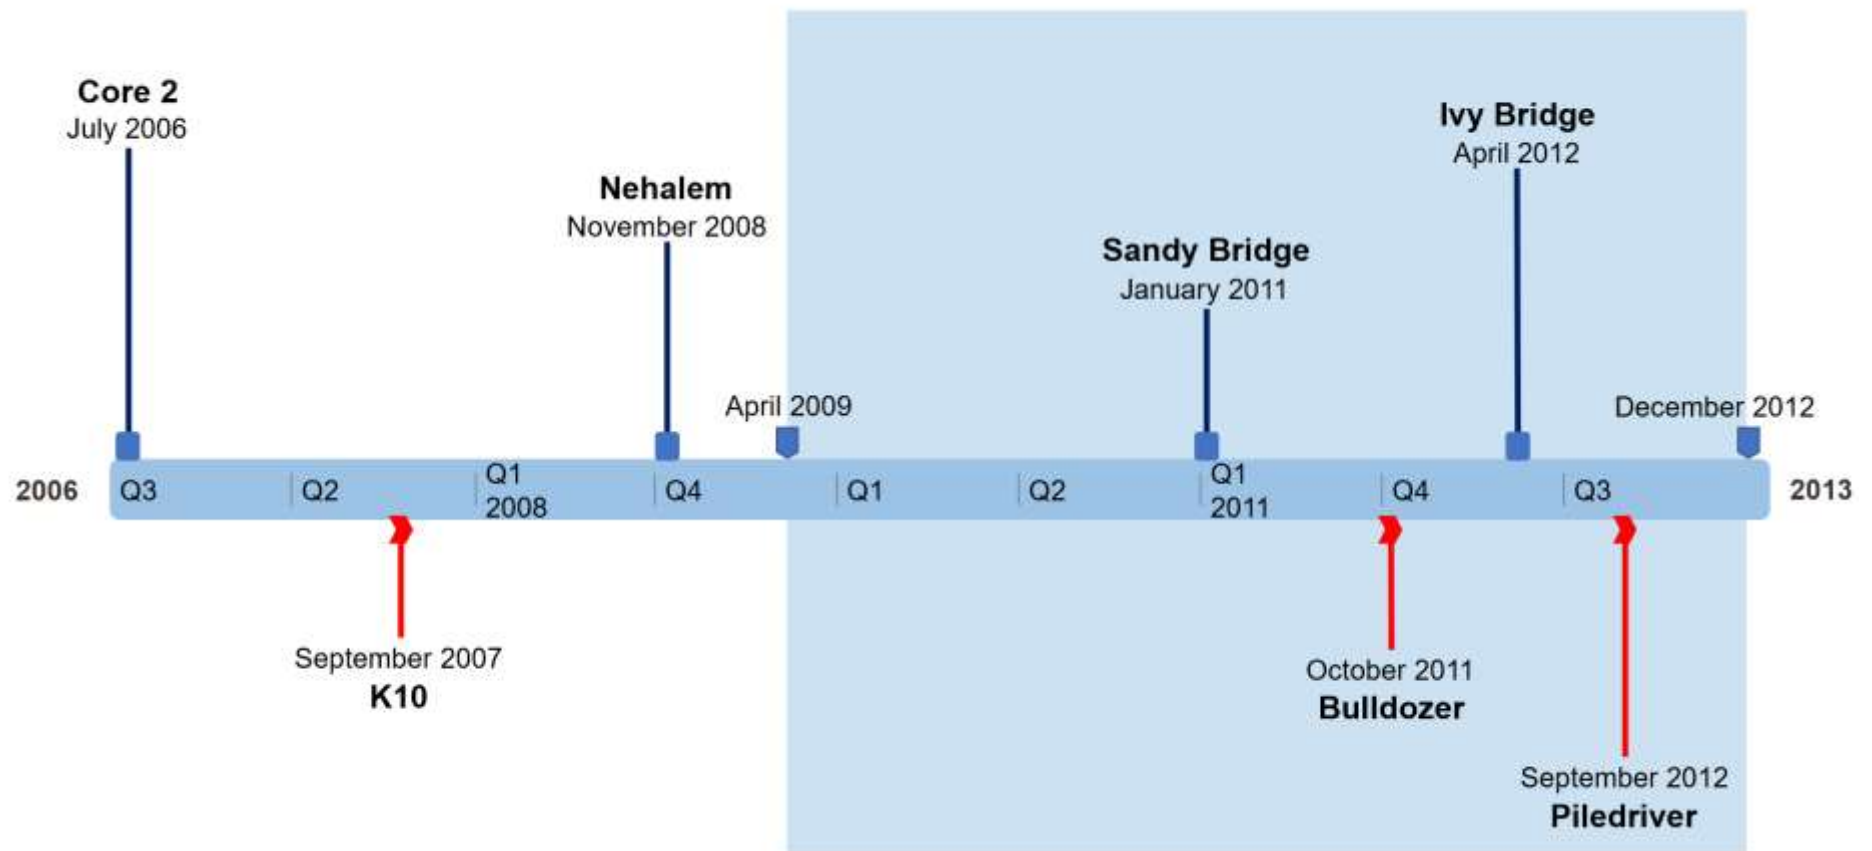
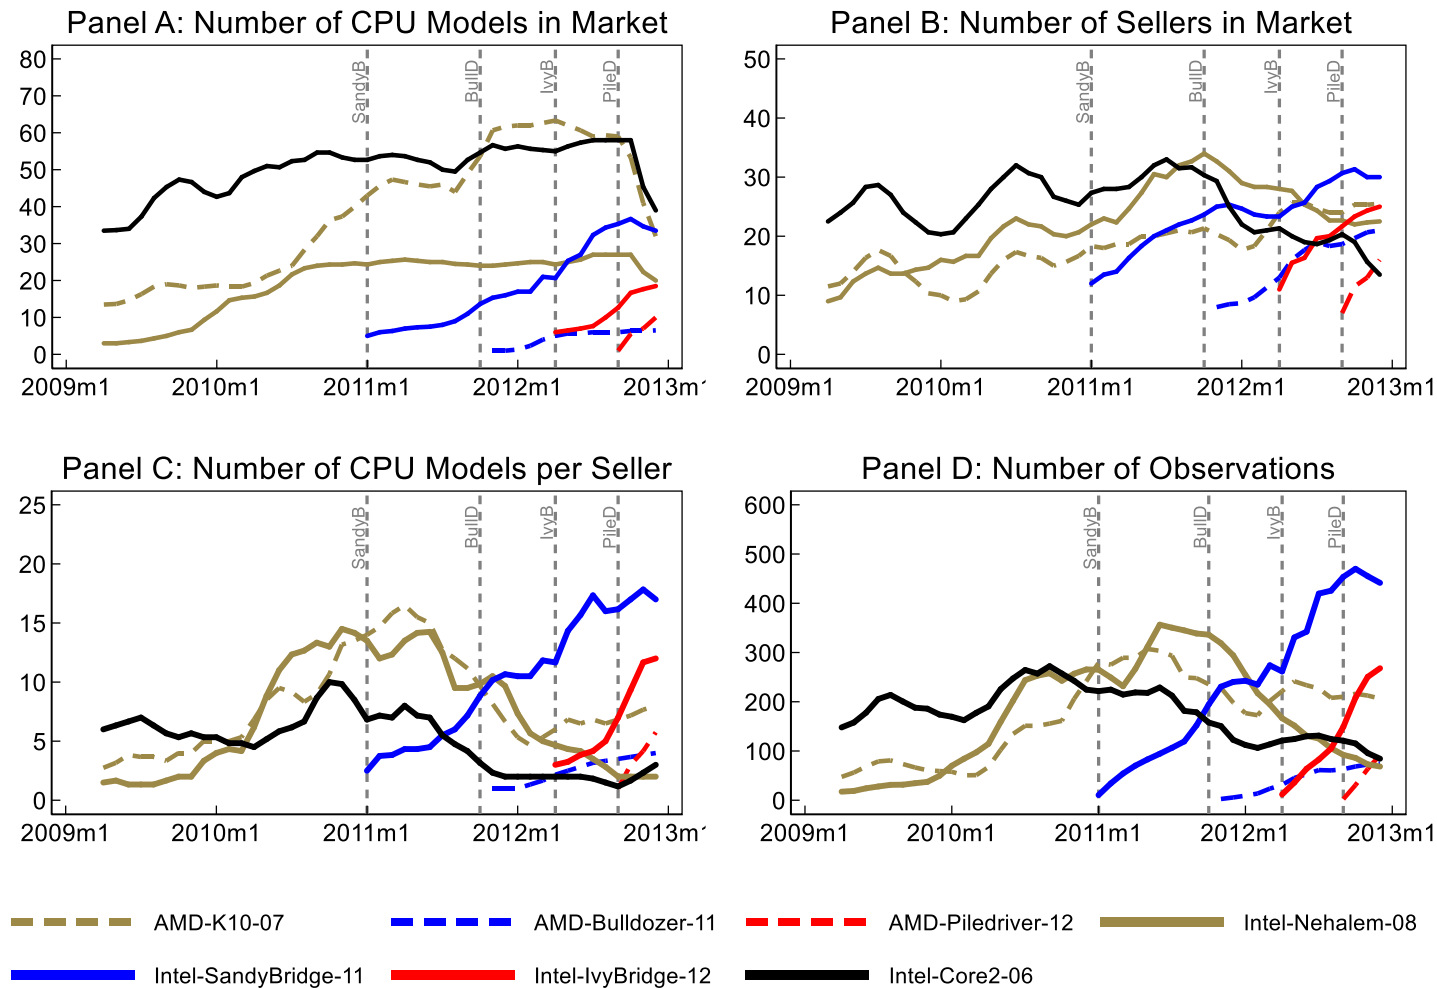
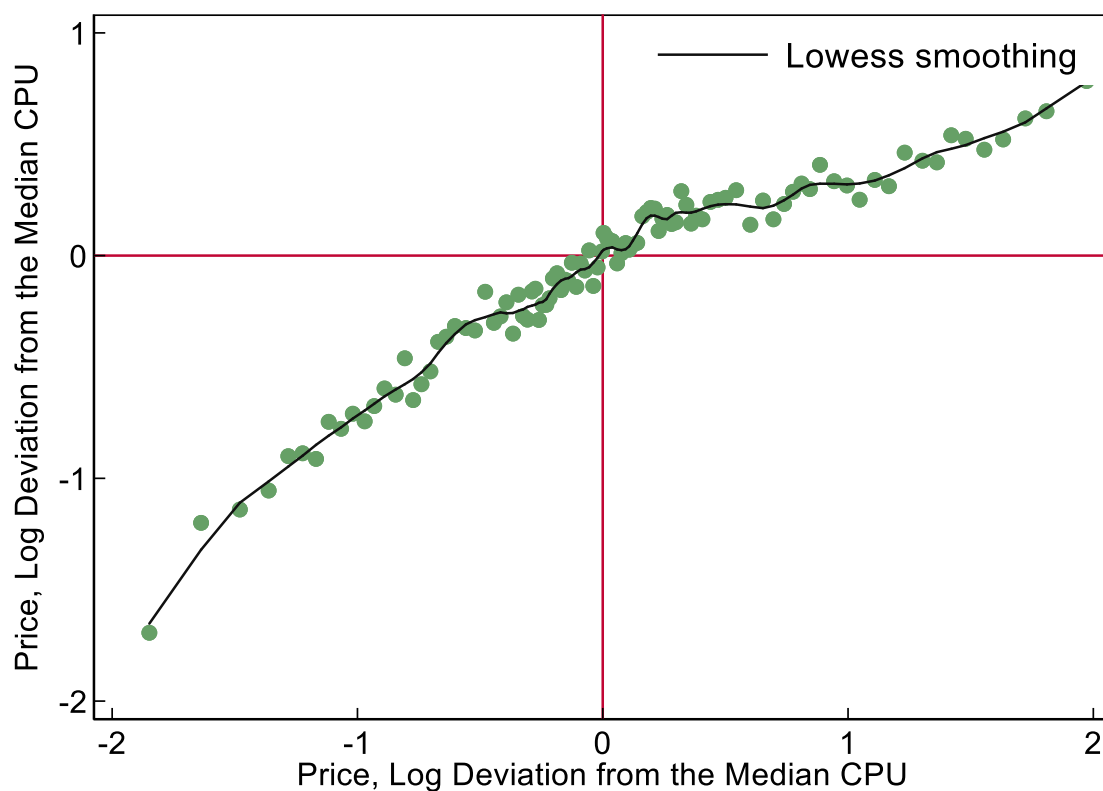


Figure 2. CPU Market Composition.



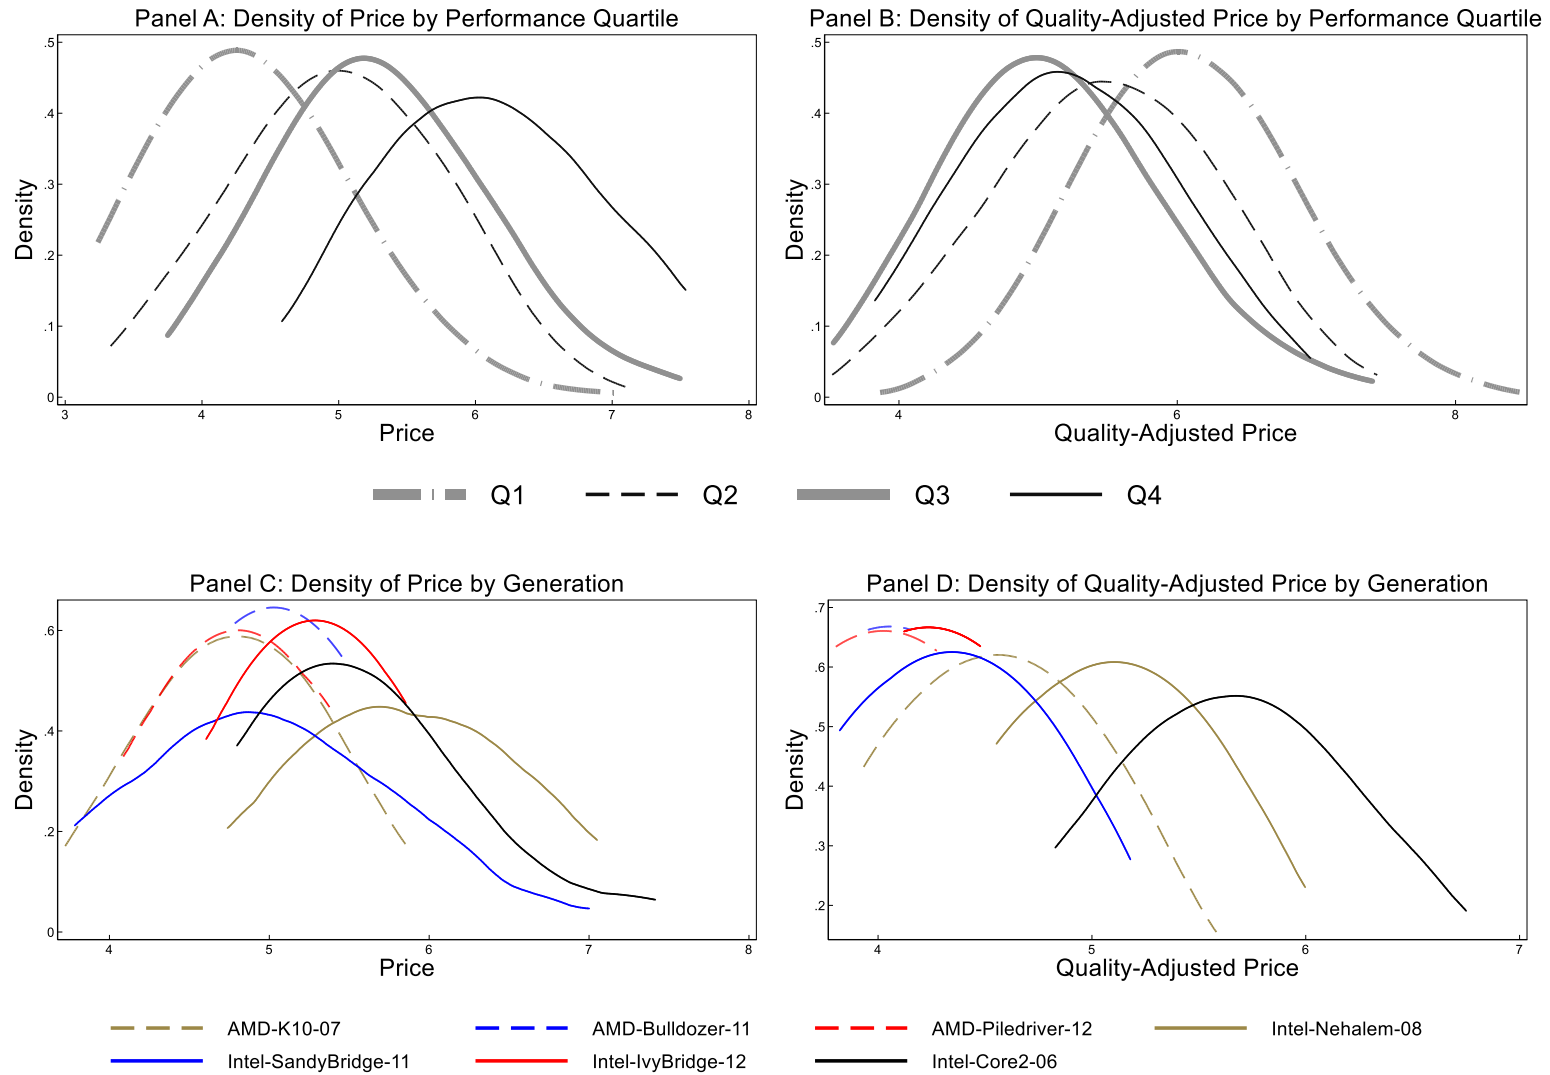
Note: In panel A-D, each line shows the monthly number of CPU models, number of sellers, number of CPUs models per seller, and number of price quotes of a CPU generation. In all panel, each grey dashed vertical line marks the release month of the corresponding CPU generation. The label of each line includes information of producer, name, and released year of the CPU generation. For example, *Intel-Nehalem-08* is *Nehalem* CPU generation that is produced by Intel and was first released in 2008.

Figure 3. Prices and Performances, Log Deviation from the Median-CPU.



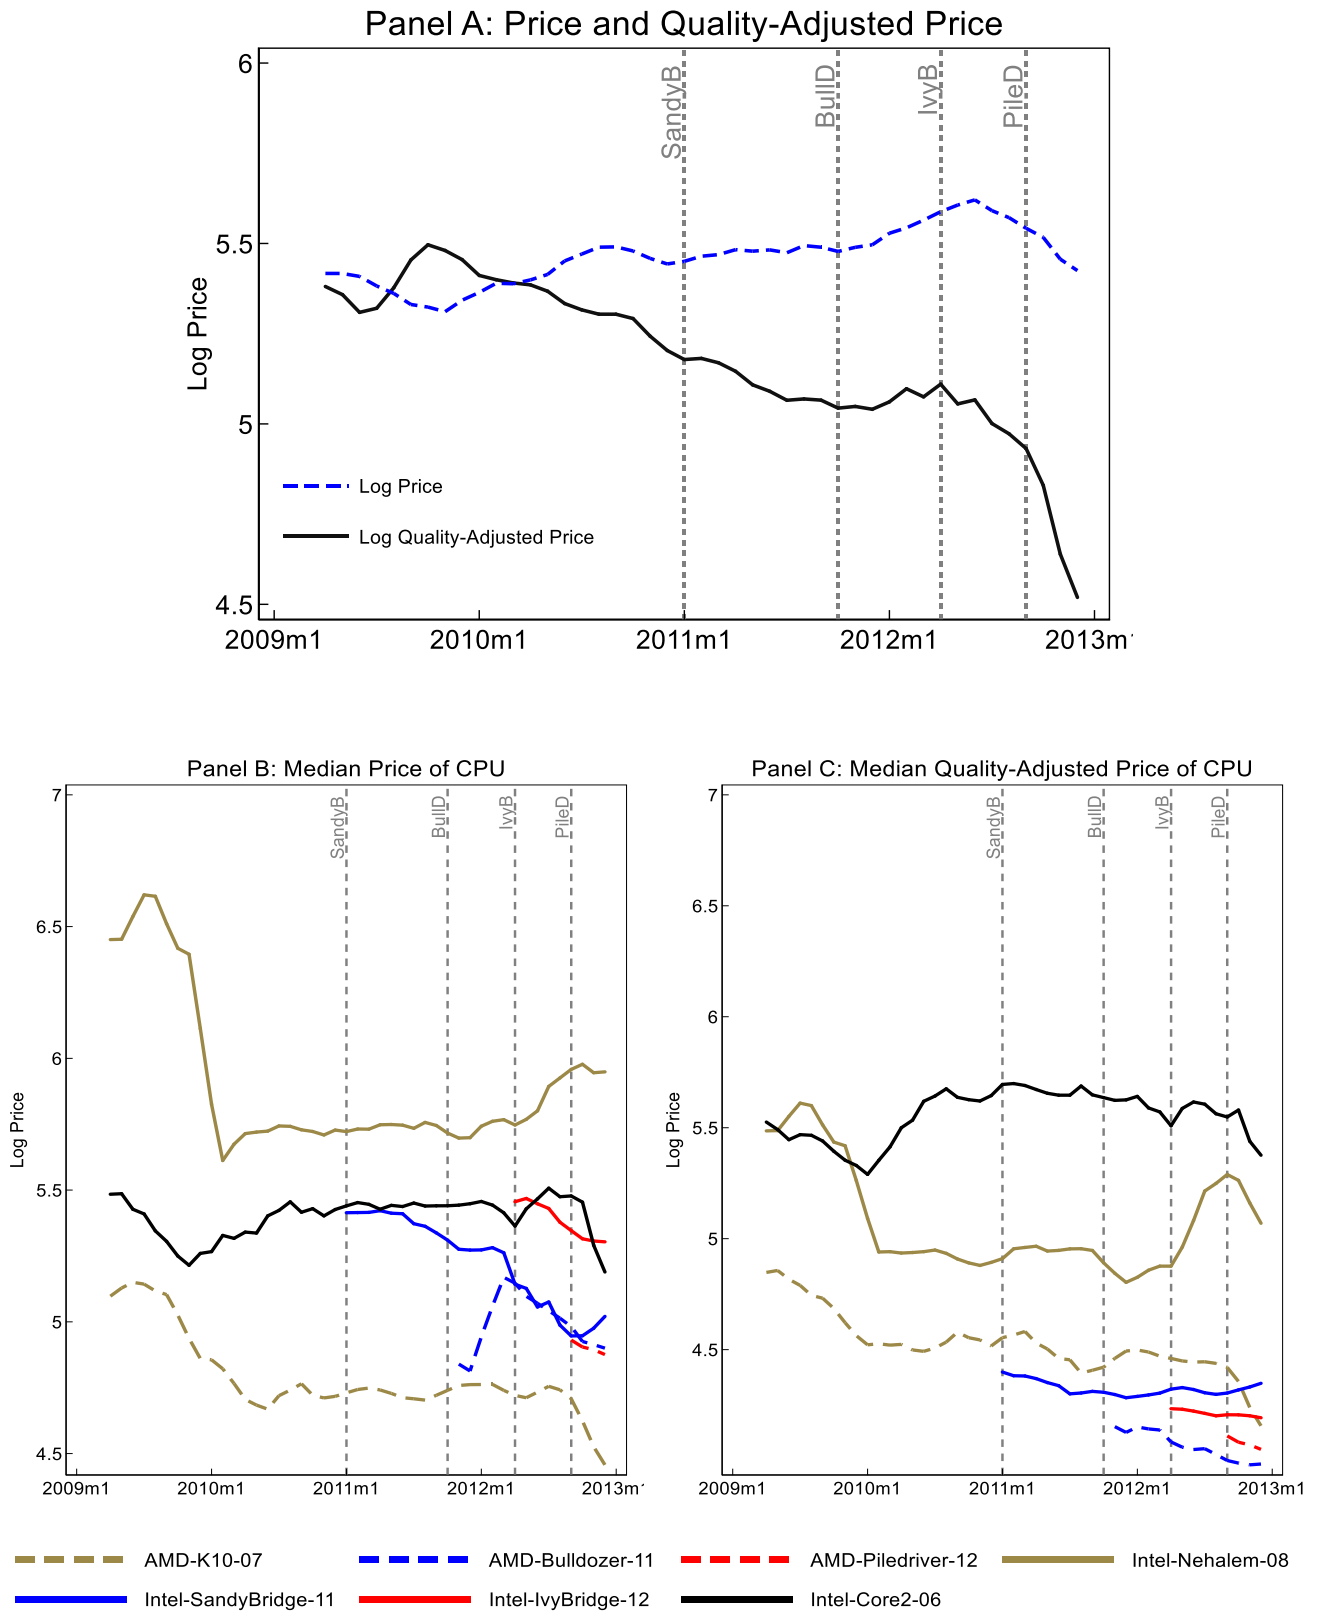
Notes: In this figure, the dots are data points averaged within bins based on percentiles of the log-deviation of price. Lowess smoother is computed with a 0.1 bandwidth.

Figure 4. Quality-Adjusted Price Distribution.



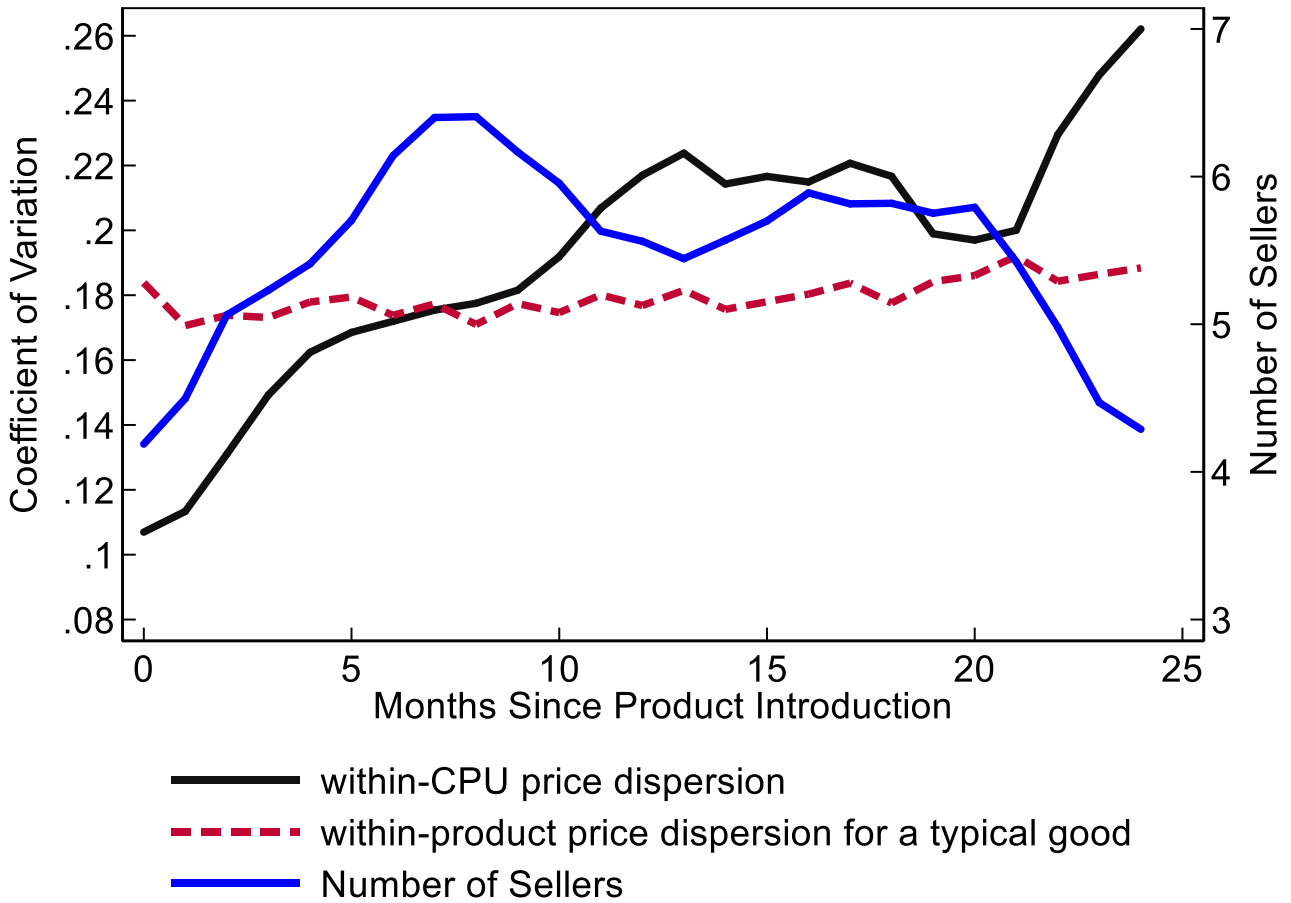
Note: Panel A and B illustrate the distribution of price and quality-adjusted price for each quality quartile, respectively. Panel C and D illustrate the distribution of price and quality-adjusted price for each CPU generation, respectively.

Figure 5. Dynamics of Price and Quality-Adjusted Price.



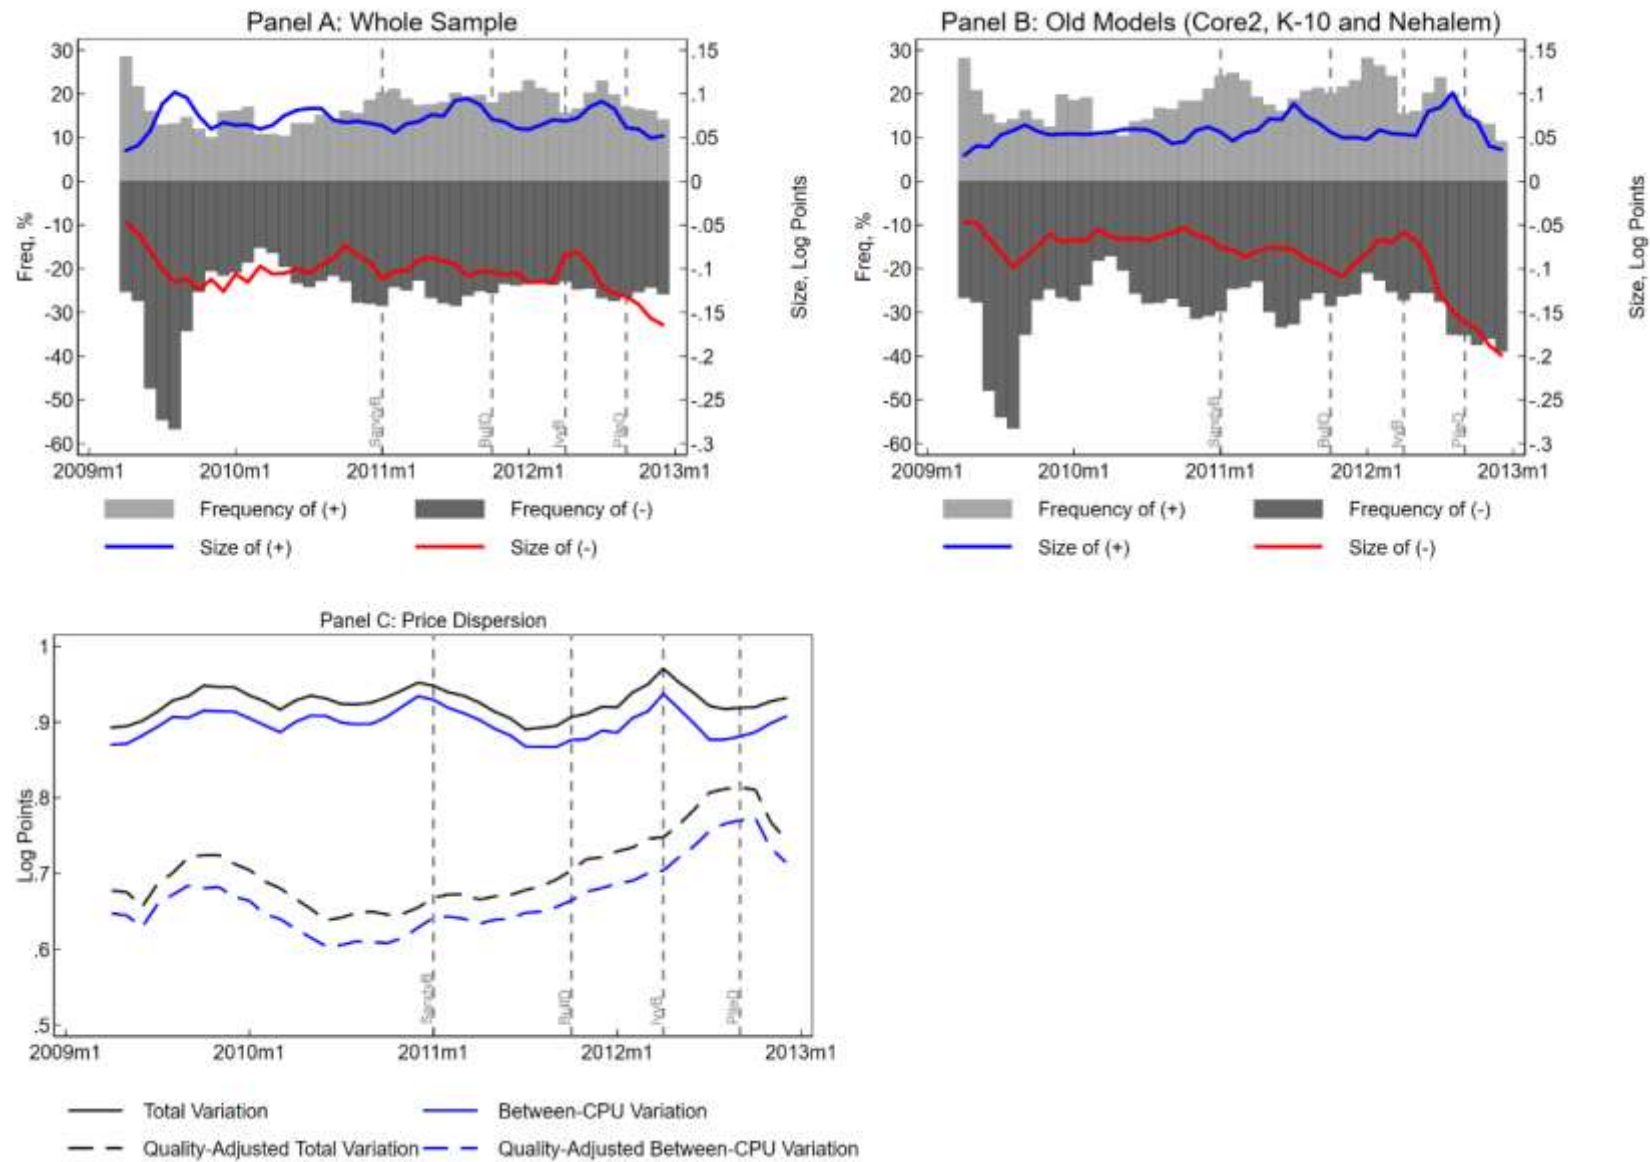
Note: In Panel A, the dashed blue line shows the monthly median log price. The black solid line shows the monthly median log quality-adjusted price. In panel B, each line shows the monthly median log price of a CPU generation. In panel C, each line shows the monthly median log quality-adjusted price of a CPU generation. In all panel, each grey dashed vertical line marks the release month of the corresponding CPU generation.

Figure 6. Average Price Dispersion over CPU life.



Note: The figure plots the non-weighted average (over CPUs) of the coefficient of variation for posted prices against the month passed since the product introduction. CPUs introduced during the first quarter are removed to account for truncated observations, and only CPUs with more than a year of life duration are considered. The red, dashed line reports the evolution of within-product price dispersion for a typical good sold online (Gorodnichenko et al., 2018). The blue, solid line shows the average number of sellers over product lifetime.

Figure 7. Pricing moments and technology shocks.



Note: Panel A and B shows the monthly frequency and absolute size of positive and negative price adjustments for the whole sample and for sample of old models, respectively. In Panel A and B, the vertical axis on the left is the frequency of price changes (%), and the vertical axis on the right is the size of price changes (log points). Panel C shows the monthly total variation and between-CPU variation of raw prices and quality-adjusted prices. In Panel C, the vertical axis on the left is the standard deviation of prices (log points). In all cases, each grey dashed vertical line marks the release month of the corresponding CPU generation.

Table 1. Descriptive Statistics for Performance.

	Mean	SD	Percentiles				
			5%	25%	50%	75%	95%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
All CPUs	2905	2632	237	867	2049	4185	8473
Panel A: Intel CPU Generations							
<i>C2</i>	2079	997	1048	1327	1731	2671	4167
<i>Nehalem</i>	4809	2017	2573	3220	4823	5700	8819
<i>Sandy Bridge</i>	4804	3034	1894	2422	3701	6308	12626
<i>Ivy Bridge</i>	6123	1836	3372	4321	6269	7007	9412
Panel B: AMD CPU Generations							
<i>K10</i>	2788	1034	1607	2000	2662	3379	4901
<i>Bulldozer</i>	5541	1313	4061	4169	5376	6593	7622
<i>Piledriver</i>	4936	2261	2005	3991	4492	6383	8950

Note: Column (1) and (2) present the mean and standard deviation of the log performance score of a CPU (q_i); column (3)-(7) present the mean performance score for each percentile of the CPU's performance score (\bar{Q}_i).

Table 2. Descriptive Statistics for Prices, USD.

	Log Price		Price Percentiles				
	Mean	SD	5%	25%	50%	75%	95%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
No quality adjustment, P	5.23	0.90	53.33	99.95	190.43	334.93	1106.91
Quality adjusted, P^{adj}	5.37	0.84	65.02	111.28	211.79	413.55	1019.73
Representative sample of goods sold online	3.37	1.53	4.00	11.00	25.00	71.00	474.00

Note: Column (1) and (2) present the mean and standard deviation of the average log price for a CPU (\bar{p}_i); column (3)-(7) present the mean price for each percentile of the CPU's price (\bar{P}_i).

Table 3. Monthly Frequency and Size of Price Changes.

	CPU's	Representative sample of goods sold online
	(1)	(2)
All price changes		
Median Frequency, %	39.37	40.12
Implied Duration, Months	2.00	2.54
Median Absolute Size, Log Points	8.10	10.90
Price increases		
Median Frequency, %	15.15	13.87
Median Absolute Size, Log Points	5.57	11.80
Prices decreases		
Median Frequency, %	22.38	15.64
Median Absolute Size, Log Points	9.20	10.30

Note: Column (1) shows the raw frequency and size of price adjustments. Column (2) shows the corresponding moments for a representative sample of goods sold online in the U.S. (Gorodnichenko et al., 2018), converted to a monthly frequency. All results are for regular prices (i.e., excluding sales).

Table 4. Predictor of regular-price stickiness.

Predictors	Frequency of price change, percent			Absolute Size of Price Changes, log points		
	All price changes	Price increases	Price decreases	All price changes	Price increases	Price decreases
	(1)	(2)	(3)	(4)	(5)	(6)
Ln Number of Sellers	-0.000 (0.001)	0.002*** (0.001)	-0.002** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.005*** (0.001)
Ln Performance Scores	0.021* (0.012)	-0.014** (0.006)	0.035*** (0.010)	-0.030*** (0.007)	-0.033*** (0.009)	-0.030*** (0.007)
Ln Median Price	0.210*** (0.068)	0.099*** (0.036)	0.112** (0.055)	0.119*** (0.041)	0.155*** (0.054)	0.084* (0.043)
Ln Median Price Squared	-0.021*** (0.006)	-0.008*** (0.003)	-0.012** (0.005)	-0.009** (0.004)	-0.013*** (0.005)	-0.007* (0.004)
Share of Price Points	-0.081*** (0.031)	-0.078*** (0.016)	-0.003 (0.025)	0.084*** (0.020)	0.114*** (0.027)	0.074*** (0.021)
Stability of Sellers	-0.473*** (0.100)	-0.066 (0.053)	-0.406*** (0.081)	-0.003 (0.057)	-0.013 (0.075)	-0.055 (0.060)
Intel CPU	0.054*** (0.016)	0.032*** (0.008)	0.022* (0.013)	-0.005 (0.009)	-0.015 (0.012)	0.025*** (0.009)
R ²	0.385	0.362	0.303	0.451	0.354	0.450
N	608	608	608	607	572	585

Note: This table shows the regression results of the frequency of price changes in column (1), frequency of positive price changes in column (2), frequency of negative price changes in column (3), absolute size of price changes in column (4), absolute size of positive price changes in column (5), and absolute size of negative price changes in column (6) for regular prices on the set of dependent variables above. Time fixed effects and the constant are included in all regressions but not reported. All variables are unweighted and measured at product level. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Table 5. within-CPU Dispersion of Regular-Price across Sellers.

Row		CPU _s	Representative sample of goods sold online
		(1)	(2)
Panel A. Measures of price dispersion			
(1)	Coefficient of Variation (CV)	0.22	0.21
(2)	Std(log price)	0.23	0.24
(3)	Value of information (VI)	0.25	0.24
(4)	Interquartile range (IQR)	0.32	0.35
(5)	Range	0.47	0.41
(6)	Gap	0.22	0.28
(7)	Std(ε)	0.18	0.21
Panel B. Spatial versus Temporal Price Dispersion			
<i>Less than 5% of the time in a given quartile</i>			
(8)	1st Quartile	0.42	0.43
(9)	2nd Quartile	0.39	0.46
(10)	3rd Quartile	0.37	0.44
(11)	4th Quartile	0.50	0.60
<i>More than 95% of the time in a given quartile</i>			
(12)	1st Quartile	0.10	0.16
(13)	2nd Quartile	0.014	0.044
(14)	3rd Quartile	0.024	0.057
(15)	4th Quartile	0.073	0.065
(16)	Mean SD of Time in Each Quartile	0.29	0.31

Note: Row (1) – (7) in Panel A report the average dispersion of regular prices measured with: the coefficient of variance (CV), which is computed as the standard deviation divided to the mean; $\text{std}(\log p)$, which is the standard deviation of the log price; value of information (VI), which is computed as the log difference between the average and the minimum price; interquartile range (IQR) equal to the log difference between the 75th and 25th percentile; range is the log difference between the highest and lowest price; gap is the log difference between the two lowest prices; and $\text{std}(\varepsilon)$, in which ε is the error term in the regression of $\log p$ on good and seller fixed effects; respectively, for CPU online-market in the US. Furthermore, for each price line, we calculate the proportion of time that the price line stays in each quartile of the cross-seller price distribution. In Panel B, row (8) – (11) and (12) – (15) report the proportion of price lines that almost never (less than 5% of the time) or almost always (more than 95% of the time) fall into a given quartile, respectively. The bottom line shows the average (across price lines) standard deviation of the proportion of time spent in each quartile. Under perfectly temporal dispersion, this measure is 0, while under perfectly spatial dispersion, it is approximately 0.43. Column (2) shows the corresponding statistics for a representative sample of goods sold online (Gorodnichenko et al., 2018) for the U.S.

Table 6. Predictors of within-CPU Regular-Price Dispersion (at Product Level).

Predictors	Dependent variable: average within-product price dispersion	
	log price, $\bar{\sigma}_i(p_i)$	log price net of seller fixed effects, $\bar{\sigma}_i(\varepsilon)$
	(1)	(2)
Ln Number of Sellers	-0.003** (0.001)	-0.001 (0.001)
Ln Performance Scores	-0.089*** (0.019)	-0.078*** (0.014)
Ln Median Price	0.059*** (0.015)	0.056*** (0.011)
Share of Price Points	0.024 (0.052)	-0.024 (0.040)
Frequency of Regular Price Changes, %	0.215*** (0.069)	0.122** (0.053)
Size of Regular Price Changes, Log Points	0.009*** (0.001)	0.008*** (0.001)
Seller Stability	0.212 (0.137)	0.117 (0.104)
Intel CPU	0.005 (0.021)	-0.006 (0.016)
R ²	0.524	0.555
N	496	496

Note: This table shows the results of the regression of the standard deviation of log price in column (1), and the regression results after removing seller fixed effects in column (2) on the set of dependent variables above. Time fixed effects and the constant are included in all regressions but not reported. All the reported variables in this table are unweighted and measured at product level. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Table 7. Predictors of between-CPU price variation.

Predictors	Dependent variable: \bar{p}_{it} , average CPU (log) price			
	(1)	(2)	(3)	(4)
Ln Performance scores	0.496*** (0.007)		0.613*** (0.009)	0.477*** (0.041)
Ln Number of Seller		0.153*** (0.009)	-0.185*** (0.009)	-1.136*** (0.067)
Share of Price Points		-0.190*** (0.034)	-0.143*** (0.028)	0.650*** (0.195)
Stability of sellers		-0.193*** (0.041)	0.327*** (0.034)	1.422*** (0.290)
Intel CPUs		0.303*** (0.020)	0.475*** (0.018)	-0.837*** (0.177)
Ln Seller \times Ln performance				0.123*** (0.008)
SPP \times Ln performance				-0.107*** (0.026)
Seller Stability \times Ln performance				-0.154*** (0.037)
Intel \times Ln performance				0.165*** (0.023)
R ²	0.313	0.063	0.390	0.409
N	9590	9590	9590	9590

Note: The regression reported in column (1) uses only quality measure as independent variables. The regression reported in column (2) uses market/product characteristics as independent variables. The regression reported in column (3) uses quality measure and market/product characteristics as independent variables. The regression reported in column (4) uses quality measure, market/product characteristics, and interaction terms between quality and market/product characteristics as independent variables. In all panels, the dependent variable is $\ln(\text{price})$. All the reported variables in this table are unweighted and measured at product-month level. The constant is included in all regressions but not reported. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Table 8. Price adjustment in response to technological change.

Predictors	Frequency of price change, percent			Absolute Size of Price Changes, log points		
	All price changes	Price increases	Price decreases	All price changes	Price increases	Price decreases
	(1)	(2)	(3)	(4)	(5)	(6)
Ln Number of Seller	0.002 (0.006)	0.002 (0.005)	0.000 (0.006)	-0.033*** (0.003)	-0.022*** (0.004)	-0.036*** (0.004)
Ln Median Price	0.657*** (0.076)	0.380*** (0.055)	0.277*** (0.067)	-0.054 (0.065)	0.093 (0.109)	-0.085 (0.074)
Ln Median Price Squared	-0.069*** (0.007)	-0.029*** (0.005)	-0.040*** (0.006)	-0.002 (0.006)	-0.007 (0.010)	-0.003 (0.007)
Share of Price Points	-0.134*** (0.013)	-0.078*** (0.010)	-0.056*** (0.011)	0.019** (0.009)	0.012 (0.013)	0.019* (0.010)
Stability of Sellers	-0.030* (0.017)	-0.039*** (0.013)	0.009 (0.015)	0.036*** (0.011)	0.018 (0.014)	0.047*** (0.013)
Upgrading	0.075 (0.629)	0.015 (0.347)	0.061 (0.525)	-0.901* (0.463)	-0.628 (0.653)	-1.343** (0.554)
Downgrading	0.016 (0.277)	0.067 (0.202)	-0.051 (0.264)	0.049 (0.186)	-0.083 (0.195)	0.127 (0.228)
Ln Number of Seller × Upgrading	-0.007 (0.029)	0.014 (0.015)	-0.021 (0.026)	0.007 (0.016)	-0.019 (0.027)	0.007 (0.018)
Ln Number of Seller × Downgrading	-0.054*** (0.009)	-0.022*** (0.007)	-0.032*** (0.009)	-0.015*** (0.005)	-0.018** (0.008)	-0.014** (0.006)
Ln Median Price × Upgrading	-0.002 (0.237)	-0.002 (0.134)	0.000 (0.199)	0.419** (0.192)	0.313 (0.278)	0.605*** (0.232)
Ln Median Price × Downgrading	0.005 (0.103)	0.001 (0.076)	0.004 (0.099)	0.016 (0.065)	0.042 (0.066)	-0.003 (0.080)
Ln Median Price Squared × Upgrading	0.001 (0.022)	-0.002 (0.013)	0.003 (0.019)	-0.043** (0.019)	-0.035 (0.029)	-0.061*** (0.023)
Ln Median Price Squared × Downgrading	0.002 (0.010)	0.000 (0.007)	0.001 (0.009)	-0.003 (0.006)	-0.004 (0.006)	-0.002 (0.007)
Share of Price Points × Upgrading	-0.034 (0.072)	0.012 (0.036)	-0.046 (0.063)	-0.094 (0.073)	-0.087 (0.086)	-0.057 (0.077)
Share of Price Points × Downgrading	0.061* (0.033)	0.025 (0.024)	0.037 (0.029)	0.009 (0.025)	0.055 (0.036)	-0.020 (0.029)
Stability of Sellers × Upgrading	-0.106 (0.136)	-0.034 (0.066)	-0.072 (0.115)	-0.118* (0.070)	0.038 (0.074)	-0.149** (0.073)
Stability of Sellers × Downgrading	0.086* (0.052)	-0.030 (0.043)	0.116** (0.050)	-0.017 (0.035)	0.025 (0.044)	-0.024 (0.044)
R ²	0.285	0.143	0.217	0.324	0.324	0.348
N	14448	14448	14448	8498	5216	6744

Note: This table shows the regression results of the frequency of price changes in column (1), frequency of positive price changes in column (2), frequency of negative price changes in column (3), absolute size of price changes in column (4), absolute size of positive price changes in column (5), and absolute size of negative price changes in column (6) for regular prices on the set of dependent variables above. The regression is at product-month level and includes product fixed effects. All variables are unweighted and measured at product-month level. The constant is included in all regressions but not reported. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Table 9. Inflation in CPU prices, 2009 – 2012.

Year	Inflation				
	(1)	(2)	(3)	(4)	(5)
2010	0.253*** (0.035)	-0.042 (0.028)	0.136*** (0.036)	0.037 (0.027)	0.033 (0.027)
2011	-0.032 (0.027)	-0.114*** (0.022)	-0.023 (0.027)	-0.106*** (0.021)	-0.132*** (0.021)
2012	-0.016 (0.021)	0.052*** (0.018)	0.061*** (0.022)	-0.043** (0.019)	-0.063*** (0.019)
Memorandum: cumulative inflation	20.5%	-10.4%	17.4%	-11.2%	-16.2%
Controls					
Performance scores	No	Yes	No	Yes	Yes
Market/product characteristics	No	No	Yes	Yes	Yes
Interactions included	No	No	No	No	Yes

Note: In this table, the dependent variable is \bar{p}_{it} . The regression is run separately for each overlapping two-year period between 2009 and 2012. All the reported variables in this table are unweighted and measured at product-month level. The constant is included in all regressions but not reported. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Appendix A

Additional Figures and Tables

Appendix Table A.1. Monthly Frequency and Size of Sales.

	One-month Two-sided Sales Filter		
	Mean Frequency	SD Frequency	Median Size
	(1)	(2)	(3)
CPU	1.88	3.64	2.28
Representative sample of goods sold online	7.82	16.54	10.5

Note: Column (1) shows the monthly average of sales frequency across CPUs (%). Column (2) reports the standard deviation of sales frequency across CPUs. Column (3) shows the absolute size of sales for the median CPU, in which the absolute size of sales equal to the gap between the log of sales price and the log of regular price (multiple by 100). Column (4) shows the number of CPUs. A sales is identified by using the one-month, two-sided sales filter. The second row shows the corresponding moments for a representative sample of goods sold online (Gorodnichenko et al., 2018) for the U.S., converted to a monthly frequency.

Appendix Table A.2. Price adjustment in response to technological change, alternative specification.

Predictors	Frequency of Price Changes, % (1)	Frequency of Positive Changes, % (2)	Frequency of Negative Changes, % (3)	Absolute Size of Price Changes, Log Points (4)	Absolute Size of Positive Changes, Log Points (5)	Absolute Size of Negative Changes, Log Points (6)
Ln Number of Seller	-0.013** (0.006)	0.001 (0.005)	-0.014** (0.006)	-0.034*** (0.003)	-0.024*** (0.004)	-0.036*** (0.004)
Ln Median Price	0.634*** (0.076)	0.349*** (0.054)	0.284*** (0.067)	-0.081 (0.064)	0.081 (0.106)	-0.109 (0.073)
Ln Median Price Squared	-0.068*** (0.007)	-0.026*** (0.005)	-0.041*** (0.006)	0.000 (0.006)	-0.006 (0.010)	-0.001 (0.007)
Share of Price Points	-0.124*** (0.013)	-0.072*** (0.009)	-0.051*** (0.011)	0.024*** (0.009)	0.021* (0.012)	0.021** (0.010)
Stability of Sellers	-0.036** (0.017)	-0.053*** (0.013)	0.017 (0.015)	0.035*** (0.011)	0.024* (0.013)	0.042*** (0.013)
Ln Number of Higher CPU Enter	-0.012*** (0.004)	0.004 (0.003)	-0.017*** (0.003)	0.012*** (0.002)	0.008*** (0.002)	0.011*** (0.002)
Ln Number of Lower CPU Enter	0.057*** (0.005)	0.033*** (0.003)	0.024*** (0.004)	-0.004 (0.003)	-0.002 (0.003)	-0.001 (0.004)
Ln Number of Higher CPU Exit	-0.014*** (0.005)	0.011*** (0.004)	-0.025*** (0.005)	0.008*** (0.003)	0.001 (0.003)	0.012*** (0.004)
Ln Number of Lower CPU Exit	0.011*** (0.004)	0.011*** (0.003)	-0.000 (0.004)	0.003 (0.002)	0.005** (0.002)	-0.000 (0.003)
R ²	0.280	0.146	0.214	0.323	0.322	0.347
N	14448	14448	14448	8498	5216	6744

Note: This table shows the regression results of the frequency of price changes in column (1), frequency of positive price changes in column (2), frequency of negative price changes in column (3), absolute size of price changes in column (4), absolute size of positive price changes in column (5), and absolute size of negative price changes in column (6) for regular prices on the set of dependent variables above. Product fixed effects and the constant are included in all regressions but not reported. All variables are unweighted and measured at product-month level. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Appendix Table A.3. Price adjustment in response to technological change, alternative specification.

	Frequency of Price Changes, %	Frequency of Positive Changes, %	Frequency of Negative Changes, %	Absolute Size of Price Changes, Log Points	Absolute Size of Positive Changes, Log Points	Absolute Size of Negative Changes, Log Points
Predictors	(1)	(2)	(3)	(4)	(5)	(6)
Ln Number of Seller	-0.005 (0.006)	-0.000 (0.005)	-0.004 (0.006)	-0.035*** (0.003)	-0.025*** (0.004)	-0.038*** (0.004)
Ln Median Price	0.656*** (0.075)	0.381*** (0.055)	0.275*** (0.067)	-0.060 (0.064)	0.098 (0.106)	-0.093 (0.073)
Ln Median Price Squared	-0.069*** (0.007)	-0.029*** (0.005)	-0.040*** (0.006)	-0.002 (0.006)	-0.007 (0.010)	-0.002 (0.007)
Share of Price Points	-0.126*** (0.013)	-0.075*** (0.009)	-0.052*** (0.011)	0.021** (0.009)	0.021* (0.012)	0.017 (0.010)
Stability of Sellers	-0.021 (0.017)	-0.042*** (0.013)	0.021 (0.015)	0.035*** (0.011)	0.023* (0.013)	0.044*** (0.012)
Upgrading	-0.013 (0.024)	-0.053*** (0.012)	0.039* (0.021)	0.000 (0.015)	0.028 (0.020)	-0.009 (0.017)
Downgrading	0.131*** (0.010)	0.045*** (0.008)	0.086*** (0.009)	0.023*** (0.006)	0.017*** (0.006)	0.023*** (0.006)
R ²	0.282	0.142	0.215	0.322	0.321	0.345
N	14448	14448	14448	8498	5216	6744

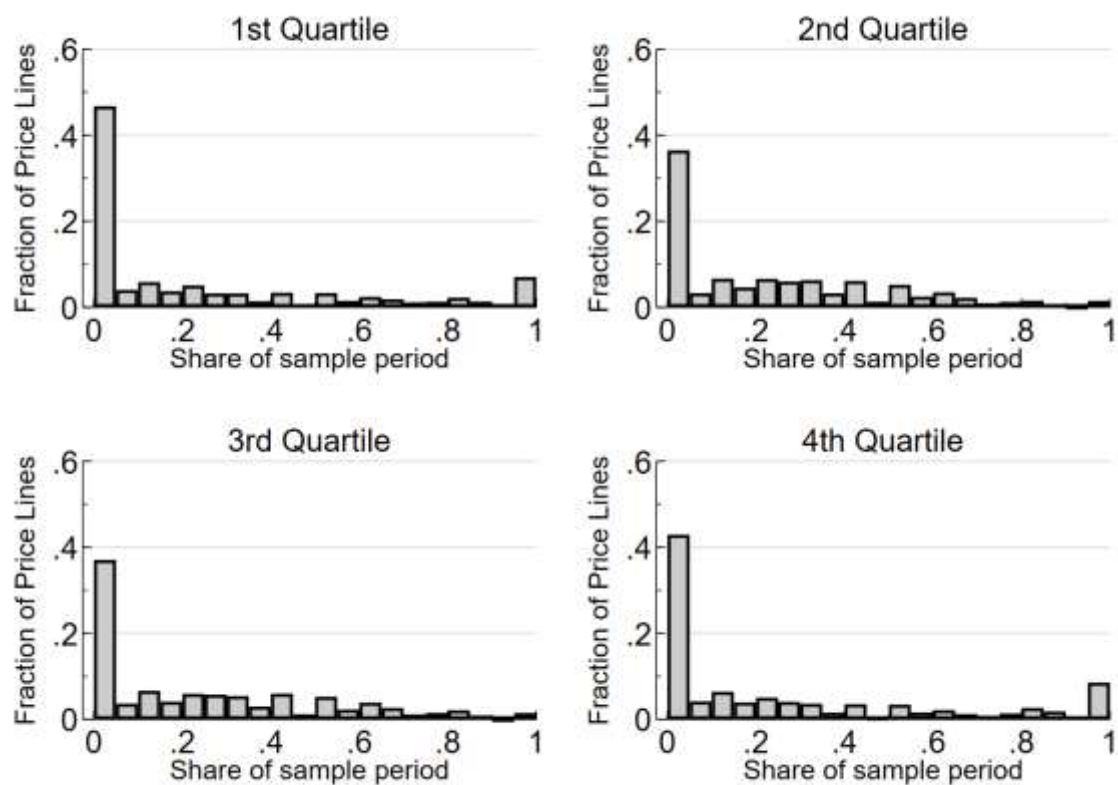
Note: This table shows the regression results of the frequency of price changes in column (1), frequency of positive price changes in column (2), frequency of negative price changes in column (3), absolute size of price changes in column (4), absolute size of positive price changes in column (5), and absolute size of negative price changes in column (6) for regular prices on the set of dependent variables above. Product fixed effects and the constant are included in all regressions but not reported. All variables are unweighted and measured at product-month level. Robust standard errors are in parentheses. *Significant at 10% level; **Significant at 5% level; ***Significant at 1% level.

Appendix Table A.4. Fictitious examples of temporal and spatial price dispersion across stores.

	Period								
	0	1	2	3	4	5	6	7	8
Panel A. Temporal price dispersion									
Price for Seller A	1	2	1	2	2	1	2	1	1
Price for Seller B	2	1	2	1	1	1	1	2	2
Panel B. Spatial price dispersion									
Price for Seller A	1	1	1	1	1	1	1	1	1
Price for Seller B	2	2	2	2	2	2	2	2	2

Notes: Panel A shows that which store charges the high price varies over time (temporal price dispersion), while Panel B shows that which store charges the high price is stable over time (spatial price dispersion). We do our analysis of spatial vs. temporal price dispersion to establish whether one can use outlets for sampling prices interchangeably. Specifically, if price dispersion is temporal, the identity of the seller is less important because on average one can recover the average price in the market by following this seller. In contrast, if price dispersion is spatial, the identity of the seller is important for measuring the level of prices in the market.

Appendix Figure A1. Fraction of Price Lines in each Quartile of the CPU Price Distribution.



Note: For each monthly product–seller price quote, the portion that the price quote stayed in each of the 4 quartiles of cross-seller price distribution is calculated. The figure illustrates the distribution across price quotes.

Appendix B

Additional details for notation

A. Frequency of Price Changes

Following previous studies, we determine the frequency of price adjustment as the proportion of non-zero price changes to the total number of price changes observed within our dataset. Particularly, we consider a price change that is smaller than 0.1% as a zero-price change, which means it is not counted as a non-zero price change. In other words, if we use $\varphi_{ist} = \mathbb{I}\{S_{is,t} > 0\} \mathbb{I}\{S_{is,t-1} > 0\}$ to identify a price adjustment that is observed (for both zero and non-zero price changes); the number of observed price adjustment per CPU-seller quote is $\Pi_{is} = \sum_t \varphi_{ist}$; the conditional function of a non-zero change is $\chi_{ist} = \mathbb{I}\{|\Delta \log P_{ist}| > 0.001\}$. Hence, the formula of the frequency of price adjustments of CPU i sold by seller s is:

$$f_{is} = \frac{\sum_t \chi_{ist}}{\Pi_{is}} \quad (3)$$

After that, we collapse the result to CPU-level by computing the raw average frequency (\bar{f}_i) for each CPU, which has more than two observations, as follow:

$$\bar{f}_i = \frac{1}{\sum_{s \in \mathcal{S}_i} \mathbb{I}\{\Pi_{is} > 2\}} \sum_{s \in \mathcal{S}_i} f_{is} \mathbb{I}\{\Pi_{is} > 2\} \quad (4)$$

Then, we compute the no weights average (\bar{f}) and between-CPU weights average (\bar{f}^b) for both posted and regular price changes frequency, which is reported in Table 4, as below:

$$\bar{f} = \sum_{i \in \mathcal{C}} \frac{1}{C} \sum_{s \in \mathcal{S}_i} f_{is} \mathbb{I}\{\Pi_{is} > 2\} * \frac{1}{\sum_{s \in \mathcal{S}_i} \mathbb{I}\{\Pi_{is} > 2\}} \quad (5)$$

Lastly, we compute the implied duration of price spell as:

$$\bar{d}_i = -\frac{1}{\ln(1-\bar{f}_i)} \quad (6)$$

B. Size of Price Changes

Using the same symbols in the notation part, our formula to calculate the average absolute size of price adjustment for CPU i is:

$$\overline{|\Delta \log P_i|} = \frac{1}{\sum_{s \in \mathcal{S}_i} \sum_t \chi_{ist}} \sum_{s \in \mathcal{S}_i} \sum_t |\Delta \log P_{ist}| \cdot \chi_{ist} \quad (7)$$