

Karadi, Peter; Schoenle, Raphael; Wursten, Jesse

Working Paper

Measuring price selection in microdata: It's not there

ECB Working Paper, No. 2566

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Karadi, Peter; Schoenle, Raphael; Wursten, Jesse (2021) : Measuring price selection in microdata: It's not there, ECB Working Paper, No. 2566, ISBN 978-92-899-4752-7, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/510642>

This Version is available at:

<https://hdl.handle.net/10419/237705>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

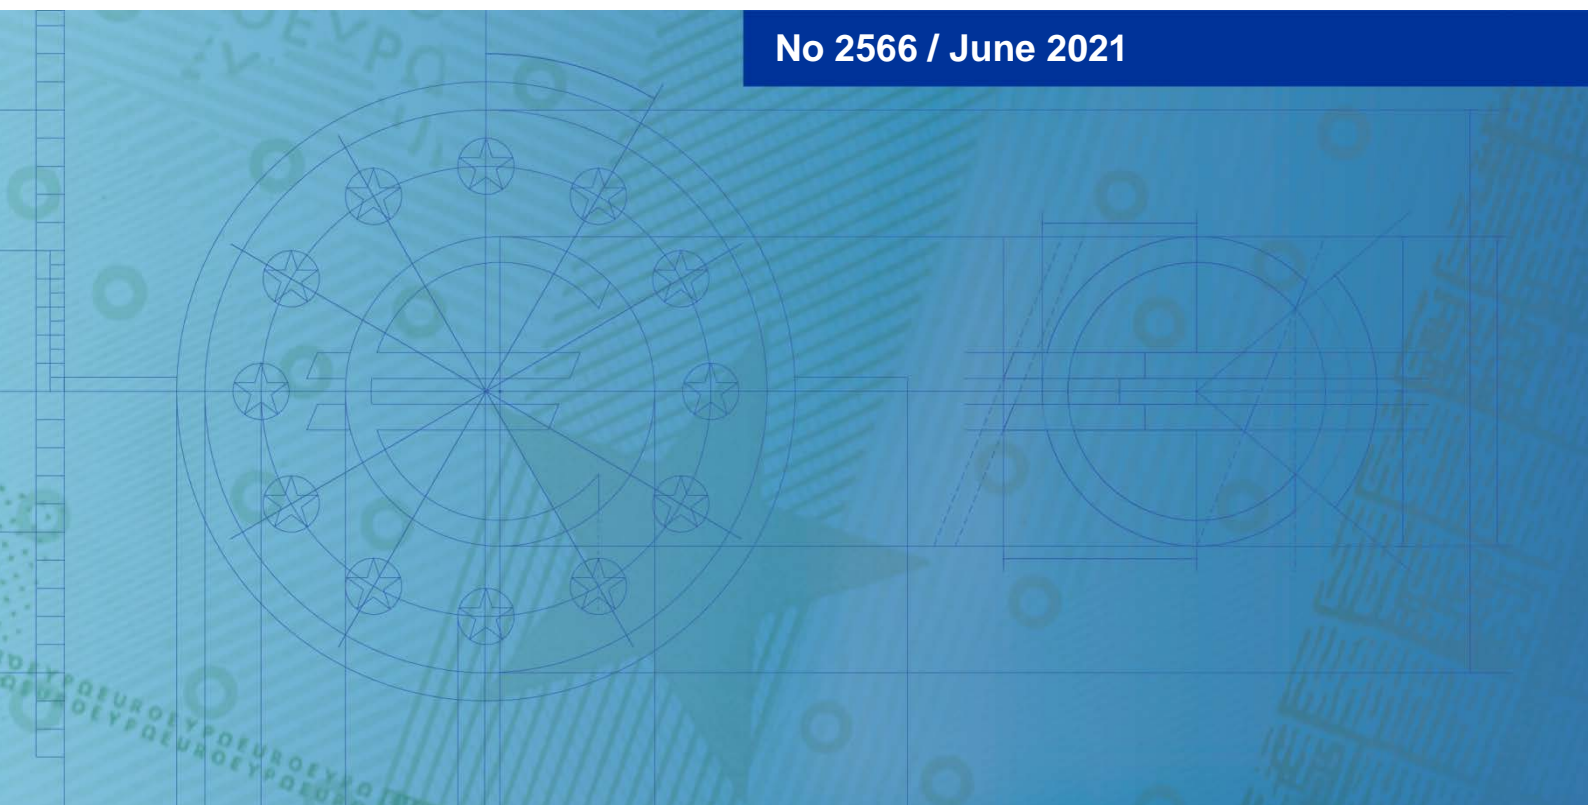
EUROSYSTEM

Working Paper Series

Peter Karadi, Raphael Schoenle,
Jesse Wursten

Measuring price selection
in microdata:
it's not there

No 2566 / June 2021



Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

We use microdata to estimate the strength of price selection – a key metric for the effect of monetary policy on the real economy. We propose a product-level proxy for mispricing and assess whether products with larger mispricing respond with a higher probability to identified monetary and credit shocks. We find that they do not, suggesting selection is absent. Instead, we detect state-dependent adjustment on the gross extensive margin. Our results are broadly consistent with second-generation state-dependent pricing models and sizable effects of monetary policy on the real economy.

JEL codes: E31, E32, E52

Keywords: monetary non-neutrality, state-dependent pricing, identified credit and monetary policy shocks, price-gap proxy, scanner data, PPI microdata

1 Non-Technical Summary

The ability of monetary policy to stabilize business cycles depends on the flexibility of the price level. If prices are very flexible, they will absorb shocks, if they are sticky, activity will adjust instead. However, the price level can be flexible even if only a few prices adjust, as long as the prices that react do so in a disproportionate manner. Such ‘self-selection’ of large price changes tends to reduce the real effects of monetary policy shocks in realistic price-setting frameworks. In this paper, we use micro price data to measure the strength of price selection.

We carefully measure the selection effect in the data. In our main analysis, we use a detailed, weekly panel of barcode-level prices in U.S. supermarkets between 2001 and 2012. The granularity of the data in the cross-section allows us to identify product-level pricing pressures, while the long time series dimension allows us to identify aggregate shocks. Selection is present if, following an aggregate shock, prices that are further away from their optima respond with higher probability than those that are closer. There are two key challenges. First, optimal prices and therefore price gaps are unobserved. Our approach is to rely on the pricing behavior of close substitutes to measure unobserved optimal prices (Gagnon et al., 2012). In particular, we measure the difference between a price and the average price of the exact same good sold by competing supermarkets, after we control for permanent store-level characteristics (coming from heterogeneity among regions and amenities). We show that our price gap proxy indeed predicts future price changes indicating that it is a valid measure of product-level price pressures. The second challenge is to identify aggregate shocks. We identify aggregate credit shocks using conventional timing restrictions (Gilchrist and Zakrajšek, 2012). We show that the aggregate shocks also affect the probability of price changes, indicating that the shock we identify is relevant for price setting.

To measure the strength of selection, we assess whether our price-gap proxies interact with identified aggregate shocks to influence the probability of price adjustment. We approach this question in a linear-probability panel-regression framework. Our dependent variables are respectively the probability of price increases and the probability of price decreases over the 24 months following the aggregate shock. We test whether the interaction term of the price gap and the aggregate shock significantly influences the probability of those price changes. If the coefficient of the interaction term is high, then, conditional on an aggregate shock, prices with higher price gaps respond with higher probability. As controls, we add the direct effects of both the price gap and the aggregate shock as well as the age of the price and numerous fixed effects.

Our main finding is that selection appears to be absent in the data as the interaction term stays insignificant in all of our regressions. These results do not change if we use alternative price-gap proxies, different aggregate shocks, different datasets, and different specifications. This is true, even though prices do respond directly to both aggregate and product-level price pressures in all our specifications. Our evidence support frameworks with a micro-macro dichotomy: even though prices respond to product-level pressures quite flexibly, the aggregate price level is sticky and monetary policy can stabilize business cycles.

2 Introduction

The ability of monetary policy to stabilize business cycles depends on the flexibility of the price level. If prices are very flexible, they will absorb shocks, if they are sticky, activity will adjust instead. However, the price level can be flexible even if only a few prices adjust, as long as the prices that react do so in a disproportionate manner. Such ‘self-selection’ of large price changes tends to reduce the real effects of monetary policy shocks in price-setting frameworks where frictions to price adjustment are microfounded by small, fixed (menu) costs (Golosov and Lucas, 2007). In this paper, we use micro price data to measure the strength of price selection.

Our model-free finding is that selection is absent in the data. We arrive at this result by analyzing the probability of price adjustments in response to an identified aggregate shock. Even though prices do respond to the shock directly, the probability of adjustment in response to a shock is not a function of the extent of their mispricing, which selection would require. Instead, we detect a uniform shift from price increases towards price decreases after a policy tightening. These results imply that price adjustment in macroeconomic models generally should feature a state-dependent adjustment through the gross extensive margin and weak selection, implying a sizable monetary non-neutrality. We characterize conditions which state-dependent pricing models with random menu cost (such as Dotsey et al., 1999; Luo and Villar, 2017; Alvarez et al., 2020) or rational inattention (Woodford, 2009) are consistent with our findings.

We measure selection by considering how the probability of price adjustment depends on an interaction of a macro-level and a micro-level component. The macro-level component is an aggregate shock which hits the economy and exerts a general adjustment force. The micro-level component is the pre-existing product-level mispricing, which is present because sticky prices failed to keep up with changes in their optimal levels. A non-zero interaction term of the second kind implies that prices of goods with large mispricing are more likely to adjust when a shock hits the economy. Such a relation would indicate the presence of a selection effect: Prices respond to a shock disproportionately strongly to make up for any already existing product-level mispricing. As a result, the aggregate price level becomes disproportionately flexible. If the prices of goods with larger mispricing do not react with a higher probability to a shock, then selection is absent. This definition of selection can be formalized by a suitable generalization of the accounting framework of Caballero and Engel (2007), as we explain below.

The empirical implementation of our approach to measure selection is straightforward. We use micro price data and regression analysis. There are two key challenges. The first is to obtain a suitable proxy for mispricing. We achieve this by using the information on the price-setting of close substitutes, as we explain below. The second is to identify an aggregate shock, which generates an uniform price-adjustment impetus. Our baseline aggregate shock is a credit shock, which we identify using timing restrictions as in Gilchrist and Zakrajšek (2012). Our results are also robust to monetary policy shocks identified using high-frequency surprises in interest rates around policy announcements (Gertler and Karadi, 2015; Jarociński and Karadi, 2020). With these proxies in hand, we can measure selection by estimating how the probability of price adjustment statistically depends on the interaction of the aggregate shock and the product-level price-gap proxy. We approach our question in a linear-probability panel-regression framework. Our dependent variables are the probability of price increases and price decreases, respectively, over the 24 months following a credit shock. Besides the interaction term of the current aggregate shock and the lagged price gap proxy that is the focus of our analysis, we also include the previous-period product-level price gap and the current-

period aggregate shock as explanatory variables. Furthermore, we satiate the regressions with a rich set of fixed effects, control for the age of prices, and cluster the standard errors across time and across product categories.

On the data side, we construct our baseline measure of the price gap using a detailed, weekly panel of barcode-level prices in U.S. supermarkets between 2001 and 2012, compiled by the marketing company IRI.¹ In a complementary analysis, we also show that our results are robust to using the producer-price microdata that underlies the U.S. producer-price index (PPI).

Our baseline measure of the price gap calculates the distance of the price from the average of close competitors, after controlling for permanent differences in prices coming from variation in geography and amenities (Gagnon et al., 2012). This ‘competitor-price gap’ is a relevant measure of product-level price pressures as long as sellers strive to keep their prices close to competitors’ prices to maintain profitability and market share. We show that our results are robust to using an alternative ‘reset-price’ gap measure. Here, we build on the approach proposed by Bils, Klenow and Malin (2012), who define reset prices as the unobserved optimal prices that a store would set if pricing frictions were temporarily removed. We follow the algorithm of Bils, Klenow and Malin (2012) to approximate these.

We find that our price gap measure is a valid proxy of product-level price pressures because first, they raise the probability of future price adjustments and second, the average size of the adjustment is proportional to the distance between optimal and actual prices. Furthermore, the aggregate shocks are also relevant, because they significantly change the probability of price increases and decreases. However, we find no evidence of price selection because the interaction term is consistently insignificant in all of our regressions. These results are robust to using the different price gap proxies (competitor price gaps or reset-price gaps), different aggregate shocks (credit shock versus monetary policy shock), different datasets (supermarket scanner versus producer-price microdata) and different specifications (linear versus non-linear probability models). Selection is always absent.

Our results pose a challenge to conventional price-setting models. This result arises as we interpret our evidence through a suitable generalization of the flexible accounting framework of Caballero and Engel (2007). According to this generalization, the price-level response to an aggregate shock can be decomposed into three adjustment margins. The first is the intensive margin which is the only one present in time-dependent models. It accounts for the change in the magnitude of price changes, when adjusting firms incorporate the impact of the aggregate shock to their price changes. The second is a gross extensive margin, which measures the adjustment through the uniform shift between price increases versus price decreases. The third is a selection effect, which measures whether goods with larger mispricing respond with higher probability to the aggregate shock. The second and the third margins can be jointly referred to as the extensive margin effects (as in Caballero and Engel, 2007). Our evidence, on the one hand, points to the presence of the gross extensive margin, as we find robust evidence on the uniform shift between price increases versus price decreases after an aggregate shock. These results are inconsistent with standard time-dependent pricing models (Calvo, 1983), where the gross extensive margin is absent,² and implies state-dependence of price changes. On the other hand, our results show no evidence for the selection effect. This is inconsistent with standard menu

¹We would like to thank IRI for making the data available. All estimates and analysis in this paper, based on data provided by IRI, are by the authors and not by IRI.

²As we clarify in Section 8, there is a change in the relative share of measured increases vs decreases in the Calvo (1983) model. However, this is a side-effect of the intensive margin effect, and it is not uniform across levels of mispricing, which would be a requirement of the presence of the gross extensive margin.

cost models (Golosov and Lucas, 2007; Karadi and Reiff, 2019; Bonomo et al., 2019), where selection is generally strong. We argue that second-generation state-dependent price-setting models with random menu costs (Dotsey et al., 1999; Luo and Villar, 2017; Alvarez et al., 2020) or rational inattention (Woodford, 2009) with a linear and flat price-adjustment (hazard) function are broadly consistent with our evidence.

Related literature Our work contributes to the strand of literature that imposes minimal structure to estimate the strength of price selection in microdata. A subset of these papers relies on Caballero and Engel (2007), who derive an indirect measure of the extensive margin effect from the unobservable adjustment hazard and density functions of price gaps. Berger and Vavra (2018) impose flexible functional forms and match unconditional moments of the price-change distribution to estimate hazard functions, finding a sizable extensive-margin effect (see also Petrella et al., 2019). Luo and Villar (2019) match moments conditional on trend inflation rates, and, in contrast, estimate hazard functions which imply weak extensive-margin effects, closer to our results. In contrast to these papers, we generate proxies for the price gaps and report non-parametric estimates of the hazard functions as in Gagnon, López-Salido and Vincent (2012). Akin to Gagnon et al. (2012), we find that the absolute value of the price gap increases the probability of price change and document that the relationship survives when one controls for unobserved heterogeneity and allows longer adjustment lags. Furthermore, as a novel contribution, we show that an increasing hazard exerts its influence through the gross extensive margin and does not imply active selection because the gap does not significantly interact with identified aggregate shocks.

Our paper is complementary to the work of Carvalho and Kryvtsov (2018), which finds no indication for price selection in US supermarket data, in line with our results. They compare the average price of close substitutes to the preset prices of products that eventually adjust and show that this preset-price-relative does not interact with aggregate inflation in a time-series setting. We, instead, construct price gap measures for both adjusted and unadjusted prices and show in a panel-data setting that the price gaps do not interact with identified aggregate shocks. Our work is also complementary to Dedola et al. (2019), who assess, whether *unobserved* product-level shocks generate a selection bias in a Heckman-type model using Danish producer-price data. We, instead, measure selection caused by the interaction of aggregate shocks with *observable* proxies of price gaps.

Our work is also related to the long strand of literature that uses observations about micro-level price behavior to infer the level of monetary non-neutrality in fully specified state-dependent price-setting models (Dotsey, King and Wolman, 1999; Golosov and Lucas, 2007; Gertler and Leahy, 2008; Midrigan, 2011; Alvarez, Bihan and Lippi, 2014; Alvarez, Lippi and Oskolkov, 2020). These papers predominantly use unconditional moments of price changes to calibrate their parameters and deliver predictions that are model-dependent. In contrast, we measure price-level adjustments conditional on identified aggregate shocks, and our predictions can be consistent with multiple theoretical models.³ The strength of monetary non-neutrality in these models depends predominantly on the strength of price selection. Our results favor those theoretical models that find a minimal role for price selection and tend to predict sizable monetary non-neutrality.

The paper is structured as follows. We describe the data in Section 3, and then explain the construction and key features of our baseline price-gap proxy in Section 4. In Section 5, we characterize the dynamic

³Hong, Klepaczy, Pasten and Schoenle (2019) study the informativeness of pricing moments jointly with aggregate price responses to monetary policy shocks, indirectly testing the importance of selection through moments such as kurtosis. Balleer and Zorn (2019) document a small frequency response – similarly to us – and a *decline* in the average absolute size of price changes after an identified monetary policy easing using the German PPI microdata, arguing that these conditional moments are inconsistent with a large selection.

response to aggregate credit shocks, and turn to the evidence on price selection in Section 6. We show that our results are robust to using a non-linear specification in Section 7.1, using an alternative reset-price gap measure in Section 7.2, using producer-price microdata in Section 7.4, using monetary policy shocks in Section 7.5, and across heterogeneous product categories in Section 7.6. Section 8 discusses some implications of the empirical results and Section 9 concludes.

3 Data

We use two datasets in our analysis. Our main dataset is a U.S. supermarket scanner dataset and our secondary dataset is the micro price data underlying the calculation of the U.S. producer-price index at the Bureau of Labor Statistics.

The supermarket scanner dataset is collected by marketing company IRI (Bronnenberg, Kruger and Mela, 2008), and includes weekly total sales (\$) as well as quantities sold at the barcode level for 31 food- and health-care product categories (for example, carbonated beverages) in major supermarkets across 50 U.S. metropolitan areas over the sample 2001-2012. It covers around 15 percent of the consumer expenditure survey.⁴

Given total sale revenues and quantities, we calculate posted prices in the IRI data as $P_{psw} = \frac{TR_{psw}}{Q_{psw}}$, where TR is the total revenue and Q is the quantity sold for each product p in store s in week w . We conduct some straightforward data cleaning. First, we round prices toward the nearest penny, as fractional prices reflect the impact of promotional sales during the week, not actual posted prices.⁵ Second, the product identification numbers of *private-label* products are not unique throughout our sample. In particular, the identification number changes for a subset of products in January 2007, January 2008, and January 2012. When constructing price spells, we follow a conservative approach and assume that all private-label goods were replaced with new goods on these dates. We disregard the price- and expenditure changes for these goods during the three dates. Third, in each year, we only include products that are available over the whole year. This way, we exclude entering and exiting products, which might exhibit idiosyncratic pricing behavior, for example, motivated by learning about the products' demand function at introduction (Argente and Yeh, 2020) or during a clearance sale at exit. These idiosyncratic factors would further reduce selection, so excluding these prices is a conservative choice, which raises the probability of us finding selection in the data. This step makes us drop 17% of the products (17.6% of annual expenditure).

We use the modal-price filter of (Kehoe and Midrigan, 2014) to construct a weekly series of reference prices (P_{psw}^f) that is insulated from the impact of temporary sales. In our subsequent analysis, we concentrate on reference prices but show that our results are robust to the inclusion of sales data. To construct reference prices, we start from a running 13-week two-sided modal price. Then, we iteratively update it to align the time of reference-price change with that of the actual-price change.⁶

We construct monthly observations from our weekly dataset. This helps us concentrate on lower frequency developments in prices, which are more relevant for business cycle fluctuations. The monthly price P_{pst} is defined as the mode of the weekly prices over a calendar month, choosing the highest if there are multiple

⁴The number of observations in our sample is almost 2.65 billion, it covers over 168 thousand unique products in 3187 unique stores in 169 supermarket chains.

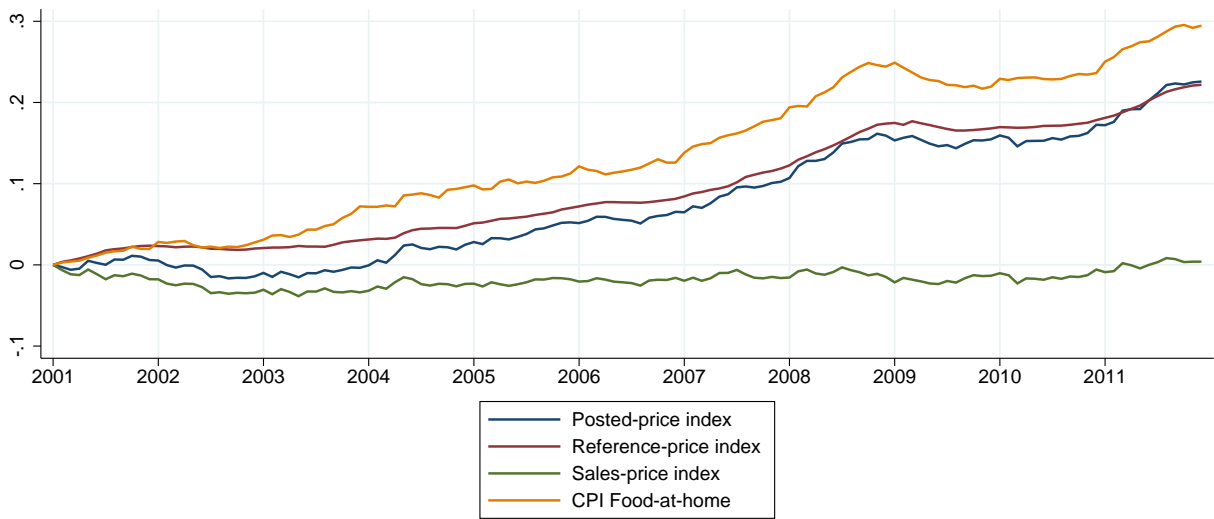
⁵The rounding influences 8.7 percent of the prices in our sample.

⁶Details of the algorithm are described in the Online Appendix.

modes. Picking the mode makes sure we do not create artificial prices and price changes, as could happen if we used averages. We calculate monthly expenditure as the sum of weekly sales.

To cross check the external validity of our scanner data, we calculate an expenditure weighted supermarket price index. Our aim is to create a chain-weighted index similarly to the consumer price index, but utilize the wider and deeper information available in our dataset. In particular, we use annual revenue weights $\omega_{psy} = TR_{psy} / \sum_p \sum_s TR_{psy}$, where subscript y reflects the year of the observation month. We measure inflation for posted- and reference prices ($i = p, f$) as $\pi_t^i = \sum_s \sum_p \omega_{pst} (p_{pst}^i - p_{pst-1}^i)$, where $p_{pst}^i = \log P_{pst}^i$. We define sales-price inflation as the difference between posted- and reference-price inflation $\pi_t^s = \pi_t^p - \pi_t^f$. We seasonally adjust the series using monthly dummies.

Figure 1: Posted, reference and the sales price index



Note: The figure depicts the evolution of posted-, reference- and sales-price indexes over our sample period, and compares them to the food-at-home subindex of the CPI. The reference-price index, which is used in the preceding analysis, tracks the food-at-home index closely at business cycle frequencies.

Figure 1 shows the evolution of the logarithm of our posted-, reference- and sales-price indices ($p_t^i = \log P_t^i, i = p, f, s$) and compares them to the evolution of the food-at-home consumer-price subindex published by the Bureau of Labor Statistics. The (log) indices are normalized to 0 in January 2001 and are constructed as cumulative sums of the inflation rates: $p_t^i = \sum_{j=0}^t \pi_j^i$. We see that our price index somewhat underestimates the overall food-at-home price-level increase over the period.⁷ The main reason is that our index relies on the price-development of existing products and – differently from the CPI – it ignores the endogenous replacement of old, exiting products with new, entering products, which tend to be more expensive and higher quality. More important for our purposes is that the fluctuations of the price index calculated from our sample at business cycle frequencies track those of the food-at-home index closely.

Our second main dataset is the micro price dataset underlying the construction of the U.S. PPI. These data have been described in detail in several papers that analyze the PPI microdata, such as Nakamura

⁷The annualized inflation in posted prices is 1.84%, and in reference prices it is 1.75%. The CPI food-at-home inflation over the same period is 2.7%.

and Steinsson (2008), Goldberg and Hellerstein (2009), Bhattarai and Schoenle (2014), Gorodnichenko and Weber (2016) or Gilchrist et al. (2017). We refer the reader for details of the PPI data to these papers and discuss below some of the key features and how those contribute to the robustness of our results.

The PPI data is a monthly dataset of transaction-based prices. Goods in manufacturing and services are the focus of the PPI data. The dataset contains prices for approximately 28,000 firms and more than 100,000 goods every month, belonging to 540 six-digit NAICS product categories. These NAICS categories are very narrow categories such as Tortilla Manufacturing (NAICS code 311830). Within this setting, goods produced by each firm are uniquely identified according to their “price-determining” characteristics. These characteristics include the type of buyer, the type of market transaction, the method of shipment, the size and units of shipment, the freight type, and the day of the month of the transaction. Goods remain in the data on average for 70 months while firms exist for approximately 7 years. Sales are very rare in PPI data, as documented by Nakamura and Steinsson (2008) so we work with unfiltered data. However, we check for the importance of product substitutions. Whenever we identify a rare product substitution through a so-called base price change, we assume a price has changed.

The main appeal of the PPI microdata is that, i) the data aim at mapping out the “entire marketed output of U.S. producers” in a representative fashion (BLS, 2020). Such producer data rather than retail data is particularly suitable for studying how all kinds of firms set prices and how they react to aggregate shocks; ii) the data range from 1981 to 2015 providing a long time series that spans multiple business cycles and provides variation in monetary policy and credit frictions and iii) the producer price data cover a wide range of sectors in manufacturing but also three-quarters of services in the U.S. economy, lending more general validity to our analysis.

4 A price-gap proxy

In order to measure selection we need a proxy for product-level mispricing, in particular the optimal price a firm would like to charge. We can then assess whether the distance to this optimal price interacts with the aggregate shock, as selection would require. This section describes how we calculate the price gap $x_t = p_t - p_t^*$, which is the distance of the price from its unobserved optimal level. In a wide class of state-dependent models the price gap is indeed the relevant idiosyncratic state variable driving the incentive to adjust prices (Golosov and Lucas, 2007; Alvarez et al., 2014). The further a price is from its optimum, the higher are the incentives to adjust it.

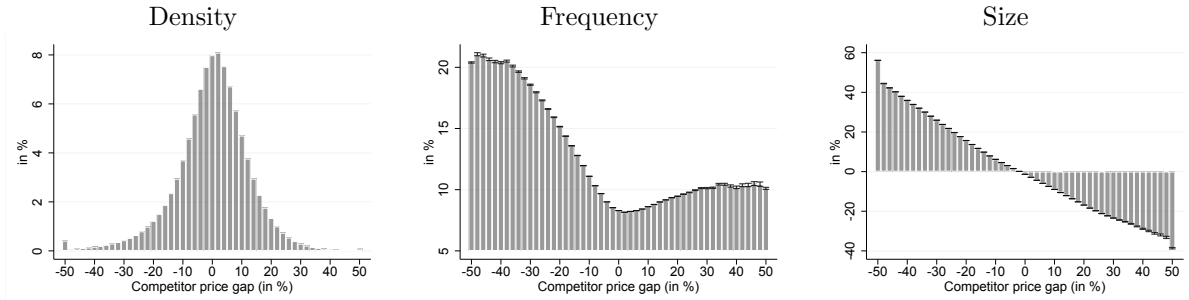
Our baseline price-gap proxy is the competitor-price gap, which measures the distance of the price from a suitably adjusted average price of close competitors. In Section 7.2 and 7.3, we show that our results are robust to using two alternatives: (i) the reset-price gap, which measures the distance of a price from its reset-price (Bils et al., 2012) and (ii) the competitor-reset-price gap, which measures the distance from the average price of those competitors that changed their prices in the particular month. We first describe the details of the construction of the competitors’ price gap. The subsequent section then describes the properties of our main identified shock before finally we are ready to discuss selection, the interaction of gaps, and the shock. A reader mainly interested in selection may directly skip ahead to Section 6.

4.1 Competitor-price gap

One of the primary concerns of firms in their price-setting decisions is figuring out how far the price of an item is from the average price of the competitors. Our data allows us to answer this question at a barcode-level granularity⁸ after controlling for temporary sales and store (and therefore also location) fixed effects. The fixed effects help control for permanent differences between the price-level of stores as a result of differences in demand conditions or costs.

Formally, we formulate the competitor price gap for product p in store s in month t as $\tilde{x}_{pst} = p_{pst}^f - \bar{p}_{pt}^{*f}$, where p_{pst}^f is the logarithm of the reference price and \bar{p}_{pt}^{*f} is the average reference price of the same product across stores. We deal with the persistent heterogeneity across stores (i.e. chains, locations) by subtracting the average store-level gap α_s and reformulate the price gap as $x_{pst} = \tilde{x}_{pst} - \alpha_s$.

Figure 2: Competitor-price gap density, and the frequency and size of subsequent reference price changes as a function of the gap



Note: The panels show the unconditional density and frequency- and size responses of subsequent price change in the baseline supermarket dataset. The density of the competitor-price gaps have fat tails (first panel); the frequency of subsequent price adjustment increases with the absolute size of the gap, but the probability of adjustment is higher for negative gaps than for positive gaps, the probability of adjustment is positive at 0 gap, and stays below 20% even for large (over .5 log points) gaps (second panel), and the size of average subsequent adjustments are close to a (minus) one-on-one relationship with the gap (third panel).

Figure 2 shows the density of the competitor-price gap x_{pst} in our baseline IRI supermarket dataset pooled across products and time, and the probability and the size of price adjustment as a function of the price-gap. The density shows that there is a sizable dispersion of price gaps even after we control for temporary sales and permanent differences in the price level of stores.

The second panel of Figure 2 shows the probability of price changes as a function of our measure of the price gap, the price-adjustment hazard. The probability is increasing with the distance away from zero in a V-shaped pattern that is consistent with state-dependent pricing models. This suggests that our proxy indeed captures an important component of the unobserved theoretical price gap (as argued also by Gagnon et al., 2012). Notably, the price change probability is positive even at a zero lagged price gap and it does not exceed 20 percent even at gaps as low as -50 percent. In Section 7.1 we argue that when one controls for cross-sectional heterogeneity and allows for a gradual price adjustment, the estimated slope of the hazard function increases substantially. Furthermore, the probability is asymmetric: the probability of a change for

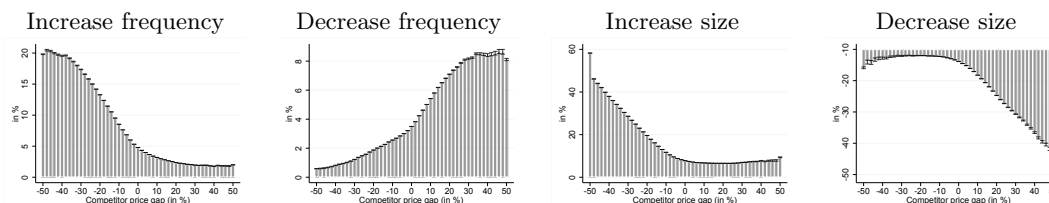
⁸In our sample, the average product is sold in 104 stores in an average month. A large number of products, however, are sold in only a few stores. Table 12 in the appendix shows that our results are robust if we restrict our sample to product-months for which 50 stores are available.

negative gaps is higher than the probability of a change for positive gaps. All the three features (positive hazard at 0, lower than 1 at high gaps, and asymmetry) are in line with the results of [Luo and Villar \(2019\)](#), who estimate flexible hazard functions by matching price-setting moments of U.S. consumer-price inflation.

The third panel of Figure 2 shows that the average size of reference price adjustments conditional on the lagged competitors' price gap. The figure shows that there is a nearly one-to-one relationship between the lagged price gap and the inverse of the reference price change. This confirms again that distance from competitors' prices indeed proxies an important feature of the theoretical price gap, and actual reference price changes aim at closing this gap, on average.

The first two panels of Figure 3 separate the probability of price increases and those of price decreases. The relationship with the gap measures is clearly present in both cases and economically intuitive. The last two panels of Figure 3 separate the positive and the negative price changes and confirm that for negative price gaps the average size of a price increase is in a tight, close to a linear negative relationship with the gap, and conversely for positive gaps the average size of a decrease is in a close to linear negative relationship with the gaps. The figures also reveal that the competitors' price gap is not the only reason firms change their prices, as we observe price increases for positive gaps and price decreases for negative gaps, though their size seems to be mostly unrelated to the price gaps.

Figure 3: Size and frequency as functions of price gap



Note: The panels show the subsequent increase and decrease probabilities, and increase and decrease sizes of price changes as a function of competitor-price gaps. The panels are consistent with the view that negative gaps play a primary role in determining the frequency and size of price increases and, vice versa, positive gaps play a primary role in the frequency and size of price decreases. The non-diminishing role of increase (decrease) frequencies at positive (negative) gaps suggest that the competitor-price gap is one, but not the only determinant of price changes. The average size of price changes is mostly unrelated to the gaps in these regions, conforming with this view.

5 Dynamic impact of credit shocks

This section discusses the properties of our baseline identified shock, the aggregate credit shock. We present the response to a monetary policy shock in our robustness section. A complementary purpose of this section is to establish that in response to a contractionary shock, reference prices adjust primarily through a shift from price increases to price decreases.

We measure financial conditions using the excess bond premium measure, which is a time series of corporate bond spreads purged from the impact of firms' idiosyncratic default probabilities ([Gilchrist and Zakrajšek, 2012](#)). We identify credit shocks with standard exclusion restrictions. To establish the dynamic properties of the credit shock in the economy, we run some simple regressions. We implement the restrictions in a local projection framework by including the contemporaneous values of the excluded variables as controls

(Plagborg-Møller and Wolf, 2019).

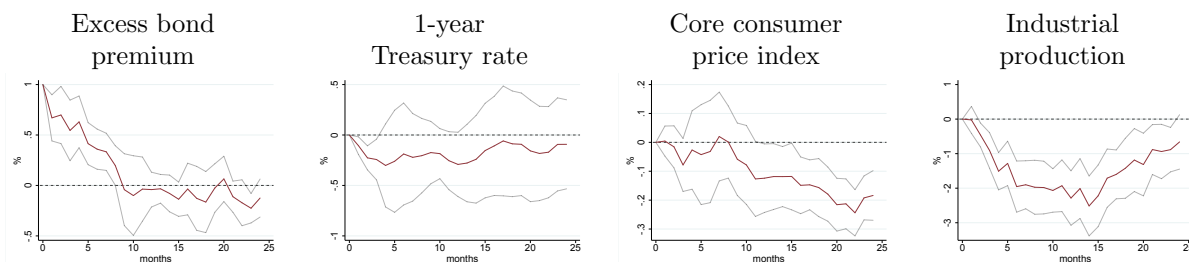
In particular, we run a series ($h = 0, \dots, 24$ months) of regressions of the form:

$$x_{t+h} - x_t = \alpha_h + \beta_h \text{ebp}_t + \Psi_h X_t + \Gamma_h \Phi(L) X_{t-1} + u_{t,h}, \quad (1)$$

where ebp_t denotes the excess bond premium and X_t bundles the one to twelve month lags of the 1-year Treasury rate, the core consumer price index, industrial production and the excess bond premium as well as their contemporaneous values.

The key object of interest is the coefficient β_h on the credit shock. We plot it for all variables of interest for $h = 0, 1, \dots, 24$ along with 95% confidence bands using the Newey and West (1987) heteroscedasticity and autocorrelation-consistent standard errors (Stock and Watson, 2018). Figure 4 shows the impulse responses over the sample from 2001 to 2012. The figures show that even though the credit shock is accompanied by a quick monetary policy easing, it is associated with a sizable decline in industrial production and the core CPI.

Figure 4: Impulse responses of key macroeconomic variables to a credit shock, 2001-2012



Note: The panels show impulse responses to an identified credit shock over the sample 2001-2012 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The figures show that the credit tightening causes a sizable drop in activity and the price index despite sizable monetary policy easing.

We next show that the adjustment of the reference-price level happens primarily through a shift between price increases and price decreases, rather than through changes in the size of average price increases or decreases. This allows us to concentrate on the probability of price adjustment rather than its size in our analysis below. To establish the result, we first decompose the cumulative reference-price inflation between months $t-1$ and $t+h$ into the frequency ($\xi_{t,t+h}$) and the size ($\psi_{t,t+h}$) of price increases and price decreases as follows:

$$p_{t+h} - p_{t-1} = \pi_{t,t+h} = \xi_{t,t+h}^+ \psi_{t,t+h}^+ + \xi_{t,t+h}^- \psi_{t,t+h}^-. \quad (2)$$

Here, the cumulative frequencies of reference-price⁹ increases and decreases between months $t-1$ and $t+h$ are defined as

$$\xi_{t,t+h}^\pm = \sum_i \bar{\omega}_{it,t+h} I_{it,t+h}^\pm, \quad (3)$$

where $I_{it,t+h}^+$ and $I_{it,t+h}^-$ are indicators that take the value 1 if the reference price of item i (a product in a particular store) increased or decreased, respectively, between month t and month $t+h$ and 0 otherwise. The weight $\bar{\omega}_{it,t+h}$ is measured as the average weight of the product between t and $t+h$. The average cumulative

⁹We suppress the superscript f for notational convenience.

size of price increases and decreases are defined as

$$\psi_{t,t+h}^{\pm} = \frac{\sum_i \bar{\omega}_{it,t+h} I_{it,t+h}^{\pm} (p_{it+h} - p_{it-1})}{\xi_{t,t+h}^{\pm}}. \quad (4)$$

Table 1: Average moments

Annualized inflation		Frequency		Reference frequency		Reference size	
Posted	Reference	Posted	Reference	Increase	Decrease	Increase	Decrease
1.84%	1.75%	36.2%	10.8%	6.6%	4.2%	12.5%	-15.1%

Note: The table lists some relevant moments of posted and reference prices. It confirms that posted price changes are very frequent, but they are mostly driven by temporary sales. The probability of reference price changes in a month is around 11%. The size of non-zero reference price change is large, over 10%.

Table 1 lists some of these moments for our supermarket prices. We can observe around 1.8 percent inflation in our sample. The inflation rate in reference prices is very close to that of the posted prices with 1.75% versus 1.84% while average sales price inflation is only 0.05 percent and not shown. Posted prices change very frequently: every month more than one-third of them change. But the majority of price changes are temporary and only 11 percent of reference prices change each month. Out of these reference-price changes, 6.6 percent are increases and 4.2 percent are decreases. Price changes are large when they occur: the average increases are 12.5 percent while the average decreases are 15.1 percent.

Figure 5 summarizes the results from regressions of these three cumulative reference price outcome variables on the credit shock and shows the importance of the frequency responses.^{10,11} We find the following results:

First, as the first panel in the first row of Figure 5 shows, there is a persistent decline in the price level. As the second and the third figures in the first row show, both price increases and price decreases contribute to this decline.

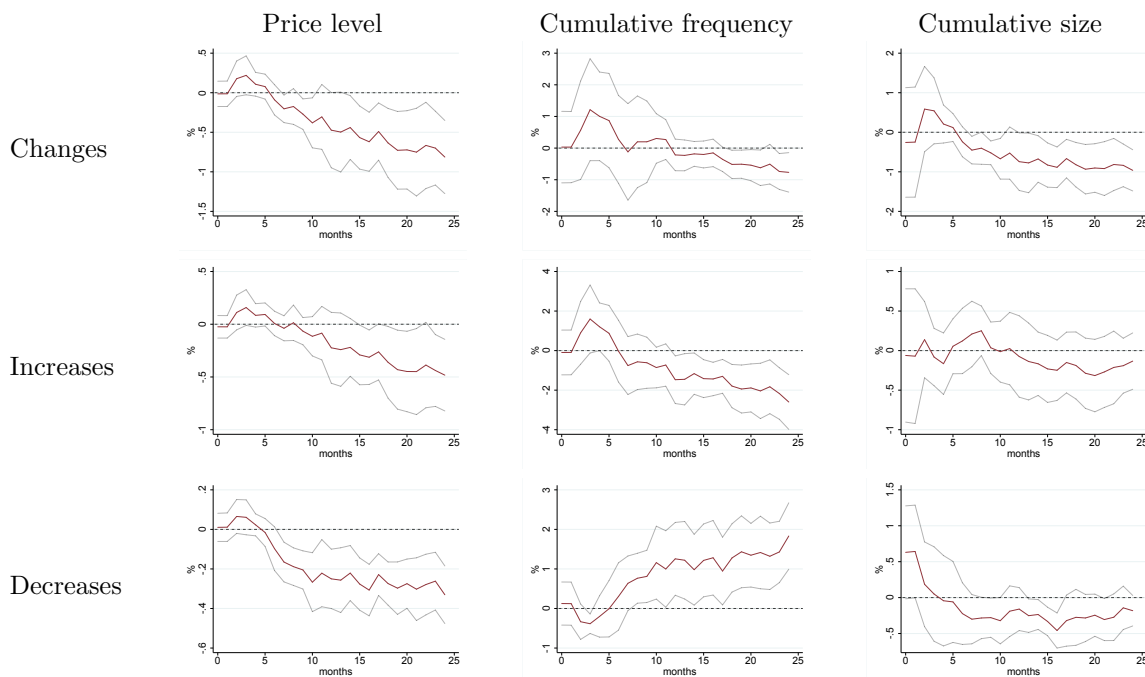
Second, and most importantly, we find a strong shift from price increases to price decreases within a year of the shock. There is a decline in the cumulative frequency of reference-price increases (second panel in the second row), and an increase in the cumulative frequency of price decreases (second panel in the third row). The decline in the cumulative price increases is not much larger than the increase in cumulative price decreases, so the aggregate frequency declines, but only marginally significantly if so at all. Both of these changes contribute to the decline in the price level, and to the reduction in the average size of price changes (third panel in the first row). The average size of price increases stays unchanged and the absolute size of the

¹⁰The local-projection analysis controls for a horizon-dependent constant. Therefore, it effectively assesses the impact of the monetary policy shock on the deviation of the dependent variables from their steady state value. The decomposition of the deviation of cumulative inflation from its steady state becomes

$$\begin{aligned} \pi_{t,t+h} - \bar{\pi}_h = & \left(\xi_{t,t+h}^+ - \bar{\xi}_h^+ \right) \bar{\psi}_h^+ + \left(\psi_{t,t+h}^+ - \bar{\psi}_h^+ \right) \bar{\xi}_h^+ + \left(\psi_{t,t+h}^+ - \bar{\psi}_h^+ \right) \left(\xi_{t,t+h}^+ - \bar{\xi}_h^+ \right) + \\ & \left(\xi_{t,t+h}^- - \bar{\xi}_h^- \right) \bar{\psi}_h^- + \left(\psi_{t,t+h}^- - \bar{\psi}_h^- \right) \bar{\xi}_h^- + \left(\psi_{t,t+h}^- - \bar{\psi}_h^- \right) \left(\xi_{t,t+h}^- - \bar{\xi}_h^- \right). \end{aligned}$$

¹¹We concentrate on reference-prices, because although sales-price inflation responds significantly to the credit shock in our baseline regression (not shown), it is only a feature of the large credit shock of the Great Recession, and it disappears if this period is excluded from the analysis (or only if monetary policy shocks are analysed, which are smaller on average).

Figure 5: Adjustment on the extensive and intensive margins to a credit shock



Note: The figure shows the impulse responses of supermarket reference-price levels and their components to an identified credit shock, and 95% confidence bands using Newey-West standard errors. The panels illustrate how the negative price-level response (first figure in the first row) is predominantly driven by the decline in the price-increase frequency (second panel in the second row) and the increase in the price-decrease frequency (second panel in the third row), while the sizes of price increases and decreases stay mostly constant (third panels in the second and third rows). This adjustment in the gross extensive margin explains the decline in the cumulative size of the price changes (third panel in the first row). The cumulative frequency (net extensive margin) declines, but only marginally significantly (second panel in the first row).

price decreases increases only marginally. This finding motivates our choice to concentrate on the frequency of reference-price changes in the subsequent analysis.

As we make clear in the next section, the shift between the price-increase- and price-decrease frequencies per se does not provide sufficient information about the relative importance of various adjustment channels. Notably, such a shift is present in time-dependent models (Calvo, 1983), which do not permit adjustment on the extensive margin. Rather, the shift in these models is the consequence of the adjustment on the intensive margin, as all adjusting firms reduce the size of their price change relative to a counterfactual scenario without the contractionary aggregate shock. The reduction in size necessarily shifts some firms on the margin to decrease their prices, which would have increased their prices otherwise. Even though the shift changes the relative share of price increases and price decreases, we should not categorize it as adjustment through the (gross) extensive margin. As we show below, extensive-margin adjustments should be assessed within groups determined by the pre-shock size of the price gap. Therefore, to separately identify the channels of adjustment, the analysis needs to explicitly control for the price gap, exactly as we do it in our upcoming empirical implementation.

6 Selection

This section measures the strength of the selection effect. Our main finding is that selection appears to be absent in the data even though both the price gap and the aggregate shocks do have direct effects on price-setting.

To arrive at these results, we estimate the extent to which idiosyncratic price adjustment pressures interact with the credit shock. We approach this estimation in a linear-probability panel-data setting using a random 10% subsample of our IRI supermarket dataset.¹² We control for heterogeneity among items using product-store fixed effects and use two-way clustering across product-category and time. Our main object of interest is the impact of the interaction of the aggregate shock and the price-gap measures on the probability of price adjustment.

We preface our empirical analysis with a detailed explanation of selection. A reader less interested in the technical details of selection may skip right away to its measurement in the data in Section 6.2.

6.1 What is selection?

As is well known for the class of state-dependent pricing models, *which* prices change in response to a nominal shock can be as important for the real effects of the shock as *how many* prices adjust (Goloso and Lucas, 2007). In fact, if prices adjust that are far from their optimal levels, then nominal shocks and consequently monetary policy can be completely neutral even if only a small subset of prices adjust (Caplin and Spulber, 1987). Selection measures how far new adjusters are from their optimal levels when an aggregate shock hits. Figure 6 illustrates the idea of selection graphically, with a description in the text under the figure. The main text presents a mathematically more formal discussion of selection.

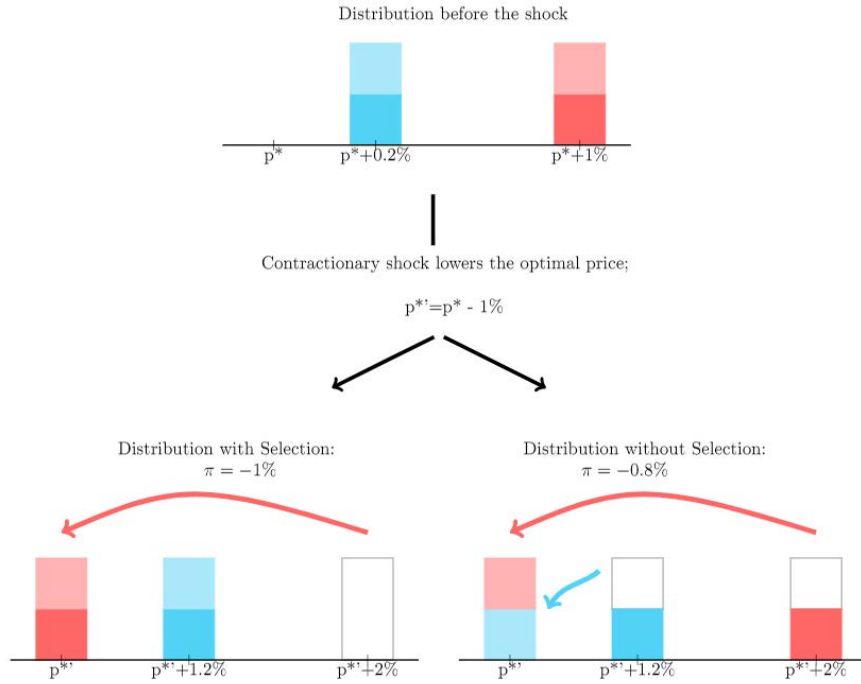
To more formally define selection, we rely on the general price-setting framework of Caballero and Engel (2007). Both the time-dependent Calvo (1983) model, and the strongly state dependent models with an (S,s) adjustment rule are special cases of the framework, as well as a continuum of intermediate cases in line with random menu costs (Dotsey et al., 1999), and rational inattention (Woodford, 2009).

In the framework, there are a continuum of firms each producing a single product i . Firms set the (log nominal) prices of their product (p_{it}) subject to a price-adjustment friction. If these frictions were temporarily absent, the optimal price in period t were p_{it}^* . The optimal price is driven by both aggregate and idiosyncratic factors $p_{it}^* = m_t + \nu_{it}$. For simplicity, we assume that shocks to both m_t and ν_{it} are permanent. The aggregate shock m_t shifts the optimal nominal price of all firms, whereas the idiosyncratic shock ν_{it} affects only firm i . The gap between the price and its optimal value $x_{it} = p_{it} - p_{it}^*$ is the relevant state variable and is sufficient to characterize each firms' price-setting choice.

The firms' price adjustment decision can be described by a generalized hazard function $\Lambda(x)$. The function takes values between 0 and 1, and its value expresses the probability of price adjustment for a firm with a price gap x . The hazard function is constant in the time-dependent Calvo (1983) model: there, the probability of adjustment is independent of the price gap. At the other extreme, in the fixed menu cost model (Caplin and Spulber, 1987; Goloso and Lucas, 2007), the hazard function is a step function that takes the value 0 when the gap is within the inaction band, and 1 outside. Caballero and Engel (2007) shows

¹²We limit the size of our IRI sample because of computational constraints. At the 10 percent level, the estimates are invariant to further sample increases. We subsequently show the robustness of our results in the full PPI sample, to monetary policy shocks, using alternative price gap proxies and non-linear probability models.

Figure 6: An illustrative example



Note: The figure illustrates the mechanism of price selection. Imagine two supermarkets, setting the prices of 4 distinct product types. Half of the prices are much higher (+1%) than their optimal price (p^ , price of a dominant competitor), and the other half are only slightly higher (+0.2%). Then, an aggregate shock – such as a uniform price cut by a dominant competitor – reduces the optimal prices of all goods by 1%. Changing prices is costly, so both supermarkets reset half of their prices, but they follow different price-setting rules: one with selection, another without selection. The one with selection (left panel) chooses to reset the prices that are furthest from their optimal prices. This rule generates an interaction between the aggregate shock and the product-level mispricing: The mispricing substantially amplifies the impact of the shock. Indeed, that supermarkets price index declines by the same amount as the size of the uniform shock even though only half of the products change prices. Such behavior is typical in state-dependent price-setting models, where selection is high, the aggregate price level is flexible and monetary policy is close to neutral. The other supermarket without selection (right panel) chooses to adjust its prices according to a predetermined rule that is independent of the mispricing. For example, it resets a certain number of prices in each aisle. Thus, it ends up picking half of the prices with smaller and half of the prices with larger mispricing. This generates a supermarket price index that is strictly smaller than the uniform shock. This behavior is representative of time-dependent price-setting models, with no selection, sluggish aggregate price level, and non-neutral monetary policy.*

that a continuum of intermediate hazard functions can arise when the menu cost takes value from an iid random variable (from a cumulative density function $G(\omega)$, as in Dotsey et al. (1999)), and when the firm is subject to rational inattention friction as in Woodford (2009) (see also Alvarez et al., 2020).

In this economy, inflation is simply

$$\pi = \int -x\Lambda(x)f(x)dx, \quad (5)$$

where $f(x)$ is the density of price gaps across firms, and we suppressed subscripts for notational convenience.

The expression is intuitive: the inverse price gap is the size of price adjustment and the hazard is its probability. Their product summed across the gap distribution and weighted by the density is the inflation rate.

We further decompose inflation (π) as the sum of price changes (increases, π^+) driven by negative gaps and price changes (decreases, π^-) driven by positive gaps.

$$\pi = \int -x\Lambda(x)f(x)dx = F(0) \int_{x<0} -x\Lambda(x)g(x)dx + (1-F(0)) \int_{x\geq 0} -x\Lambda(x)h(x)dx = F(0)\pi^+ + (1-F(0))\pi^-, \quad (6)$$

where $g(x) = f(x)/\int_{x<0} f(x)dx$, $h(x) = f(x)/\int_{x\geq 0} f(x)dx$ are new functions that are proper densities over negative gaps and positive gaps, respectively. $F(0) = \int_{x<0} f(x)dx$ is the share of products with negative gaps, and, naturally, its complement $(1-F(0))$ measures the share of positive gaps. The decomposition is novel, and we propose it to highlight the relevance of the gross extensive margin, which emerges from the shift in the relative share of price increases versus price decreases as a response to an aggregate shock within the positive and the negative price-gap groups. The gross extensive margin can be active even though the overall frequency of price changes stays constant (net extensive margin is inactive). We argue later that the gross extensive margin is conceptually different from selection.

In the expressions below, we concentrate on the positive price gaps (and price decreases). The behavior of inflation caused by products with negative price gaps is analogous. As in [Costain and Nakov \(2011\)](#), we find it useful to decompose the positive-gap inflation into a mean and a covariance term.

$$\pi_t^- = \int_{x\geq 0} -x\Lambda(x)h(x)dx = (-\bar{x}^-)\bar{\Lambda}^- + \int_{x\geq 0} -x(\Lambda(x) - \bar{\Lambda}^-)h(x)dx, \quad (7)$$

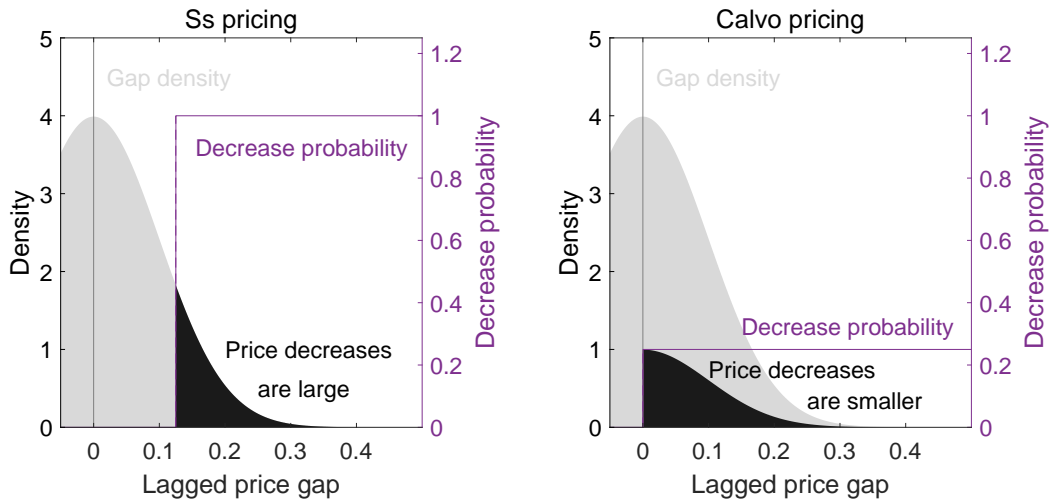
where $\bar{x}^- = \int_{x\geq 0} xh(x)dx$ is the average positive price gap, and $\bar{\Lambda}^- = \int_{x\geq 0} \Lambda(x)h(x)dx$ is the average probability of price adjustment among products with positive price gaps. The decomposition implies that the positive price-gap inflation is the sum of the the product of the inverse average positive price gap and the average decrease frequency and a covariance term between the inverse gap and the probability of price decreases. This decomposition is useful, because the covariance term is a theoretical measure of the extent of state-dependence in price setting. The term disappears in constant-hazard time-dependent models ([Calvo, 1983](#)). It is straightforward to see why: in these models $\Lambda(x)$ is constant, therefore $(\Lambda(x) - \bar{\Lambda}^-) = 0$ for each x . In contrast, the covariance term is necessarily non-zero (negative for positive gaps, and positive for negative gaps) in the relevant state-dependent price setting models with increasing hazard: the larger the absolute value of the gap, the higher the probability of price adjustment.

To illustrate the inflation decomposition more plastically, it is instructive to contrast two pricing frameworks: a state-dependent model such as [Golosov and Lucas \(2007\)](#) with a fixed menu cost and an (S,s) type adjustment and the time-dependent model of [Calvo \(1983\)](#). As discussed above, the generalized hazard function ($\Lambda^-(x)$) is a step function in the former and constant in the latter. First, we discuss the key features of pricing in these two models. Figure 7 focuses on the main objects of interest for the state-dependent-model (left panel) and for the time-dependent model (right panel). The density of price price gaps ($f(x)$) is depicted by the grey shaded area and the density of price decreases as the black shaded area. The latter comes as the product of the adjustment hazard and the density of gaps.

The panels in Figure 7 illustrate the big contrast between the size distribution of price decreases in the

two pricing frameworks: While price changes are on average large in the state-dependent model (left panel), they are on average small in the time-dependent model (right panel). The left panel also illustrates two general features of state-dependent models: The probability of price adjustment increases with the size of the price gap (non-decreasing hazard). In our subsequent empirical implementation, we concentrate on the first feature and measure the extent of state dependence with the tightness of the relationship between the probability of adjustment and the price gap. We next hammer out the meaning of selection in the context of these adjustment mechanisms both formally and graphically.

Figure 7: State-dependence



Note: The figure shows the density of the price-gap, the hazard over positive price gaps and the price-decrease density as a function of gaps in a state-dependent menu cost model (left panel) and a time-dependent model (right panel). The difference between the hazard functions in the models implies large difference in the predicted relationship between the gap and the probability of price decreases.

We measure price flexibility with the impact effect of a marginal aggregate shock as [Caballero and Engel \(2007\)](#)

$$\frac{\partial \pi}{\partial m} = F(0) \frac{\partial \pi^+}{\partial m} + (1 - F(0)) \frac{\partial \pi^-}{\partial m}, \quad (8)$$

where π^+ and π^- are inflation among the group of products with negative and positive pre-shock gaps, respectively.¹³ From equation (7), the marginal impact of an aggregate shock on the positive-gap group becomes

¹³We define positive-gap and negative-gap inflation based on gaps excluding the effect of the aggregate shock to simplify the algebra and stay closer to upcoming our empirical implementation. Using these definitions, we do not need to worry about how the aggregate shock shifts the distribution between the negative- and positive-gap groups. Instead, we can concentrate on how the size of the adjustment and the probability of adjustment changes as a response to the aggregate shock within pre-determined groups. Importantly, the tight link between the gaps and the direction of price changes breaks with this definition: it is not true any more that firms with negative pre-shock gaps necessarily increase their prices after the shock or vice versa. This means that we need to measure pre-shock gaps, as in our empirical implementation below, to separate the two groups, and separating price increases from price decreases post shock will not generate the same groups.

$$\begin{aligned}
\frac{\partial \pi^-}{\partial m} = & \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \underbrace{-\bar{x}^- \frac{\partial \bar{\Lambda}^-}{\partial m}}_{\text{gross extensive}} + \underbrace{\int_{x \geq 0} -x \left(\frac{\partial \Lambda(x)}{\partial m} - \frac{\partial \bar{\Lambda}^-}{\partial m} \right) h(x) dx}_{\text{selection}} = \\
& \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \underbrace{\bar{x}^- \int_{x \geq 0} \Lambda'(x) h(x) dx}_{\text{gross extensive}} + \underbrace{\int_{x \geq 0} x \left(\Lambda'(x) - \int_{x \geq 0} \Lambda'(x) h(x) dx \right) h(x) dx}_{\text{selection}},
\end{aligned} \tag{9}$$

The second equation uses the insight derived in [Caballero and Engel \(2007\)](#) that $\partial \Lambda(x)/\partial m = -\Lambda'(x)$ for all x where Λ is differentiable. The result holds, because the aggregate shock shifts all the gaps uniformly by a unit. Consequently, $\partial \bar{\Lambda}^-/\partial m = -\int_{x \leq 0} \Lambda'(x) h(x) dx$. In words, the marginal change in the average price-decrease frequency equals to the average slope of the hazard function over the positive price gaps. The expression is analogous in the negative-gap group, and the margins of adjustment for the overall inflation rate are the sums of the respective margins in the negative-gap and the positive-gap inflation rates weighted by their respective shares.

This expression decomposes the impact effect into three terms. The first is the intensive-margin effect, which is driven by higher post-shock gaps over the whole distribution. The size of the effect takes a particularly easy form: it equals to the average frequency of price adjustment in the positive gap group ($\bar{\Lambda}^-$). It is intuitive: all firms that would have changed their price anyway, now change their price marginally more by a term $-\partial \bar{\Lambda}^-/\partial m = 1$. Overall, the intensive margin effect is $\bar{\Lambda} = F(0)\bar{\Lambda}^+ + (1 - F(0))\bar{\Lambda}^-$, which equals to the frequency of price changes as in [Caballero and Engel \(2007\)](#). This margin of adjustment is the only one present in the constant-hazard time-dependent models.

The second term is the gross extensive-margin effect. This measures the impact of the shock on the relative share of price changes between the positive-gap and the negative-gap groups. A policy tightening reduces the share of price decreases, and increases the share of price increases. We argue that effect is conceptually different from selection, and it is worthwhile not to bundle them together as the extensive margin effect as in ([Caballero and Engel, 2007](#)). The reason is that measuring the two adjustment channels separately can help us to distinguish between state-dependent pricing models with very different aggregate implications, as we do in our paper. The strength of the effect depends on the average size of the gap and the average impact on the frequency of price changes in the positive-gap group. This latter term is positive within the relevant class of models with the increasing hazard property, so the gross extensive margin effect will amplify the impact of the aggregate shock. The overall impact of the gross extensive margin also leads to amplification: an tightening reduces price increases and increases (in absolute value) price decreases. Formally, $F(0)(-\bar{x}^+) \partial \bar{\Lambda}^+/\partial m + (1 - F(0))(-\bar{x}^-) \partial \bar{\Lambda}^-/\partial m \leq 0$. In constant hazard model $\partial \bar{\Lambda}^-/\partial m = \int_{x \leq 0} \Lambda'(x) h(x) = 0$ is zero, so the gross extensive margin is absent.¹⁴

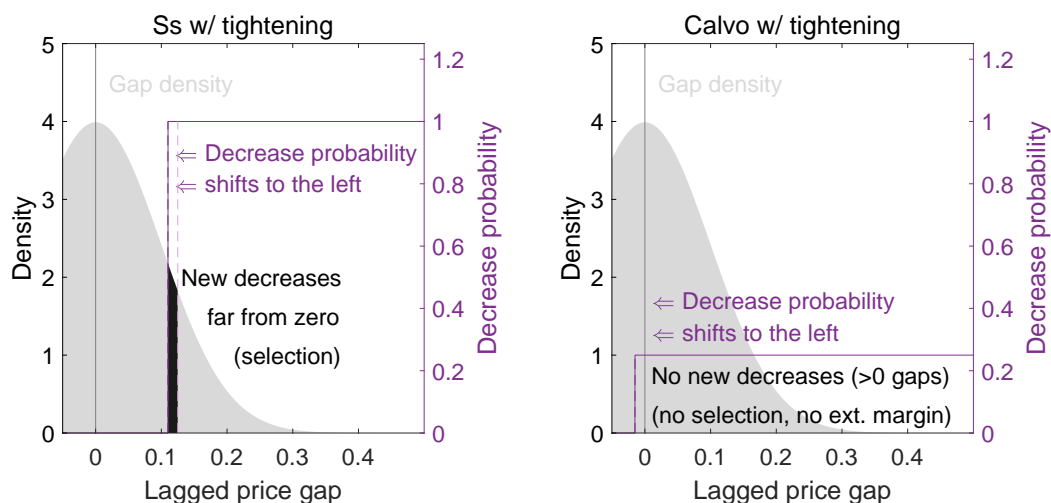
The third term is the selection effect, which will be one of the focus of our empirical exercise. It measures the contribution arising from any shift in the distribution caused by unusual (selected) position of new price

¹⁴Note that the measurement of gross extensive margin effect requires an indicator of the price gap, as in our empirical implementation. The intensive margin channel also shifts the relative share of price increases versus price decreases, so the observable frequency cannot be directly used to assess the strength of the gross extensive margin.

changes. In particular, it is the inverse of the average position of new changes ($\partial\Lambda(x)/\partial m$) controlling for the average frequency change ($\partial\bar{\Lambda}/\partial m$). Notably, for selection what matters is not whether prices that are far from their optimal levels are *regularly* adjusted with higher probability (state dependence), rather it evaluates how large are the *new* price changes that are triggered or canceled (i.e. endogenously selected) by the aggregate shock.

Figure 8 illustrates this point. It depicts a policy tightening, which reduces the optimal price of each firm, raising their after-shock gap, or, equivalently, shifts the adjustment hazard to the left as a function of the pre-shock gap, which we use on our horizontal axis. The black shaded areas on the figures show the new adjusters among the group with positive pre-shock gap.¹⁵ The outcome is very different in the two pricing models. In the time-dependent model (right panel) there are no new adjusters within the relevant group, both the gross extensive margin and the selection effects are missing. In contrast, in the state-dependent model (left panel), both channels are present. First, there are new adjusters within the positive pre-shock gap group. Second, the new adjusters have large gaps, therefore they are far from their optimal prices when the aggregate shock hits. Consequently, when they decrease their prices, they will decrease it by a lot so as to release their accumulated price pressures. The idiosyncratic price pressures, therefore, amplify the impact of the aggregate shock and make the price level flexible, i.e. the selection effect is high.

Figure 8: Selection



Note: The figure shows the density of the before-shock price-gaps, the price-decrease hazards and the distribution of new price changes as a function of gaps in a state-dependent menu cost model (left panel) and a time-dependent model (right panel). The figure illustrates a large difference between selection in the two frameworks: while the distribution of price changes do not change in the time-dependent model, so selection is absent, new price decreases are far from their optima in the menu cost model implying large selection.

This definition of selection motivates our empirical specification, which we detail in the next section. The empirical framework aims at assessing the impact of the *interaction* of the aggregate shock and our price gap measure on the probability of price changes. In other words, it estimates whether prices that are further

¹⁵As discussed above, we examine extensive-margin changes within the pre-shock positive gap group. Therefore, we ignore new price decreases which emerge from products with just-below-zero gaps, which would have been price increases without the shock. We allocate these shifts to the intensive margin effect.

from their optima change with higher probability than average *when* the aggregate shock hits: The selection effect.

6.2 Empirical specification

Our baseline specification is a linear-probability panel-data model. It regresses an indicator variable of price increases and decreases on the interaction of the aggregate shocks and the price gap proxy. We control for the impact of the price gap in normal times (by including the level of the gap directly) and for the impact of the shock on the average frequency (by including the aggregate shock directly), as well as time-dependence (by including the age of the price) and product-store fixed effects.

$$I_{pst,t+h}^{\pm} = \beta_{xih}^{\pm} x_{pst-1} \hat{\text{ebp}}_t + \beta_{xh}^{\pm} x_{pst-1} + \beta_{ih}^{\pm} \text{ebp}_t + \gamma_h^{\pm} T_{pst-1} + \Gamma_h^{\pm} \Phi(L) X_t + \alpha_{psh}^{\pm} + \alpha_{mh}^{\pm} + \varepsilon_{psth}^{\pm}, \quad (10)$$

where $I_{pst,t+h}^{\pm}$ is an indicator of a price increase or respectively a decrease of product p in store s between periods t and $t+h$. x_{pst-1} is the price gap proxy. ebp denotes the level of the excess bond premium as a measure of aggregate financial conditions. $\hat{\text{ebp}}_t$ is the aggregate credit shock, and we explain its construction below. T_{pst} denotes the logarithm of the age of a price, which we measure as the months since the last price change. Finally, $\Gamma_h^{\pm} \Phi(L) X_t$ are aggregate controls. We include contemporaneous values and up to six lags of industrial production as a measure of economic activity, the core consumer price index (CPI) as a measure of the price level and the 1-year treasury yield as a monetary policy indicator.

To show the robustness of our analysis, we present two additional variants of our main specification. First, we analyze a specification with time-fixed effects, without the direct impact of the identified shock. Second, we analyze a specification with separate coefficients for positive and negative gaps.

We identify the credit shock by exclusion restrictions analogously to the local projections shown in section 5. We measure financial conditions by the excess bond premium (Gilchrist and Zakrajšek, 2012). The direct impact of the shock on the probability of price adjustment is measured by β_{ih}^{\pm} , as the regression controls for the impact of the current period activity, prices, and treasury yields. Plagborg-Møller and Wolf (2019) have shown that this is equivalent to an identification with Cholesky factorization ordering the credit shock last. To obtain a credit shock measure ($\hat{\text{ebp}}$) for the interaction term, we appropriately purge the excess bond premium by regressing it on the current month industrial production, the core consumer price level, and the 1-year treasury yield and up to six lags of its own level and the other aggregate variables. Our shock measure is the residual of this regression. As before, this is equivalent to identifying the credit shock using a Cholesky factorization and ordering the credit shock last.¹⁶

All our specifications include the price gap that is lagged by one month $x_{pst-1} = p_{pst-1} - p_{pst-1}^*$. The advantage of the lagged measure is that it is predetermined, therefore unaffected by the contemporaneous aggregate shock $\hat{\text{ebp}}_t$. Their independence simplifies the interpretation of the empirical results. At the same time, the measure ignores the impact of contemporaneous idiosyncratic shocks, which, together with the contemporaneous aggregate shock, affect the contemporaneous optimal price (p_{pst}^*), which, in turn, determine the size of the price adjustment. Most of the variation in the contemporaneous gap, however,

¹⁶The credit shock is a generated regressor. However, as the interaction term is insignificant in our regressions, we do not need to adjust our standard errors (Wooldridge, 2010)

is explained by the lagged gap, and, furthermore, it is in an (inverse) one-to-one relationship with the actual size of the price change. Therefore, if selection were present, the price change probability should be significantly influenced by the interaction of the *lagged* gap and the contemporaneous aggregate shock, as in our specification. The impact of unobserved idiosyncratic shocks on price selection is the focus of the analysis of [Dedola et al. \(2019\)](#), who find some statistically significant, but economically small impact using Danish producer-price data.

Our specifications also control for the direct impact of the price gap β_{xh} on the price adjustment probabilities. The expected sign is negative for price increases (and positive for price decreases), as a larger positive gap indicates the current price is too high (less incentive to increase prices) and a larger negative gap that the current price is too low (more incentive to increase prices). In the baseline specification, we also control for the direct impact of the aggregate shock. Its coefficient (β_{ih}) has a negative expected sign for price increases (and positive expected sign for price decreases, as a credit tightening ($\hat{e}bp_t > 0$) implies fewer increases (and more decreases) and a credit easing ($\hat{e}bp_t < 0$) implies more increases (and fewer decreases).

Our specification also takes into consideration the time-dependent features of price-setting. First, we measure the age of the price as the number of months since the last reference-price change and include its logarithm as a control variable. Additionally, we also include calendar-month fixed effects in our baseline regression, to control for the seasonality of price changes.

The dependent variable in our baseline regression is an indicator variable of a reference-price change in the upcoming 24 months $h = 24$. The horizon is motivated by the peak price-level impact of the shock (see Section 5). On average, the probability of a reference price increase and a decrease at this horizon is 54% and 24%, respectively.

Our main results are twofold. Our first result presents strong evidence for state dependence in the data as the probability of price adjustment increases with the price gap, and the aggregate credit shock affects the probability of price adjustment. Table 2 illustrates these results in our baseline regressions for the competitor price gap proxy in response to a credit shock. Column (1) is the main specification for price increases and Column (4) for price decreases. Columns (2) and (5) show specifications with time fixed effects while (3) and (6) show results for positive and negative gaps separately.

The coefficients show that the effects are generally economically meaningful. First, a one standard deviation credit tightening (30 basis points) reduces the probability of price increases by around 1 percentage point and increases the probability of a price decrease by the same amount. Second, the probability of a price increase for a product at the first quartile of the price-gap distribution relative to the third quartile is 26 percentage points lower; and for the price decrease, it is 23 percentage points higher. These results are inconsistent with the time-dependent model of [Calvo \(1983\)](#), which ignores product-dependent pressures as a factor in price-adjustment probability.

Our second, surprising main finding in the paper is that there is no evidence of selection: Conditional on the aggregate shock, the chance of new price adjustments is independent of our measure of the price gap. None of the coefficients on the interaction term are statistically significantly different from 0. This is inconsistent with menu cost models with high selection ([Golosov and Lucas, 2007](#); [Karadi and Reiff, 2019](#)). As we argue in Section 8, models with random menu costs and weak selection ([Dotsey et al., 1999](#); [Luo and Villar, 2017](#)) are, however, broadly consistent with such evidence.

Table 2: Estimates, scanner data, competitor-price gap, credit shock

	(1)	(2)	(3)	(4)	(5)	(6)
	Price increase ($I_{pst,t+24}^+$)			Price decrease ($I_{pst,t+24}^-$)		
Gap (x_{pst-1})	-1.75*** (0.06)	-1.75*** (0.06)		1.55*** (0.06)	1.55*** (0.06)	
Shock (ebp_t)	-0.03*** (0.01)		-0.04*** (0.01)	0.03*** (0.01)		0.03*** (0.01)
Selection ($x_{pst-1} \hat{ebp}_t$)	-0.00 (0.04)	-0.00 (0.04)		0.01 (0.05)	0.01 (0.04)	
Age (T_{pst-1})	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.00** (0.00)	0.01*** (0.00)	0.01*** (0.00)
Pos. gap (x_{pst-1}^+)			-2.26*** (0.13)			2.29*** (0.10)
Neg. gap (x_{pst-1}^-)			-1.44*** (0.07)			1.10*** (0.06)
Pos. sel. ($x_{pst-1}^+ \hat{ebp}$)			0.04 (0.06)			-0.04 (0.05)
Neg. sel. ($x_{pst-1}^- \hat{ebp}$)			-0.03 (0.06)			0.04 (0.07)
Product x store FE	✓	✓	✓	✓	✓	✓
Calendar-month FE	✓	✗	✓	✓	✗	✓
Time FE	✗	✓	✗	✗	✓	✗
N	16.1M	16.1M	16.1M	16.1M	16.1M	16.1M
within R^2	18.5%	16.6%	18.9%	17.3%	16.4%	18.2%

Note: The table shows estimation results from a linear-probability panel model using scanner data. The regressions are run separately using an indicator with value 1 for reference-price increases (columns 1-3) and an indicator with value 1 for reference-price decreases (columns 4-6). The regressions include product-store and calendar-month fixed effects, control for the age of the price (time spent since last change), and use standard errors with two-way clustering. The baseline regressions (columns 1 and 4) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities, which is evidence for state-dependence, their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with time-fixed effects (columns 2 and 5) and a specification with separate coefficients for positive and negative gaps (columns 3 and 6).

*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

7 Robustness

In this section, we show that our results remain robust when we use a non-linear specification, when we consider the reset-price gap, when we use producer price (PPI) microdata instead of retail data, and when we consider a monetary instead of a credit shock. Selection remains absent in all cases.

7.1 Non-linearity

A potential concern may be found in the linear relationship imposed by our baseline specification between the price-change probabilities and the price gap, as well as the price gap conditional on an aggregate shock size. If these relationships were non-linear, our rejection of the presence of selection could be a rejection of the linearity assumption. This section conducts a robustness check using a non-parametric approach and rules out such concerns.

Our main non-parametric approach begins by assigning price gaps into 5 approximately equal-sized¹⁷ bins. The bin that serves as a reference group includes items with small price gaps, in particular price gaps between $-1\% \leq x_{pst-1} < 1\%$. We compare the price-setting behavior of the reference bin to two bins with medium sized negative ($-4\% \leq x_{pst-1} < -1\%$) and positive ($1\% \leq x_{pst-1} < 4\%$) gaps and two bins with large negative ($x_{pst-1} < -4\%$) and positive ($4\% \leq x_{pst-1}$) gaps.

As before, we run separate regressions for positive and negative price changes with item- and time fixed effects and two-way clustering. Instead of the size of the price gap, we now include bin-dummies both as direct regressors and in the interaction terms. We exclude the reference group from among the direct regressors and the interaction terms, so the estimated coefficients compare the average impact of the groups on the price-change probabilities directly, and conditional on an aggregate shock relative to the reference group.

Our results are robust to this non-parametric specification as Table 3 shows. In particular, the probability of price change increases with the average absolute size of the price gaps in the bin, but the interaction terms are not significantly different from the interaction term of the reference bin, implying no detectable selection.

As an additional robustness check, Figure 9 depicts the estimated coefficients of an analogous specification with 15 equal-sized bins instead of 5.¹⁸ Again, the specification includes item- and time fixed effects, and the standard errors are clustered by category and time. The red lines show the impact on price-increase probabilities and price-decrease probabilities of the price-gap groups relative to the group with small gaps. The panels show that the size of the gap has a large impact on the probability of price adjustment in the next 2 years: For example, if the negative gap exceeds 20 percent, the probability of a future price decrease goes up by almost 50 percentage points. The figure also shows that the relationship between the gaps and the adjustment probabilities is monotone and close to linear, justifying our linear baseline specification. The blue lines show the impact of the gap-bin-credit-shock interaction terms on the price-change probabilities. The panels show that the additional impact brought about by the aggregate shock does not significantly vary with the gap, so selection is undetectable. The panels also show that the insignificant results in our baseline specifications are not the consequence of us imposing linearity as there is no detectable interaction between the aggregate shock and the gaps at any gap levels.

7.1.1 A multinomial-probit and an ordered-probit specification

Our baseline results rely on a linear-probability specification. The framework is suitable to assess the average impact of explanatory variables on the binary dependent variables (Wooldridge, 2010), especially if the actual

¹⁷The bins are not equal-sized to maintain symmetry. In particular, we generate 5 equal-sized bins with negative gaps and 5 equal-sized bins with positive gaps. Then we merge the largest negative and the smallest positive bins and each consecutive negative and positive bin. Because the price gap distribution is approximately symmetric, we obtain 5 approximately equal-sized bins.

¹⁸The horizontal axis of the figures show the threshold competitor-price gaps of the bins.

Table 3: Non-linear specification, 5 groups

	(1)	(2)
	$I_{pst,t+24}^+$	$I_{pst,t+24}^-$
Large negative gap ($x_{pst-1} << 0$)	0.35*** (0.01)	-0.28*** (0.01)
Medium negative gap ($x_{pst-1} < 0$)	0.15*** (0.01)	-0.13*** (0.01)
Medium positive gap ($x_{pst-1} > 0$)	-0.15*** (0.01)	0.13*** (0.01)
Large positive gap ($x_{pst-1} >> 0$)	-0.33*** (0.01)	0.32*** (0.01)
Large negative selection ($x_{pst-1} << 0$) $\hat{e}bp_t$	0.01 (0.01)	-0.01 (0.01)
Medium negative selection ($x_{pst-1} < 0$) $\hat{e}bp_t$	0.00 (0.01)	-0.00 (0.00)
Medium positive selection ($x_{pst-1} > 0$) $\hat{e}bp_t$	0.00 (0.00)	-0.00 (0.00)
Large positive selection ($x_{pst-1} >> 0$) $\hat{e}bp_t$	0.01 (0.01)	-0.01 (0.01)
Age (T_{pst-1})	0.02*** (0.00)	0.01*** (0.00)
Product x store FE	✓	✓
Time FE	✓	✓
N	16.1M	16.1M
within R^2	16.6%	16.5%

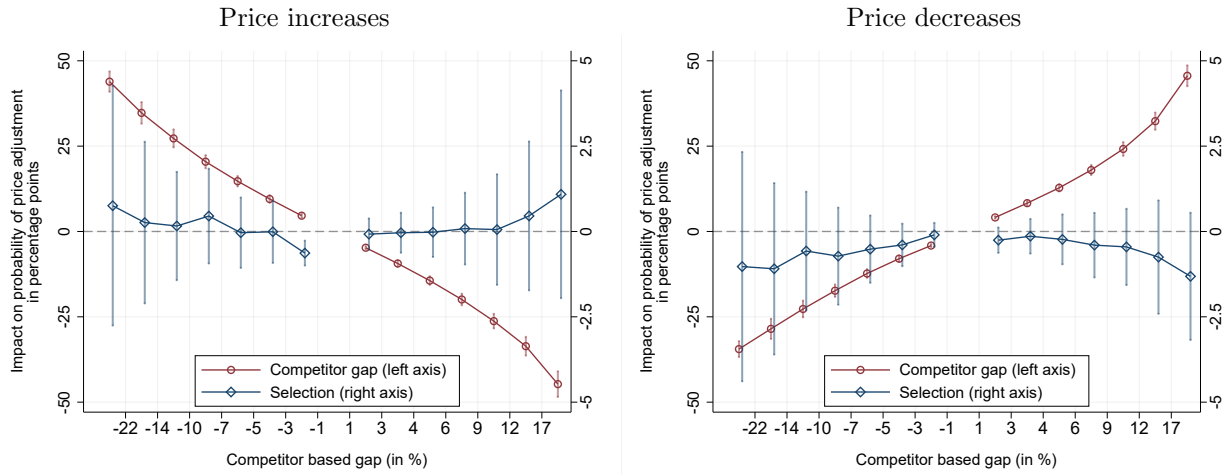
Note: The table shows estimation results from the linear-probability panel model using scanner data. The regressions are run separately on an indicator with value 1 for price increases (columns 1) and an indicator with value 1 for price decreases (columns 2). The regressions include product-store and time fixed effects and calculates standard errors with two-way clustering. The table shows that even though groups with larger absolute gaps increase the price-change probability, the interaction-term with the aggregate shock always stay insignificantly different from 0. The results confirm that selection is undetectable in the data.

*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

relationship for common values of the dependent variables is indeed close to linear as our evidence just above indicates. However, the linear-probability specification does not take into account some general restrictions, which apply to probabilities as for example the true sum of the probability of price increase, price decrease and no price change equals one, so the sum of the probabilities of a price increase and a price decrease at a given point in the state space cannot exceed one.

In this section, we show that our results are robust to specifications which explicitly take into account

Figure 9: Non-linear specification, 15 groups



Note: The figures depict the impact on price-increase (left panel) and price-decrease (right panel) probabilities the estimated differential impact of 15 equal-sized groups with various price gaps (red lines) and price-gap-credit-shock interaction terms (blue lines) relative to the group with close to 0 gaps. The specification includes item- and time fixed effects, and the standard errors are clustered along categories and time. The horizontal lines show 95-percent confidence bands. The figure shows that while the gap in itself significantly influences the price-adjustment probability, it has no significant interaction with the aggregate shock at any gap sizes, so selection is not detectable.

such restrictions on our discrete dependent variables. In particular, we run a multinomial- and an ordered probit specification. The multinomial specification assumes that the firm compares the relative advantages of a price increase, no price change, and a price decrease. The specification allows the explanatory variables to differently influence the gains of price increases relative to no change, compared to the gains from price decrease relative to the no-change decision. Consequently, we report separate coefficients for price increases versus price decreases, similarly to our baseline specification. In contrast, the ordered-probit specification assumes the three decisions are ordered on a single line. The coefficients then relate changes in variables to movement along this line and the position on the line determines the likelihood of each outcome. As such, there is just one set of coefficients determining the probability of each outcome, whereas the multinomial model has outcome-specific coefficients.

We cannot include fixed effects in either specification due to the incidental parameters problem of non-linear models (Hahn and Newey, 2004), nor can we turn to conditional maximum likelihood models (Chamberlain, 1980; Pfaff, 2014) as current implementations are computationally infeasible beyond a few thousand observations and our sample contains millions. Instead, we control for cross-sectional heterogeneity in the probability of price adjustments by including the average frequency of price increases and price decreases of close substitutes (see below),¹⁹ and standardize price gaps at the item level (i.e. we subtract the mean and divide by the standard deviation).

We measure the average frequency of price increases and decreases of close substitutes as follows. We calculate reference price increases and decreases of each product p in a particular market M , but we ignore

¹⁹These variables would be crowded out by the item fixed effects in our baseline specifications.

own changes in store s . Formally, the frequency of increases (decreases are analogous) is

$$\xi_{psM}^+ = \sum_{t=1}^T \frac{\sum_{z \in M \setminus s} \omega_{pzt} I_{pzt}^+}{\sum_{z \in M \setminus s} \omega_{pzt}}, \quad (11)$$

where $M \setminus s$ is the set of stores in market M except store s , ω_{pzt} is the annual expenditure weights of product p in store z in month t , and I_{pzt}^+ is an indicator function taking the value 1 if there is a price increase for product p in store z at month t .

Table 4: Multinomial and ordered probit estimates, scanner data, competitors' price gap, credit shock

	(1)	(2)	(3)
	Multinomial probit		Ordered probit
	Incr. $\left(I_{pst,t+24}^+\right)$	Decr. $\left(I_{pst,t+24}^-\right)$	Change $\left(I_{pst,t+24}\right)$
Gap (x_{pst-1})	-3.15*** (0.02)	3.37*** (0.04)	-4.24*** (0.04)
Shock (ebp_t)	-0.11*** (0.03)	0.05*** (0.01)	-0.10*** (0.02)
Selection $(x_{pst-1} \hat{ebp}_t)$	-0.05 (0.06)	-0.21** (0.11)	0.04 (0.12)
Age (T_{pst-1})	0.01* (0.00)	-0.03*** (0.00)	0.02*** (0.00)
Freq. incr. (ξ_{psM}^+)	5.17*** (0.03)	2.91*** (0.02)	1.79*** (0.03)
Freq. decr. (ξ_{psM}^-)	3.02*** (0.03)	5.84*** (0.05)	-1.33*** (0.04)
Product x store FE	✗	✗	✗
Calendar-month FE	✓	✓	✓
Time FE	✗	✗	✗
N	16.1M	16.1M	14.3M

*Note: The table shows estimation results from multinomial probit and ordered probit models using scanner data. The regressions consider 3 choices (price increase, no change, decrease). The regressions control for the age (time spent since last change) of the price, the frequency of price increases and price decreases among competitor prices in the market (excluding own change), and use standard errors with clustering across product-stores. The results are robust: the price gap and the aggregate shock significantly influence the price-change probabilities, but their interaction is never consistent with a significant selection effect. The interaction term is mostly insignificant, except for price decreases in the multinomial probit model. In this case, the interaction term is significantly different from 0, but has a counterintuitive sign: a higher positive gap and aggregate credit tightening imply fewer (not more) price decreases. Consequently, this estimate is also inconsistent with selection. Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

Table 4 presents the results. Columns 1 and 2 show the estimated coefficients from the multinomial probit specification and column 3 the estimated coefficients from the ordered probit specification. We find

that the price gap and the aggregate shock significantly influence the price-change probabilities, but their interaction is never consistent with a significant selection effect. The interaction term is mostly insignificant, except for price decreases in the multinomial probit model. In this case, however, while the interaction term is significantly different from 0, it has a counterintuitive sign: A higher positive gap and aggregate credit tightening imply less (not more) price decreases. Consequently, this estimate is also inconsistent with selection.

7.2 Reset-price gap

In this section, we show the robustness of our result to using reset-price- rather than competitor-price gaps. We define the reset price as the counterfactual price a firm would charge if price-adjustment frictions were temporarily absent. In a wide class of state-dependent price-setting models, the reset price is the key product-level attractor that drives the probability and the size of price adjustment.

Bils et al. (2012) offers an iterative algorithm to obtain a proxy for the reset price. The algorithm relies on two key assumptions. First, when a firm adjusts its price, it sets it at the reset price. Second, when the firm does not adjust its price, its reset price evolves with the reset-price inflation of its close competitors, which can be measured from the changes in the reset prices of the adjusting competitors. Formally, the logarithm of the reset-price of item i in month t is

$$p_{it}^* = \begin{cases} p_{it} & \text{if } I_{it} = 1 \\ p_{it-1}^* + \pi_{ct}^* & \text{otherwise} \end{cases} \quad (12)$$

where I_{it} is an indicator function that takes the value of 1 if the price of product i changes in the month t , and π_{ct}^* is the reset-price inflation in the product's category c . Reset-price inflation, in turn, is given by

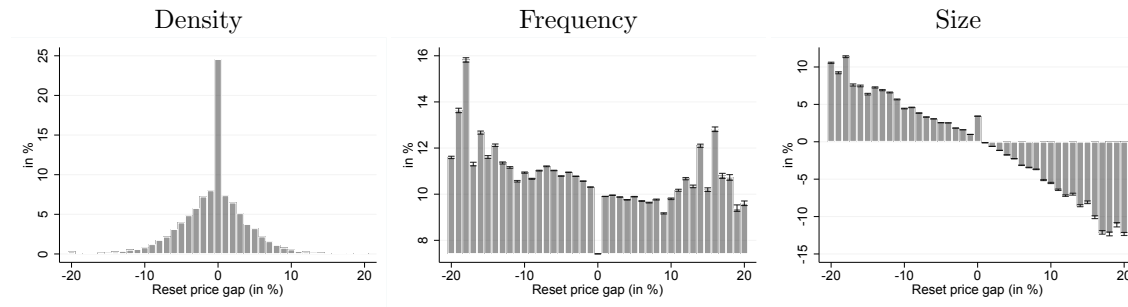
$$\pi_{ct}^* = \sum_{i \in c} \frac{\omega_{it} I_{it} (p_{it}^* - p_{it-1}^*)}{\sum_{i \in c} \omega_{it} I_{it}}, \quad (13)$$

where ω_{it} denotes the expenditure weight of item i .

The reset-price gap is simply the distance of the logarithm of the price from the logarithm of its reset price $x_{it} = p_{it} - \pi_{it}^*$. We assess whether these reset price gaps truly proxy actual price gaps by looking at the average size of price changes conditional on the reset-price gap in the previous month. If the proxy is good, there should be a strong correlation between the size of the price change and the lagged reset price gap. The left panel of Figure 10 shows the histogram of the reset price gaps. The distribution has a negative median, fat tails, and is left-skewed. The right panel of Figure 10 shows the average size of subsequent price adjustment for each reset-price gap bins. The relationship is clearly negative.

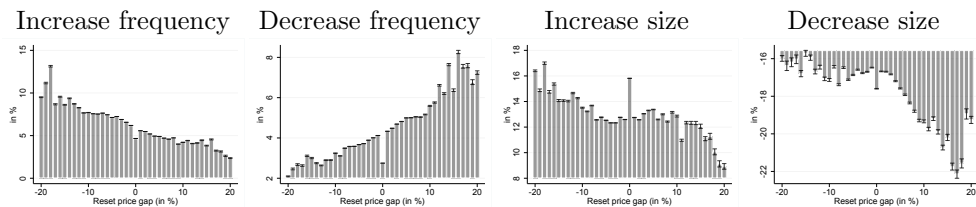
Figure 11 presents the size and frequency histograms separately for price increases and decreases. The tight negative relationship between the reset-price gap and the moments are salient in the graphs. At the same time, the panels of the figure reveal that the reset-price gap is not the sole factor driving firms price-setting decisions, because we see a non-negligible fraction of firms increasing their prices even though their reset-price gap is positive, and conversely, decrease them even though the gap is negative. Furthermore, the probability of price increases drops down from its peaks after the price gaps become lower than 20 percent (see the bottom left panel on Figure 11).

Figure 10: Reset-price gap density and the subsequent frequency and size of price changes as a function of the gap, scanner data



Note: The figures show the unconditional density of the reset-price gap and frequency and size of subsequent price changes as a function of the gap in the baseline supermarket dataset. The figures show that the density of the reset-price gaps has fat tails (first panel); the frequency of subsequent price adjustment increases with the absolute size of the gap (the second panel, see also next figure); and the size of average subsequent adjustments are inversely related with the gap (third panel).

Figure 11: Frequency and size of subsequent price changes as a function of reset-price gap



Note: The figures show the subsequent increase and decrease probabilities and increase and decrease sizes as a function of reset-price gaps. The figures show the negative relationship between gaps and the increase-frequency (except for large negative gaps) and the positive relationship between the gaps and the decrease-frequency. The figures also confirm a similar relationship with the average size of subsequent price changes, with the relationship losing its monotonicity for positive gaps and increase sizes and negative gaps with decrease sizes.

Table 5 shows that we obtain qualitatively similar effects using reset-price gaps. A one standard deviation credit tightening (30 basis points) reduces the probability of price increases by around 1 percentage point and increases the probability of a price decrease by the same amount. The probabilities are less sensitive to the reset price gaps than to our baseline competitor-price gap but nevertheless qualitatively similar as the probability of a price increase for a product at the first quartile relative to the third quartile is 2.25 percentage points lower, and for the price decreases it is 1.6 percentage points higher. The interaction terms remain insignificantly different from zero, continuing to imply the absence of any selection effect.

7.3 Competitor-reset-price gap

In this section, we show that our results are robust to using a third gap measure, which we call the competitor-reset-price. To obtain a proxy for the optimal price, this measure focuses on those competitors' prices, which changed their price in a particular month. In a wide class of theoretical models, the price that competing stores set indeed reveal the unobserved optimal price (see, for example, Dotsey et al., 1999).

Table 5: Estimates, scanner data, reset-price gap, credit shock

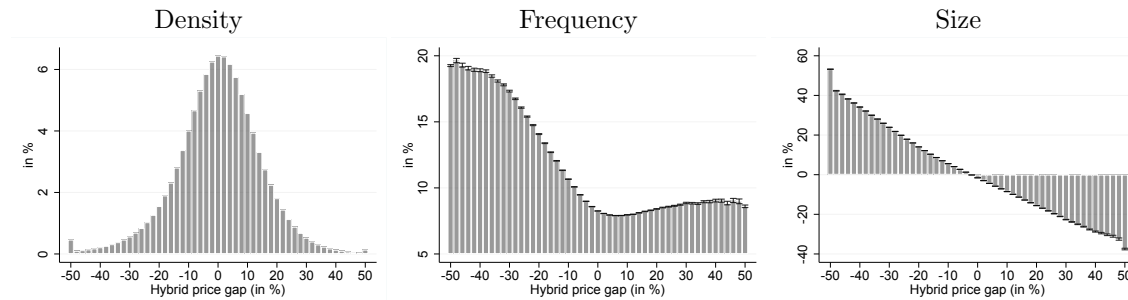
	(1)	(2)	(3)	(4)	(5)	(6)
	Price increases ($I_{pst,t+24}^+$)			Price decreases ($I_{pst,t+24}^-$)		
Gap (x_{pst-1})	-0.45*** (0.07)	-0.48*** (0.06)		0.34*** (0.04)	0.37*** (0.04)	
Shock (\hat{ebp}_t)	-0.04*** (0.01)		-0.04*** (0.01)	0.03*** (0.01)		0.03*** (0.01)
Selection ($x_{pst-1}\hat{ebp}_t$)	-0.14 (0.14)	-0.13 (0.12)		0.12 (0.12)	0.14 (0.10)	
Age (T_{pst-1})	0.01*** (0.00)	0.01*** (0.00)	0.01*** (0.00)	0.01*** (0.00)	0.02*** (0.00)	0.01*** (0.00)
Positive gap (x_{pst-1}^+)			-0.39*** (0.07)			0.33*** (0.07)
Negative gap (x_{pst-1}^-)			-0.49*** (0.13)			0.35*** (0.07)
Pos. sel. ($x_{pst-1}^+\hat{ebp}_t$)			0.11 (0.15)			-0.03 (0.13)
Neg. sel. ($x_{pst-1}^-\hat{ebp}_t$)			-0.27** (0.13)			0.21* (0.12)
N	16.1M	16.1M	16.1M	16.1M	16.1M	16.1M
within R^2	2.6%	0.3%	2.6%	1.3%	0.3%	1.3%

Note: The table shows estimation results from the linear-probability panel model using scanner data. The regressions are run separately on an indicator with value 1 for price increases (columns 1-3) and an indicator with value 1 for price decreases (columns 4-6). The regressions include product-store and calendar-month fixed effects, control for the age of the price (months since last change), and calculate standard errors with two-way clustering. The baseline regressions (columns 1 and 4) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with time-fixed effects (columns 2 and 5, without calendar-month FE) and a specification with separate coefficients for positive and negative gaps (columns 3 and 6). Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.

Formally, we formulate the competitor-reset-price gap x_{pst}^r for product p in store s in month t in two steps. First, we calculate an unadjusted gap as $\tilde{x}_{pst}^r = p_{pst}^f - \bar{p}_{pt}^{*fr}$, where p_{pst}^f is the logarithm of the reference price and \bar{p}_{pt}^{*fr} is the average reference-reset-price of the same product across those stores that changed their prices in month t . Second, analogously to the competitor-gap, we deal with the persistent heterogeneity across stores (i.e. chains, locations) by subtracting the average store-level gap α_s^r and reformulate the price gap as $x_{pst}^r = \tilde{x}_{pst}^r - \alpha_s^r$. We restrict attention to those product-months, where the number of competing stores that changed their prices in the particular month exceed fifty.

Figure 12 shows the density of the competitor-reset-price gap x_{pst} in our baseline IRi supermarket dataset pooled across products and time, and the probability and the size of price adjustment as a function of the price-gap. The figures are qualitatively identical and quantitatively close to the analogous figures based on the competitor-price gap (i.e. average of the price in competing stores, not only those competing stores, which changed their prices).

Figure 12: Competitor-reset-price gap density, and the frequency and size of subsequent reference price changes as a function of the gap



Note: The panels show the unconditional density of the competitor-reset-price gap and frequency- and size responses of subsequent price change in the baseline supermarket dataset. As with the baseline competitor-price gap (Figure 2), the density of the competitor-reset-price gaps have fat tails (first panel); the frequency of subsequent price adjustment increases with the absolute size of the gap, and the size of average subsequent adjustments are close to a (minus) one-on-one relationship with the gap (third panel).

Table 6 compares the baseline regressions with competitor-price gap to the same regressions using the competitor-reset-price gaps. The results stay robust. The magnitude of the estimated coefficients change somewhat, but neither the sign nor the significance of the estimated impacts change relative to the baseline.

7.4 Producer-price microdata

Our results also hold when we use producer price microdata rather than the (retail) IRi the scanner data. We obtain producer price data from the microdata underlying the U.S. producer price index. It is less granular than the IRi microdata which reduces the quality of the relevant price-setting proxies, but in turn spans a much longer time period (1981-2012) and a wider set of sectors than the IRi microdata.

Consistent with the baseline results we find that the impulse response of the PPI to the credit shock is economically intuitive. Despite a monetary easing, industrial production and prices fall following a credit shock. We estimate specification (3) with the PPI on the left-hand side to show these results (see Figure 13).

Moreover, our price gap proxy is informative in a similar fashion as in the IRi retail data, establishing the general validity of our proxy measures. Figures 14 show the density of competitor-price gaps, together with the relation of the gaps and the frequency and size of subsequent price changes. They show that there is a clear negative relationship between the size of the price changes and the proxies for the gaps. The proxies clearly capture a relevant part of the theoretical price gaps, even though the proxies are not as well measured as in the more granular scanner data. The frequency of price changes increases with the price gap as the gap becomes sufficiently large, even though it declines for small price gaps. This might reflect the role of heterogeneity in price-setting frequencies across sectors, as well as the presence of measurement error.

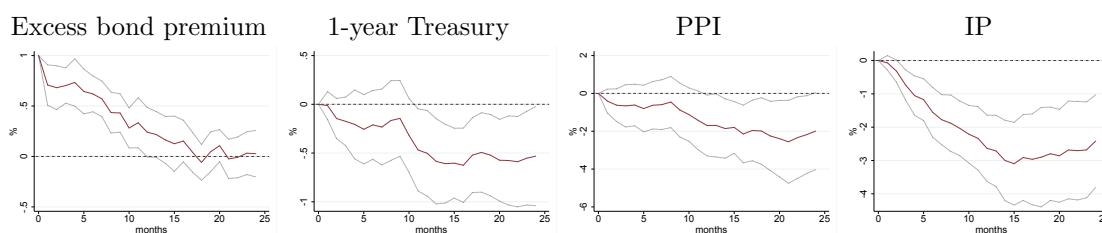
²⁰The sample size is smaller, because the competitor-reset-price gap imposes a stronger requirement on the admissible product-months than the baseline regressions. In the baseline, a product-month is admissible if there are at least 50 competitors selling the same product in the particular month. For competitor-reset-price gaps, we require the presence of at least 50 competitors selling the same good that change their reference prices in the particular month.

Table 6: Robustness to using competitor-reset gap, scanner data, credit shock

	(1)	(2)	(3)	(4)
	Increases ($I_{pst,t+24}^+$)		Decreases ($I_{pst,t+24}^-$)	
	Baseline	Competitor-reset-gap	Baseline	Competitor-reset-gap
Gap (x_{pst-1})	-1.75*** (0.06)	-1.29*** (0.04)	1.55*** (0.06)	1.19*** (0.06)
Shock (ebp_t)	-0.03*** (0.01)	-0.05*** (0.01)	0.03*** (0.01)	0.04*** (0.01)
Selection ($x_{pst-1}\hat{ebp}_t$)	-0.00 (0.04)	-0.01 (0.05)	0.01 (0.05)	0.00 (0.06)
Age (T_{pst-1})	0.02*** (0.00)	0.02*** (0.00)	0.00** (0.00)	0.00 (0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✓	✓	✓
Time FE	✗	✗	✗	✗
N	16.1M	9.3M	16.1M	9.3M
Within R^2	18.5%	15.2%	17.3%	14.5%

Note: The table shows estimation results from the linear-probability panel model using scanner data with the competitor-price gap and the competitor-reset-price gap. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions control for the age (time spent since last change) of the price and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with the competitor-reset-price gap (columns 2 and 4). Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.

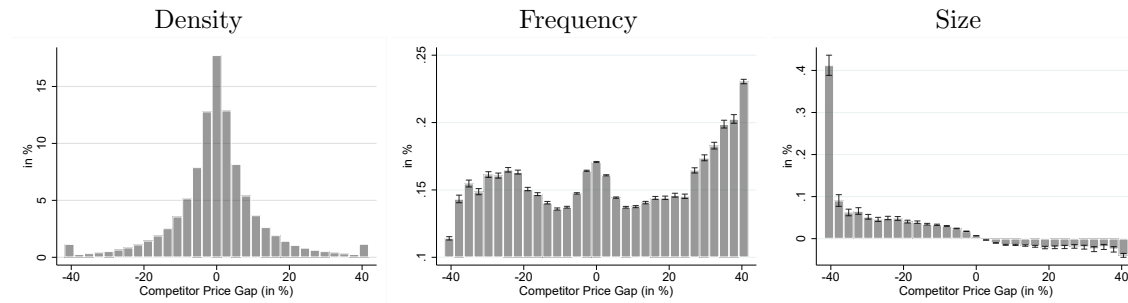
Figure 13: Impulse responses of key macroeconomic variables to a credit shock, 1985-2015



Note: The figures show impulse responses to an identified credit shock over the PPI sample 1985-2015 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The figures show that the credit tightening causes a sizable drop in activity and the price index despite sizable monetary policy easing.

Table 7 shows the linear-probability panel estimations for our baseline and time-fixed effect specifications using the PPI microdata for the competitor-price gap. We find results that are entirely consistent with the baseline. There is evidence for state dependence because larger gaps and the aggregate shocks change the probabilities of price adjustment. At the same time, there is no evidence of selection: conditional on the

Figure 14: Competitor-price gap density and the subsequent frequency and size of price changes as a function of the gap, PPI data



Note: The figures show the unconditional density, and frequency and size responses of subsequent price change in the PPI dataset. The figures show that the density of the competitor-price gaps has fat tails (first panel); the frequency of subsequent price adjustment increases with the absolute size of the gap, after the gap is large enough; and the size of average subsequent adjustments are negatively related with the gap (third panel).

aggregate shock, the new adjusters do not come from those prices that are further from their optimal levels. None of the interaction terms are statistically significant.

The impact of the price gaps and the aggregate credit shocks also remain economically significant. The probability of price increases between a product with a competitor price gap at the third quartile and at the first quartile gets smaller by 23 percentage points, and the probability of price increases gets larger by 22 percentage points. Finally, a one standard deviation credit tightening reduces the price increase probability by around 0.7 percentage points and increases the price decrease probability by a similar amount.

Overall, these results indicate the absence of a selection effect generalises from the supermarket sector to the wider US economy.

7.5 Response to monetary policy shocks

This section describes the properties of our second identified shock, the monetary policy shock, and price adjustment in response. Again, we find that prices adjust predominantly through the gross extensive margin. Additionally, we also find a response of the size and frequency of price changes that is inconsistent with time dependent models.

As we have done with the credit shock, we start by characterizing the dynamic impact of a monetary policy shock on inflation and its components using the local projection method (Jordà, 2005). The local projection framework puts minimal structure on the data generating process. As instruments for monetary policy shocks, we use changes in the 3-month-ahead federal funds futures in a 30-minute window around FOMC press statements like Gertler and Karadi (2015). The identification assumption is that because financial markets incorporate all available information into futures prices before the announcement, the change in the futures price indicates the size of the policy surprise. Furthermore, the narrow window guarantees that no other economic shock systematically contaminates the measure. We restrict our interest to announcements where the interest rate surprise and the S&P blue-chip stock price index moved in the opposite direction over the same time frame. As argued by Jarociński and Karadi (2020), such co-movement is indicative of a dominant monetary policy shock, when the impact of the central bank's contemporaneous announcements

Table 7: Estimates, PPI data, competitors' price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $(I_{pst,t+24}^+)$		Decreases $(I_{pst,t+24}^-)$	
Gap (x_{pst-1})	-0.23*** (0.02)	-0.23*** (0.02)	0.22*** (0.02)	0.22*** (0.02)
Shock (ebp_t)	-0.023*** (0.01)		0.021*** (0.01)	
Selection $(x_{pst-1}ebp_t)$	0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
Age (T_{pst-1})	0.035*** (0.00)	0.035*** (0.00)	0.01*** (0.00)	0.01*** (0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✗	✓	✗
Time FE	✗	✓	✗	✓
N	9.7M	9.7M	9.7M	9.7M
Within R^2	4.4%	3.5%	4.3%	3.7%

*Note: The table shows estimation results from the linear-probability panel model using PPI microdata. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions include product-store fixed effects, control for the age (time spent since last change) of the price, and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with time-fixed effects (columns 2 and 4). Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

about the economic outlook played a minor role. We transform these surprises to monthly variables by summing up monetary policy surprises within each calendar month. We run a series ($h = 0, \dots, 24$ months) of regressions of the following form:

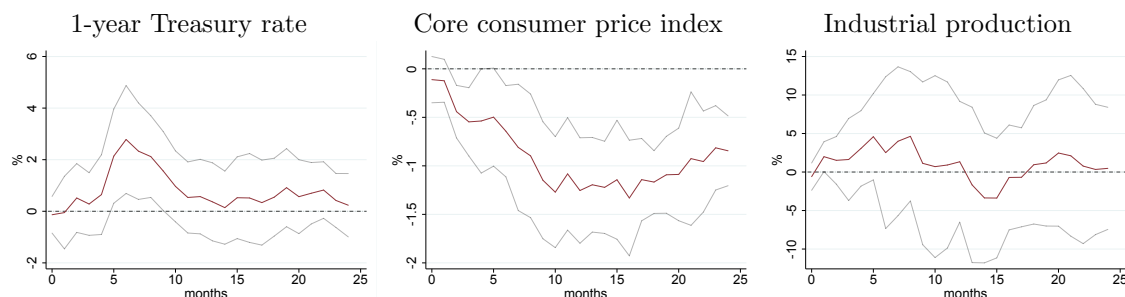
$$x_{t+h} - x_t = \alpha_h + \beta_h \Delta i_t + \Gamma_h \Phi(L) X_t + u_{t,h}, \quad (14)$$

where x_t is the variable of interest (for example the log price level), and Δi_t is our proxy for a monetary policy shock. The local projections also include a set of controls $\Gamma_h \Phi(L) X_t$, where Γ_h is a vector of parameters for each h , X_t is a vector of control variables and $\Phi(L)$ is a lag-polynomial. Unless stated otherwise, the controls we use are the 1 to 6 months lags of the 1-year Treasury rate, the consumer price index, industrial production and the excess bond premium (Gilchrist and Zakrajšek, 2012).

Our key object of interest is the coefficient β_h . In the figures below, we plot β_h , $h = 0, 1, \dots, 24$ along with 95% confidence bands.

First, consumer prices responds as expected to the monetary shock. Figure 15 plots impulse responses to some key macroeconomic variables. In particular, we plot the response to the 1-year constant maturity

Figure 15: Impulse responses of key macroeconomic variables to a monetary policy tightening



Note: The panels show impulse responses to an identified monetary policy shock in the scanner-data sample between 2001 and 2012 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The panels show that a monetary policy causes a sizable drop in the price index, but no noticeable drop in activity.

treasury rate, the response to the logarithm of the consumer price index excluding food and energy, and the response of the logarithm of the industrial production. The response of consumer prices are consistent with standard results (Gertler and Karadi, 2015): the interest-rate increase generates a delayed and hump-shaped decline in the core consumer price index. Within the short sample, we observe no noticeable drop in industrial production.

Figure 16: Impulse responses of the supermarket-price indices to a monetary policy tightening



Note: The panels show impulse responses to an identified monetary policy shock in the scanner-data sample between 2001 and 2012 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The figures show that the reference price index declines significantly at around a 6-month horizon as a response to the shock. The posted price index increases significantly at a two-year horizon contrary to standard theory, but the increase is mostly driven by an increase in the filtered-out sales-price index.

Second, when we consider the response of our supermarket prices, supermarket reference prices exhibit the expected response to a monetary tightening. Figure 16 plots the impulse responses of supermarket prices. Similarly to the aggregate price index, supermarket reference prices display a hump-shaped decline with wide confidence bands after a monetary policy tightening (see middle panel). The decomposition of posted prices to reference- and sales-price indices also reveals that the supermarkets respond to the shock by actively adjusting their reference-prices, and not by modifying their strategy on temporary sales. This finding is consistent with views that argue that temporary-sales strategies are predetermined and not an active adjustment margin at business cycle frequencies (Anderson, Malin, Nakamura, Simester and Steinsson, 2017). Consequently, we concentrate on reference-prices in our subsequent analysis.

As before, we find that adjustment of the reference-price level happens through the extensive margin (by modifying the number of price changes) rather than through the intensive margin (by changing the average size of price changes). To show this result, we decompose the cumulative reference-price inflation into the frequency ($\xi_{t,t+h}$) and the size ($\psi_{t,t+h}$) of price increases and price decreases as follows:

$$p_{t+h} - p_{t-1} = \pi_{t,t+h} = \xi_{t,t+h}^+ \psi_{t,t+h}^+ + \xi_{t,t+h}^- \psi_{t,t+h}^-. \quad (15)$$

The frequency of reference-price²¹ increases and decreases are defined as

$$\xi_{t,t+h}^\pm = \sum_i \bar{\omega}_{it,t+h} I_{it,t+h}^\pm, \quad (16)$$

where $I_{it,t+h}^+$ and $I_{it,t+h}^-$ are indicators that take the value 1 if the reference price of item i (a product in a particular store) increased or decreased between months $t-1$ and $t+h$, respectively, and 0 otherwise. The weight $\bar{\omega}_{it,t+h}$ is measured as the average weight of the product between t and $t+h$. The average size of price increases and decreases are defined as

$$\psi_{t,t+h}^\pm = \frac{\sum_i \bar{\omega}_{it,t+h} I_{it,t+h}^\pm (p_{it+h} - p_{it-1})}{\xi_{t,t+h}^\pm}. \quad (17)$$

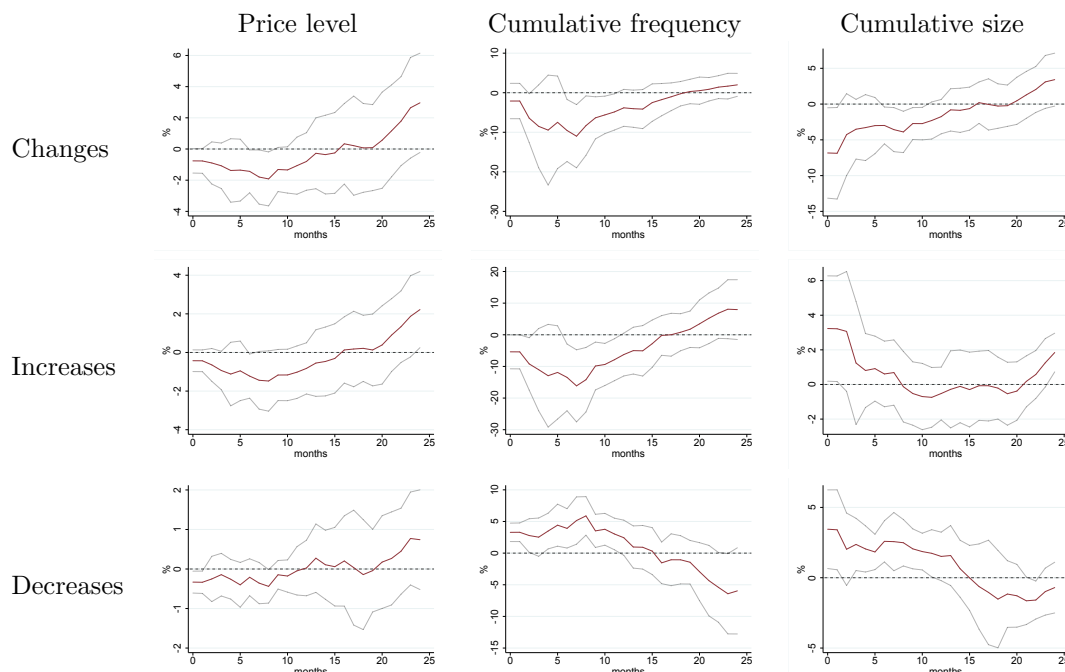
Following a monetary shock, we find that there is a strong adjustment on the extensive margin within a year of the policy shock: the cumulative frequency of reference-price increases declines, and the cumulative frequency of price decreases rises. The decline in the cumulative price increases is larger (around 15 percent at its peak) than the increase in cumulative price decreases (around 5 percent at its peak), so the aggregate frequency declines. Both of these changes contribute to the decline in the price level, and to the reduction in the average size of price changes. Figure 17 illustrates the decomposition of the response to a monetary policy shock into these adjustment margins.

By contrast, the average size of price increases rises and the absolute size of the price decreases declines. Both of them mitigate the impact of the shock on the price level rate. Such evidence is inconsistent with the underlying assumptions in time-dependent models (Calvo, 1983), which assume a constant frequency of price changes and attribute the adjustment after a monetary policy shocks to the intensive margin. Our evidence instead points to the importance of the extensive margin, as the frequency of price increases and decreases adjust significantly. It challenges the predominance of the intensive margin adjustment, which would predict *smaller* price increases and *larger* price decreases. The increase in the frequency and the decline in the average absolute size of price adjustment has been documented after large (e.g. value-added tax, exchange rate) shocks by Karadi and Reiff (2019); Auer et al. (2018), who also showed that menu cost pricing models with leptokurtic idiosyncratic productivity shock (Midrigan, 2011) are consistent with this pattern. To our knowledge, we are the first to document the same pattern after regular monetary policy shocks in U.S. data. In contemporaneous work, a similar pattern has been documented using German PPI data by Balleer and Zorn (2019).

Table 8 shows that the estimated coefficients are similar to those from our specification using the credit shock and tell the same story. There is evidence for the state-dependence of price changes as the price gaps

²¹We suppress the superscript f for notational convenience.

Figure 17: Adjustment on the extensive and intensive margins to a monetary policy tightening



Note: The figures show the impulse responses of supermarket reference-price levels and their components to an identified monetary policy shock, and 95% confidence bands using Newey-West standard errors. The figures show that the negative price-level response (first figure in the first row) is predominantly driven by the decline in the price-increase frequency (second panel in the second row) and the increase in the price-decrease frequency (second panel in the third row). This adjustment in the gross extensive margin explains the decline in the cumulative size of the price changes (third panel in the first row). The cumulative frequency (net extensive margin) also declines (second panel in the first row).

and the monetary policy shock significantly impact the probability of price changes, but there is no evidence for selection as they do not interact.

7.6 Heterogeneity across product categories

We find that our results are robust to heterogeneity across product categories when we run our baseline regression separately for the 31 different product categories. This is reassuring because our baseline regression does not differentiate between potentially heterogeneous responses to idiosyncratic and aggregate volatility across product categories. The heterogeneity in demand elasticities due to different product characteristics and market structure (e.g. alcohol vs. milk) might potentially bias our estimates.

We show in Figures 18 and 19 that indeed the results are robust across product categories. Figure 18 shows the estimated coefficients and the uncertainty surrounding them for price increases, Figure 19 for price decreases. We find that in the majority of categories a higher gap (price too high) and a credit tightening both significantly decrease the probability of a price increase and increase the probability of a price decrease. The interaction of gap and shock is never significant again indicating that selection is absent.

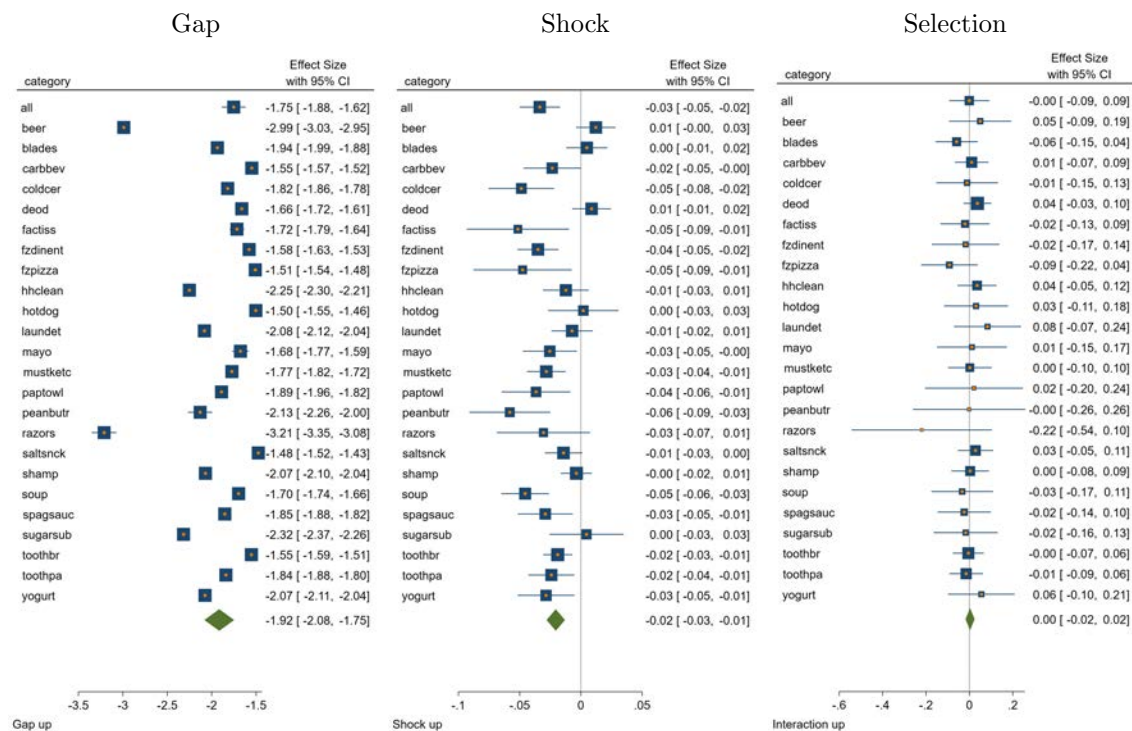
Table 8: Estimates, scanner data, competitors' price gap, monetary policy shock

	(1)	(2)	(3)	(4)	(5)	(6)
	Price increases ($I_{pst,t+12}^+$)			Price decreases ($I_{pst,t+12}^-$)		
Gap (x_{pst-1})	-1.71*** (0.06)	-1.71*** (0.06)		1.36*** (0.05)	1.36*** (0.05)	
Shock (Δi_t)	-0.03* (0.01)		-0.03 (0.02)	0.01* (0.01)		0.01* (0.01)
Selection ($x_{pst-1}\Delta i_t$)	-0.07 (0.06)	-0.07 (0.05)		0.07 (0.06)	0.07 (0.05)	
Age (T_{pst-1})	0.03*** (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.01*** (0.00)	0.01*** (0.00)	0.01*** (0.00)
Positive gap (x_{pst-1}^+)			-1.92*** (0.10)			1.93*** (0.09)
Negative gap (x_{pst-1}^-)			-1.58*** (0.06)			1.01*** (0.05)
Pos. selection ($x_{pst-1}^+\Delta i_t$)			-0.05 (0.09)			0.05 (0.05)
Neg. selection ($x_{pst-1}^-\Delta i_t$)			-0.08 (0.12)			0.08 (0.08)
Product x store FE	✓	✓	✓	✓	✓	✓
Calendar-month FE	✓	✗	✓	✓	✗	✓
Time FE	✗	✓	✗	✗	✓	✗
N	23.7M	23.7M	23.7M	23.7M	23.7M	23.7M
Within R^2	16.4%	14.7%	16.5%	13.3%	12.7%	13.8%

Note: The table shows estimation results from the linear-probability panel model using scanner data and a monetary policy shock. The regressions are run separately on an indicator with value 1 for reference-price increases (columns 1-3) and an indicator with value 1 for reference-price decreases (columns 4-6). The regressions include product-store and calendar-month fixed effects, control for the age of the price (time spent since last change), and use standard errors with two-way clustering. The baseline regressions (columns 1 and 4) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities, which is an evidence for state-dependence, their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with time-fixed effects (columns 2 and 5) and a specification with separate coefficients for positive and negative gaps (columns 3 and 6).

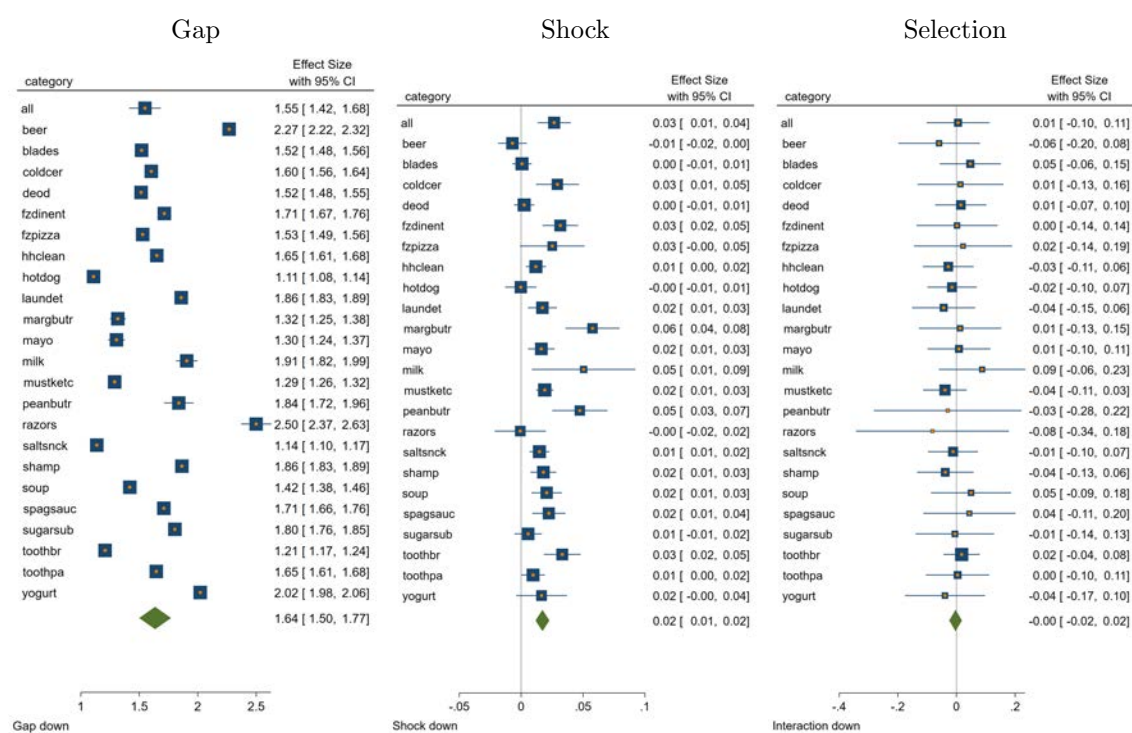
Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.

Figure 18: Estimated coefficients across product categories, price increases



Note: The figure shows estimates across product categories for price increases with 95 percent confidence bands for our baseline specification that uses the scanner data, the competitor-price gap measure, and credit shocks. The panels show that the results are robust across product categories: the higher gap significantly decreases the price-increase probability; credit tightening decreases the price-increase probability in the majority of categories, and their interaction is never significant, indicating selection is absent.

Figure 19: Estimated coefficients across product categories, price decreases



8 Discussion

What does our evidence imply about the extent of monetary non-neutrality? Which theoretical models are consistent with our evidence? While our facts may guide future modeling work, we briefly point out in this section which models are consistent with our findings, and what they imply for monetary non-neutrality.

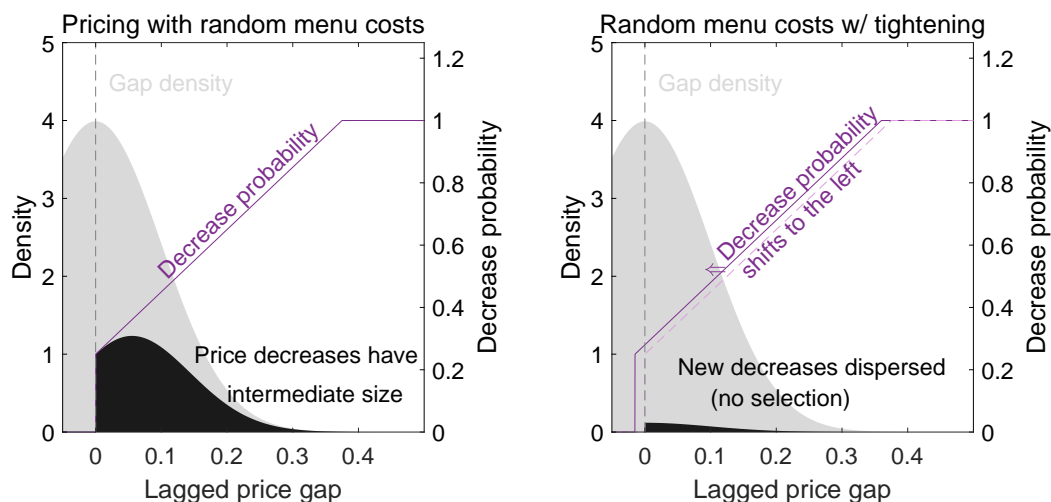
Recall that we have documented three sets of empirical facts. First, we have found strong evidence for state-dependence, as, *unconditionally*, the probability of price adjustment increases with the larger (absolute) value of the price gap. Second, we have documented the lack of evidence for selection, as, *conditional* on identified credit and monetary-policy shocks, there is no evidence that prices with larger (absolute) price gaps are adjusting with higher probability. Third, we have found evidence for the gross extensive margin: the aggregate shock shifts the relative share of price increases and price decreases conditional on the pre-shock gaps.

Which theoretical models can justify these empirical results, and what do they imply for the extent of monetary non-neutrality? We argue that a class of state-dependent models with random menu costs can be in line with our evidence (Dotsey et al., 1999; Costain and Nakov, 2011; Luo and Villar, 2017; Alvarez et al., 2020). In particular, those with a (piecewise) linear and flat generalized hazard function of the form

$$\Lambda(x) = \begin{cases} a + bx & \text{if } x \geq 0 \text{ and } a + bx \leq 1 \\ a - cx & \text{if } x < 0 \text{ and } a - cx \leq 1 \\ 1 & \text{otherwise} \end{cases} \quad (18)$$

where a, b, c are parameters. We consider the hazard function flat when it takes a value strictly below 1 in the relevant range of the gap distribution.

Figure 20: State-dependence and selection with random menu cost



Note: The figure shows the density of the before-shock price-gaps, the price-decrease hazard, and the price-decrease density as a function of gaps in a linear-hazard random-menu-cost model. The left panel illustrates state dependence, showing that the probability of price decrease is increasing with the gap size. The right panel illustrates the lack of selection, as after an aggregate tightening, which shifts the price-decrease hazard to the left, the new price decreases are dispersed along the price-gap distribution.

Table 9: Overview of active adjustment channels

	Data	Time- dependent	(S,s) & Convex hazard	Linear hazard
Intensive margin	✓	✓	✓	✓
Gross extensive margin	✓	✗	✓	✓
Selection	✗	✗	✓	✗

Note: The table overviews the presence of various adjustment channels emerged from the data through our analysis and in popular price setting models. The data is consistent with random menu cost and rational inattention models with linear hazard, but they are inconsistent with time-dependent model (Calvo, 1983) and standard fixed-menu-cost (S,s) models (Goloso and Lucas, 2007) and other models generating (strongly) convex hazard functions.

Figure 20 illustrates such a model. It is consistent with all our evidence. First, the embodied underlying model implies state dependence: the probability of price adjustment increases with the price gap ($b, c > 0$), as confirmed by the direct impact of the gap on price-change probability in our baseline regression. Second, the *linear* hazard function is consistent with the close to linear relationship found between price gaps and the probability of price adjustment imposed in our baseline regressions, and confirmed by the non-linear specification reported in Section 7.1 and Figure 9. Third, the gross extensive margin effect is active: the tightening generates a rise in the price-decrease probability, which shows up as the direct impact of the aggregate shock in our baseline regression. The same shock causes an analogous decrease in the price-increase probability: the shift in the relative share of increases versus decreases generate adjustment at the gross extensive margin. Formally, adjustment at the gross extensive margin equals $F(0)\bar{x}^+c + (1-F(0))(-\bar{x}^-)b > 0$. And fourth, the linear hazard function generates no selection. The reason is that $\Lambda'(x)$ is constant in the relevant range, and equals to b when $x \geq 0$ and to $-c$, and therefore $\Lambda'(x) - \int_{x \geq 0} \Lambda'(x)h(x)dx = 0$ when $x \geq 0$ and $\Lambda'(x) - \int_{x < 0} \Lambda'(x)g(x)dx = 0$ when $x < 0$ in the relevant range. Intuitively, the new adjusters are dispersed over the price-gap space: The increase in the probability of adjustment (distance between the hazard functions) is the same over most of the price gap space – except close to zero (which are ignored by definition) and at very high values (with very low mass). Therefore, our baseline regression finds no significant impact on the interaction of the aggregate shock and the price gap on the price-adjustment probabilities. This also means that the correlation between new adjusters and the gap (and therefore the size of price adjustment) is zero. As selection is absent, monetary non-neutrality is high, not unlike in time-dependent models.

9 Conclusion

This paper measures the price selection in supermarket-scanner and producer-price microdata. There is no evidence for selection as conditional on aggregate credit or monetary policy shocks, firms do not choose to adjust prices that are further from their optimal levels. This occurs despite the fact that unconditionally they do adjust prices with large price gaps with higher probability and they also adjust the probability of price adjustments conditional on a shock. The evidence challenges both standard time-dependent (Calvo, 1983) and state-dependent models with high selection (Goloso and Lucas, 2007; Karadi and Reiff, 2019),

but is consistent with state-dependent models with random menu costs (Dotsey et al., 1999; Luo and Villar, 2017; Alvarez et al., 2020) or rational inattention (Woodford, 2009) with linear and flat adjustment hazard, which predicts high monetary non-neutrality similar to time-dependent models.

References

- Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi (2014) “Small and Large Price Changes and the Propagation of Monetary Shocks,” Working Paper 20155, National Bureau of Economic Research.
- Alvarez, Fernando E., Francesco Lippi, and Aleksei Oskolkov (2020) “The Macroeconomics of Sticky Prices with Generalized Hazard Functions,” NBER Working Papers 27434, National Bureau of Economic Research, Inc.
- Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jon Steinsson (2017) “Informational Rigidities and the Stickiness of Temporary Sales,” *Journal of Monetary Economics*, Vol. 90, pp. 64–83.
- Argente, David and Chen Yeh (2020) “Product Life Cycle, Learning, and Nominal Shocks,” *unpublished manuscript*.
- Auer, Raphael, Ariel Burstein, and Sarah Lein (2018) “Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation,” BIS Working Papers 751, Bank for International Settlements.
- Balleer, Almut and Peter Zorn (2019) “Monetary Policy, Price Setting, and Credit Constraints,” CEPR Discussion Papers 14163, C.E.P.R. Discussion Papers.
- Berger, David and Joseph Vavra (2018) “Dynamics of the US Price Distribution,” *European Economic Review*, Vol. 103, pp. 60–82.
- Bhattarai, Saroj and Raphael Schoenle (2014) “Multiproduct Firms and Price-Setting: Theory and Evidence from US Producer Prices,” *Journal of Monetary Economics*, Vol. 66, pp. 178–192.
- Bils, Mark, Peter J Klenow, and Benjamin A Malin (2012) “Reset Price Inflation and the Impact of Monetary Policy Shocks,” *The American Economic Review*, Vol. 102, pp. 2798–2825.
- BLS (2020) “Producer Price Index Frequently Asked Questions,” Technical report, <https://www.bls.gov/ppi/ppifaq.htm>.
- Bonomo, Marco, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, and Rodolfo Rigato (2019) “Multi-Product Pricing: Theory and Evidence for Large Retailers in Israel,” unpublished manuscript, Centre for Economic Policy Research.
- Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela (2008) “Database Paper – The IRI Marketing Data Set,” *Marketing Science*, Vol. 27, pp. 745–748.
- Caballero, Ricardo J and Eduardo MRA Engel (2007) “Price Stickiness in Ss models: New Interpretations of Old Results,” *Journal of Monetary Economics*, Vol. 54, pp. 100–121.

- Calvo, Guillermo A. (1983) “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, Vol. 12, pp. 383 – 398.
- Caplin, Andrew S. and Daniel F. Spulber (1987) “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, Vol. 102, pp. 703–726.
- Carvalho, Carlos and Oleksiy Kryvtsov (2018) “Price Selection,” staff working papers, Bank of Canada.
- Chamberlain, Gary (1980) “Analysis of Covariance with Qualitative Data,” *The Review of Economic Studies*, Vol. 47, pp. 225–238.
- Costain, James and Anton Nakov (2011) “Price Adjustments in a General Model of State-Dependent Pricing,” *Journal of Money, Credit and Banking*, Vol. 43, pp. 385–406.
- Dedola, L, M Strom Krisoffersen, and G Zullig (2019) “Price Synchronization and Cost Passthrough in Multiproduct Firms: Evidence from Danish Producer Prices,” Technical report, Mimeo.
- Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999) “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *The Quarterly Journal of Economics*, Vol. 114, pp. 655–690.
- Gagnon, Etienne, David López-Salido, and Nicolas Vincent (2012) “Individual Price Adjustment along the Extensive Margin,” *NBER Macroeconomics Annual 2012, Volume 27*, pp. 235–281.
- Gertler, Mark and Peter Karadi (2015) “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, Vol. 7, pp. 44–76.
- Gertler, Mark and John Leahy (2008) “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, Vol. 116, pp. 533–572.
- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajek (2017) “Inflation Dynamics during the Financial Crisis,” *American Economic Review*, Vol. 107, pp. 785–823.
- Gilchrist, Simon and Egon Zakrajšek (2012) “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, Vol. 102, pp. 1692–1720.
- Goldberg, Pinelopi K. and Rebecca Hellerstein (2009) “How rigid are producer prices?,” Staff Reports 407, Federal Reserve Bank of New York.
- Golosov, Mikhail and Robert E. Lucas (2007) “Menu Costs and Phillips Curves,” *Journal of Political Economy*, Vol. 115, pp. 171–199.
- Gorodnichenko, Yuriy and Michael Weber (2016) “Are Sticky Prices Costly? Evidence from the Stock Market,” *American Economic Review*, Vol. 106, pp. 165–199.
- Hahn, Jinyong and Whitney Newey (2004) “Jackknife and Analytical Bias Reduction for Nonlinear Panel Models,” *Econometrica*, Vol. 72, pp. 1295–1319.

- Hong, Gee Hee, M Klepaczy, Ernesto Pasten, and Raphael Schoenle (2019) “The Real Effects of Monetary Shocks: Evidence from Micro Pricing Moments,” Working Papers Central Bank of Chile 875, Central Bank of Chile.
- Jarociński, Marek and Peter Karadi (2020) “Deconstructing Monetary Policy Surprises: the Role of Information Shocks,” *American Economic Review: Macroeconomics*, Vol. 12, pp. 1–43.
- Jordà, Òscar (2005) “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, Vol. 95, pp. 161–182.
- Karadi, Peter and Adam Reiff (2019) “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*, Vol. 11, pp. 111–46.
- Kehoe, Patrick and Virgiliu Midrigan (2014) “Prices are Sticky After All,” *Journal of Monetary Economics*.
- Luo, Shaowen and Daniel Vallas Villar (2017) “The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models,” Finance and Economics Discussion Series 2017-028, Board of Governors of the Federal Reserve System (U.S.).
- (2019) “The Price Adjustment Hazard Function: Evidence from High Inflation Periods,” Technical report, manuscript.
- Midrigan, Virgiliu (2011) “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, Vol. 79, pp. 1139–1180.
- Nakamura, Emi and Jón Steinsson (2008) “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, Vol. 123, pp. 1415–1464.
- Newey, Whitney and Kenneth West (1987) “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, Vol. 55, pp. 703–08.
- Petrella, Ivan, Emiliano Santoro, and Lasse P. Simonsen (2019) “Time-varying Price Flexibility and Inflation Dynamics,” EMF Research Papers 28, Economic Modelling and Forecasting Group.
- Pförr, Klaus (2014) “femlogit-Implementation of the multinomial logit model with fixed effects,” *Stata Journal*, Vol. 14, pp. 847–862.
- Plagborg-Møller, Mikkel and Christian K Wolf (2019) “Local Projections and VARs Estimate the Same Impulse Responses,” *Unpublished paper: Department of Economics, Princeton University*.
- Stock, James H and Mark W Watson (2018) “Identification and Estimation of Dynamic Causal Effects in Macroeconomics using External Instruments,” *The Economic Journal*, Vol. 128, pp. 917–948.
- Woodford, Michael (2009) “Information-Constrained State-Dependent Pricing,” *Journal of Monetary Economics*, Vol. 56, pp. S100–S124.
- Wooldridge, Jeffrey M (2010) *Econometric Analysis of Cross Section and Panel Data*: MIT press.

A Appendix

The appendix shows that our baseline results are robust to excluding product-store fixed effects (Table 10), to using posted prices instead of reference prices (Table 11), using only product-month combinations with at least 50 competitors (Table 12, and to ending the sample in 2007 just before the Great Financial Crisis (Table 13).

Table 10: Robustness to dropping item fixed effects, scanner data, competitors' price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $\left(I_{pst,t+24}^+\right)$		Decreases $\left(I_{pst,t+24}^-\right)$	
Gap (x_{pst-1})	-1.75*** (0.06)	-0.99*** (0.10)	1.55*** (0.06)	0.90*** (0.10)
Shock (ebp_t)	-0.03*** (0.01)	-0.04*** (0.01)	0.03*** (0.01)	0.03** (0.01)
Selection $(x_{pst-1}\hat{ebp}_t)$	-0.00 (0.04)	-0.01 (0.02)	0.01 (0.05)	0.02 (0.03)
Age (T_{pst-1})	0.02*** (0.00)	-0.01** (0.01)	0.00** (0.00)	-0.03*** (0.00)
Product x store FE	✓	✗	✓	✗
Calendar-month FE	✓	✓	✓	✓
Time FE	✗	✗	✗	✗
N	16.1M	16.1M	16.1M	16.1M
Within R^2	18.5%	8.9%	17.3%	9.3%

*Note: The table shows estimation results from the linear-probability panel model using scanner data with and without product-store fixed effects. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions control for the age (time spent since last change) of the price and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification without product-store fixed effects (columns 2 and 4). Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

Table 11: Robustness using posted prices, scanner data, competitors' price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $\left(I_{pst,t+24}^+\right)$		Decreases $\left(I_{pst,t+24}^-\right)$	
	Reference	Posted	Reference	Posted
Gap (x_{pst-1})	-1.75*** (0.06)	-1.46*** (0.05)	1.55*** (0.06)	1.25*** (0.05)
Shock (ebp_t)	-0.03*** (0.01)	-0.04*** (0.01)	0.03*** (0.01)	0.03*** (0.01)
Selection $(x_{pst-1}ebp_t)$	-0.00 (0.04)	-0.01 (0.03)	0.01 (0.05)	0.02 (0.04)
Age (T_{pst-1})	0.02*** (0.00)	0.01*** (0.00)	0.00** (0.00)	-0.01*** (0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✓	✓	✓
Time FE	✗	✗	✗	✗
N	16.1M	18.6M	16.1M	18.6M
Within R^2	18.5%	17.6%	17.3%	14.8%

Note: The table shows estimation results from the linear-probability panel model using scanner data with reference and posted prices. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions include product-store and calendar-month fixed effects control for the age (time spent since last change) of the price and aggregate variables, and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust using posted prices.

*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

Table 12: Robustness using product-months with at least 50 competitors, scanner data, competitors' price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $(I_{pst,t+24}^+)$		Decreases $(I_{pst,t+24}^-)$	
	Baseline	50+ competitors	Baseline	50+ competitors
Gap (x_{pst-1})	-1.75*** (0.06)	-1.76*** (0.06)	1.55*** (0.06)	1.56*** (0.06)
Shock (ebp_t)	-0.03*** (0.01)	-0.03*** (0.01)	0.03*** (0.01)	0.03*** (0.01)
Selection $(x_{pst-1}\hat{ebp}_t)$	-0.00 (0.04)	-0.00 (0.05)	0.01 (0.05)	0.01 (0.05)
Age (T_{pst-1})	0.02*** (0.00)	0.02*** (0.00)	0.00** (0.00)	0.00** (0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✓	✓	✓
Time FE	✗	✗	✗	✗
N	16.1M	15.3M	16.1M	15.3M
Within R^2	18.5%	18.9%	17.3%	17.7%

Note: The table shows estimation results from the linear-probability panel model using scanner data with product-month with at least 1 (baseline) and at least 50 competitors. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions include product-store and calendar-month fixed effects control for the age (time spent since last change) of the price and aggregate variables, and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust if we restrict to product-store combinations with at least 50 competitors. Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.

Table 13: Robustness to excluding the Great Recession, scanner data, competitors' price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $\left(I_{pst,t+24}^+\right)$		Decreases $\left(I_{pst,t+24}^-\right)$	
	2001-2012	2001-2007	2001-2012	2001-2007
Gap (x_{pst-1})	-1.75*** (0.06)	-1.74*** (0.07)	1.55*** (0.06)	1.50*** (0.06)
Shock (ebp_t)	-0.03*** (0.01)	-0.03*** (0.01)	0.03*** (0.01)	0.02*** (0.01)
Selection $(x_{pst-1}ebp_t)$	-0.00 (0.04)	0.06 (0.07)	0.01 (0.05)	-0.06 (0.07)
Age (T_{pst-1})	0.02*** (0.00)	0.02*** (0.00)	0.00** (0.00)	0.01*** (0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✓	✓	✓
Time FE	✗	✗	✗	✗
N	16.1M	9.9M	16.1M	9.9M
Within R^2	18.5%	17.7%	17.3%	16.5%

Note: The table shows estimation results from the linear-probability panel model using scanner data with reference and posted prices. The regressions are run separately on an indicator with value 1 for price increases (columns 1-2) and an indicator with value 1 for price decreases (columns 3-4). The regressions include product-store and calendar-month fixed effects control for the age (time spent since last change) of the price and aggregate variables, and use standard errors with two-way clustering. The baseline regressions (columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust using posted prices.

*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

Acknowledgements

All opinions expressed are personal and do not necessarily represent the view of the European Central Bank or the Eurosystem, the Federal Reserve Bank of Cleveland or the Federal Reserve Board System. Part of this research was conducted with restricted access to the Bureau of Labor Statistics (BLS) data. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS. We thank Fernando Alvarez, Florin Bilbiie, Luca Dedola, Yuriy Gorodnichenko, Cosmin Ilut, Anton Nakov, Luminita Stevens, Ivan Werning and participants at the Ifo Conference 2020 and seminars at the Central Bank of Netherlands (DNB), Duke University, European Central Bank, and University of Maryland for their comments and Lydia Cox for excellent research assistance. Jesse Wursten gratefully acknowledges financial support of the Research Foundation Flanders (FWO) under grant number 1125818N.

Peter Karadi

European Central Bank, Frankfurt am Main, Germany; CEPR, London, United Kingdom; email: peter.karadi@ecb.europa.eu

Raphael Schoenle

Federal Reserve Bank of Cleveland, Cleveland, Ohio, United States; Brandeis University, Waltham, Massachusetts, United States; CEPR, London, United Kingdom; email: raphael.schoenle@clev.frb.org

Jesse Wursten

KU Leuven, Leuven, Belgium; IRLE, UC Berkeley, Berkeley, California, United States; email: jesse.wursten@kuleuven.be

© European Central Bank, 2021

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-4752-7

ISSN 1725-2806

doi:10.2866/510642

QB-AR-21-057-EN-N