

NBER WORKING PAPER SERIES

DYNAMICS OF THE U.S. PRICE DISTRIBUTION

David Berger
Joseph Vavra

Working Paper 21732
<http://www.nber.org/papers/w21732>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2015

We are grateful to Nick Bloom, Eduardo Engel, Nir Jaimovich, Giuseppe Moscarini and Emi Nakamura for helpful comments and Rozi Ulics and Randal Verbrugge for support at the BLS. Finally, we thank Chiara Maggi and Yuta Takahashi for being wonderful RAs. All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by David Berger and Joseph Vavra. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamics of the U.S. Price Distribution
David Berger and Joseph Vavra
NBER Working Paper No. 21732
November 2015
JEL No. E3,E31,E32,E5,E52,L16

ABSTRACT

We use microdata underlying U.S. consumer, producer and import price indices to document how the distribution of price changes evolves over time. Two striking features characterize pricing at each stage of production: 1) Frequency is countercyclical. 2) Frequency is correlated with variance. Conversely, other statistics which have received recent attention, like kurtosis, do not exhibit uniform patterns across our datasets. What implications do our empirical results have for monetary policy? Using a flexible accounting framework which collapses the high-dimensional distribution of price changes into a single measure of aggregate price flexibility, we show that flexibility is highly variable and countercyclical.

David Berger
Department of Economics
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
david.berger@northwestern.edu

Joseph Vavra
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
joseph.vavra@chicagobooth.edu

1 Intro

A growing literature argues that the *microeconomic* distribution of price changes matters for *macroeconomic* price flexibility and thus monetary policy. In this paper, we extend the existing empirical literature by systematically documenting the time-series evolution of the entire distribution of U.S. price changes at various stages of production. Using the Bureau of Labor Statistics (BLS) microdata that underlies the Consumer, Producer and Import Price Indices we show that there are important common patterns in the distribution of price changes over time. We then explore the implications of this variation for aggregate price flexibility. Using a simple, flexible accounting framework we argue that price flexibility rises in recessions.

While there has been widespread attention to "first moments"² of the price change distribution, there has been much less empirical study of higher moments of the distribution and their relationship to the broader business cycle.³ Furthermore, existing studies have focused on particular moments and data sets in isolation, which makes it more challenging to identify robust features of pricing behavior.⁴ In this paper, we show that there are striking common patterns in the distribution of price changes collected at different stages of production, but there are also certain features which are unique to particular data sets.

We systematically report time-series statistics for numerous moments and percentiles that go well beyond the existing literature, and several empirical regularities emerge from this analysis: 1) There are large movements across time in all percentiles of the distribution of price changes. 2) The frequency of adjustment is positively correlated with the variance of price changes. 3) The frequency of adjustment and variance of price changes are strongly countercyclical. We show that these basic facts hold across each of our data sets and regardless of how price series are filtered.⁵ Conversely, some patterns related to higher moments

²Studies typically focus on e.g. the frequency and size of price changes and their relationship to inflation.

³Klenow and Malin (2010) and Vavra (2014) are exceptions

⁴For example, using CPI data, Vavra (2014) and Alvarez and Lippi (2014) explore the implications of the variance of price changes for monetary policy while Midrigan (2011) and Alvarez et al. (2014) focus on the implications of kurtosis. Berger and Vavra (2015) focus on the implications of variance in IPP import price data.

⁵It is important to note that many, but not all of these empirical facts are new. In particular, all of the empirical facts relating to "centered moments" of the CPI from an earlier draft of this paper were subsumed in Vavra (2014). In particular, Table 1 in Vavra (2014) documents that the frequency is countercyclical as well as the business cycle co-movement of the variance, skewness and kurtosis of the distribution of price changes. Berger and Vavra (2015) document the positive correlation between the frequency of adjustment and the standard deviation of price changes in IPP data. All remaining statistics are to the best of our knowledge new to this paper.

of the distribution of price changes differ across CPI, PPI and IPP data, or are sensitive to measurement issues. In particular: 4) Various measures of price change kurtosis are strongly procyclical and are negatively correlated with frequency in the CPI, but not in PPI or IPP data. 5) Statistics related to skewness are highly sensitive to the particular measure used and also vary substantially across data sets.

Why is it important to study the distribution of price changes and what should we take from the array of statistics computed in the first half of the paper? Microeconomic price-setting behavior influences the degree of aggregate price flexibility, which will in turn have strong implications for the real response of the economy to nominal shocks. In the second half of the paper, we introduce an accounting framework which allows us to collapse the complicated high dimensional distribution of price changes at each point in time into a single, easily interpretable measure of aggregate price flexibility.

This step necessarily requires introducing additional structure, but we try to do so in a highly flexible way. For example, in a Calvo model, firms are selected to adjust prices at random so aggregate price flexibility is completely determined by the average frequency of adjustment. At the opposite extreme, in the Caplin and Spulber (1987) model, adjusting firms change prices by such large amounts that the aggregate price level is fully flexible regardless of the underlying frequency of adjustment. Rather than taking a strong stand on a particular price-setting environment, we use a version of the generalized Ss model of Caballero and Engel (2007), which nests many of these extremes. Furthermore, we estimate this model using a highly flexible functional form which imposes minimal restrictions on the distribution of desired price changes at a point in time and no restrictions on the evolution across time.

The flexible modeling framework of Caballero and Engel (2007) is useful for summarizing our somewhat complicated pricing facts and their implications for how price flexibility varies over time. We show that greater frequency, greater variance and smaller kurtosis are all associated with greater price flexibility in this model. In contrast, the skewness of firms' desired price changes has little relationship with aggregate price flexibility. Thus, movements across time in the frequency of adjustment, variance or kurtosis of price changes should be associated with movements in aggregate price flexibility. When viewed through the lens of our model, we find that most of the time-series correlations we document in the BLS data imply time-varying flexibility which rises during recessions. That is, aggregate price flexibility is highly variable and strongly countercyclical. Furthermore, we find that a large fraction of

this time-variation in price flexibility arises from changes in the distribution of price changes rather than through time-variation in the frequency of adjustment. This implies that a Calvo model which exogenously matched the frequency of adjustment across time would substantially understate the time-variation in price flexibility in the data.

Many recent papers have used fully specified structural models to argue that the distribution of price changes has important implications for aggregate price flexibility. For example, Midrigan (2011) and Alvarez et al. (2014) show that theory assigns a large role to the price change distribution in shaping the average response of inaction to nominal shocks and Vavra (2014) argues that similar mechanisms lead to increases in price flexibility during recessions. Such structural models necessarily impose strong assumptions on the shocks which hit firms and thus on the evolution of desired price changes across time. This in turn implies that they are unable to fully replicate the complicated evolution of the distribution of observed price changes across time. In contrast, our model is flexible. It imposes no restrictions on the evolution of firms' desired price changes across time but still allows us to construct a measure of price flexibility at a point in time. This flexibility comes at a cost: our framework is less useful for making predictions (aside from the fact that pricing moments are somewhat persistent, so that knowledge of the distribution today is informative for the distribution tomorrow) or for assessing counterfactuals under alternative policy environments. We have no theory for the evolution of price gaps and instead simply estimate their distribution period by period. So while our methodology provides a useful way of summarizing the complicated distribution of price changes and how this will respond to shocks on impact, we have less to say about how variables will evolve after impact or how distributions will potentially change in response to changes in policy. That is, our framework provides a historical view of price flexibility which requires minimal structure, but is somewhat sensitive to lucas critique arguments when trying to do predictive analysis.

While specifying a full structural model is important if one wants to understand what drives firms' pricing decisions or for performing counterfactual analysis, one contribution of our paper is showing that an important component of the nominal transmission mechanism can be measured with more minimal structural assumptions. Measuring aggregate price flexibility at a particular point in time can be done without an explicit model of the evolution of price distributions across time. In particular, given a specification for the hazard of price adjustment and the distribution of firms' desired price changes at a moment in time, aggregate price flexibility is fully revealed by the observed distribution of firms' actual price changes at

that same point in time. We apply this identification procedure to BLS CPI, PPI and IPP micro data to create a time-series for price flexibility in each data set and find that in all cases it is strongly countercyclical.

Our work relates to many existing, largely empirical papers which document facts about the distribution of price changes. Typically, these papers focus on one data set, whereas we focus on the time-series properties of a broad set of statistics in multiple data sets covering different points in the supply chain. For example, Klenow and Malin (2010) document many interesting facts about prices, but concentrate solely on CPI data and do not focus on the time-series properties of higher moments of the price change distribution. Chu et al. (2015) study the distribution of price changes in the U.K., but exclusively study the CPI and do not discuss implications for price flexibility and monetary policy. Vavra (2014) is the most closely related paper, and our work is distinguished in several ways: Vavra (2014) studies only CPI data and focuses mostly on the variance of price changes rather than broader features of the price distribution studied in our analysis. On the theoretical front, as mentioned above, his model imposes much stronger structural assumptions while our analysis uses a more flexible accounting framework to describe price flexibility.

The remainder of the paper proceeds as follows: Section 2 contains our main empirical findings. Section 3 discusses the implications for time-varying flexibility using the simple, flexible structure of Caballero and Engel (2007). Section 4 lays out our main results which are that price flexibility varies significantly over time and is strongly countercyclical. Finally, Section 5 concludes.

2 Data

2.1 Data Sources

We analyze three sources of micro data collected by the BLS, and we describe each data set in brief. The restricted access *CPI research database* collected by the Bureau of Labor Statistics (BLS) contains individual price observations for the thousands of non-shelter items underlying the CPI and spans the period 1988-2012. Prices are collected monthly only in

New York, Los Angeles and Chicago, and we restrict our analysis to these cities⁶ to ensure the representativeness of our sample. The database contains thousands of individual "quote-lines" with price observations for many months. Quote-lines are the highest level of disaggregation possible and correspond to an individual item at a particular outlet. An example of a quote-line collected in the research database is 2-liter coke at a particular Chicago outlet. These quote-lines are then classified into various product categories called "Entry Level Items" or ELIs. The ELIs can then be grouped into several levels of more aggregated product categories finishing with eleven major expenditure groups: processed food, unprocessed food, household furnishings, apparel, transportation goods, recreation goods, other goods, utilities, vehicle fuel, travel, and services. For more details on the structure of the database see Nakamura and Steinsson (2008).

We use confidential micro data on import prices collected by the Bureau of Labor Statistics for the period 1994-2011. This data is collected on a monthly basis and contains information on import prices for very detailed items over time. This data set has previously been used by Gopinath and Rigobon (2008), Gopinath and Itskhoki (2010), Neiman (2010), Berger et al. (2012) and Berger and Vavra (2015). Below, we provide a brief description of how the data is collected. The target universe of the price index consists of all items purchased from abroad by U.S. residents (imports). An "item" in the data set is defined as a unique combination of a firm, a product and the country from which a product is shipped. The target universe of the price index consists of all items purchased from abroad by U.S. residents (imports). An "item" in the data set is defined as a unique combination of a firm, a product and the country from which a product is shipped. Price data are collected monthly for approximately 10,000 imported items. The BLS collects "free on board" (fob) prices at the foreign port of exportation before insurance, freight or duty are added, and almost 90% of U.S. imports have a reported price in dollars. Following the literature, we restrict our analysis to these dollar denominated prices. The BLS collects prices monthly using voluntary confidential surveys, which are usually conducted by mail. Respondents are asked for prices of actual transactions that occur as close as possible to the first day of the month. For more details about the IPP data set see Gopinath and Rigobon (2008).

The PPI Research Database contains a panel of raw data from the production firms used to construct the PPI. The earliest prices in the database are from the late 1970s. For most

⁶Prices were also collected monthly in San Francisco and Philadelphia from 1988-2012. We do not include these cities because we want to have a consistent sample for the entire life of the data set (1988-2009).

categories, however, the sample period begins some time during the early to mid 1980s. For the period 1981-2012 (the period we focus on), the PPI Research Database contains data for categories that constitute greater than 90% of the value weight for the Finished Goods PPI. For more details see Nakamura and Steinsson (2008).

Like the IPP, the PPI is collected by BLS through a representative survey of firms. This methodology introduces greater concerns about data quality than in the CPI where BLS agents actually observe prices of products on the shelf. In order to address these concerns the BLS focuses on only collecting actual transaction prices. Specifically, the BLS requests the price of actual shipments transacted within a particular time frame. It is important to note that many of the transactions for which prices are collected as part of the IPP and PPI are a part of implicit or explicit long-term contracts between firms and their suppliers. The presence of such long-term contracts makes interpreting the IPP and PPI data more complicated than interpreting CPI data. This is less of a concern in the IPP because we only use market based transactions, however, this concern remains in the PPI data.

2.2 Variable definitions

Much of the recent literature has discussed the difference between sales, regular price changes and product substitutions. In our analysis, we focus on regular price changes, excluding sales and product substitutions. We use the series excluding sales and product substitutions as our benchmark for two reasons: 1) Eichenbaum et al. (2011) and Kehoe and Midrigan (2015) argue that the behavior of sales is often significantly different from that of regular or reference prices and that regular prices are likely to be the important object of interest for aggregate dynamics. Thus, we choose to exclude sales in our benchmark analysis. However, it is important to note that sales are infrequent in IPP and PPI data, and our results are largely similar if we include sales in the CPI analysis rather than excluding them. 2) Product substitutions require a judgement on what portion of a price change is due to quality adjustment and which component is a pure price change. Thus, this introduces measurement error in the calculation of price changes at the time of product substitution. Bils (2009) shows that these errors can be substantial. For this reason, we exclude product substitutions from our benchmark analysis.

We define the price change of item i at time t as $dp_{i,t} = \log \frac{p_{i,t}}{p_{i,t-1}}$.⁷ Then, using aggregation weights provided by the BLS, it is straightforward to calculate the cross-sectional distribution of log price changes for each month and investigate how it varies over the business cycle. Following Vavra (2014), we focus separately on the distribution of non-zero price changes and frequency rather than computing statistics for the distribution of price changes including zeros.

2.3 Data facts

Figure 1 plots the distribution of (non-zero) price changes across time for the CPI, IPP and PPI. In particular, we plot the 10th, 25th, 37.5th, 50th, 62.5th, 75th and 90th percentiles of the distribution of price changes for all three data sets along with (gray) NBER recession bars. The first observation is that the average size of a price change is large in all three datasets: the mean interquartile range (the 75th percentile minus the 25th percentile) is around 7%. Second, the distribution of price changes varies significantly over time. This variation is most dramatic for the CPI, but is still substantial for the IPP and PPI.⁸ These time-series movements do not occur at random; they are correlated with the business cycle. In particular, the average size of price changes falls, and the frequency and dispersion of price changes rises during recessions.

Table 1 formally documents the business cycle properties of price-setting at quarterly frequencies.⁹ Since there is some high frequency noise in the data, and because low frequency trends can introduce spurious correlation, our preferred specifications focus on variation at business cycle frequencies. In particular, the top panel shows how bandpass filtered (BP) frequency and the first four moments of price changes vary with GDP growth rates. The middle panel reports the same correlations using a Hodrick-Prescott filter (HP) to eliminate low frequency trends and a 3-quarter moving average filter (MA) to eliminate high frequency

⁷In addition to this measure of the size of a price change, we also computed the price change size as $dp = 2 \frac{(p_t - p_{t-1})}{(p_t + p_{t-1})}$, which has the advantage of being bounded and thus less sensitive to outliers. We also investigated using residuals from a regression of the current price on the previous price as a measure of the size of price changes. Results with these two alternative measures are very similar to the results reported below and so are excluded for brevity. The results are available from authors upon request.

⁸The larger time variation in the CPI might be related to the fact that CPI is not affected by long-term contracts.

⁹In Appendix A we report results for various percentiles of the distribution. We also show that the same time-series relationships that we document below are also present in monthly data.

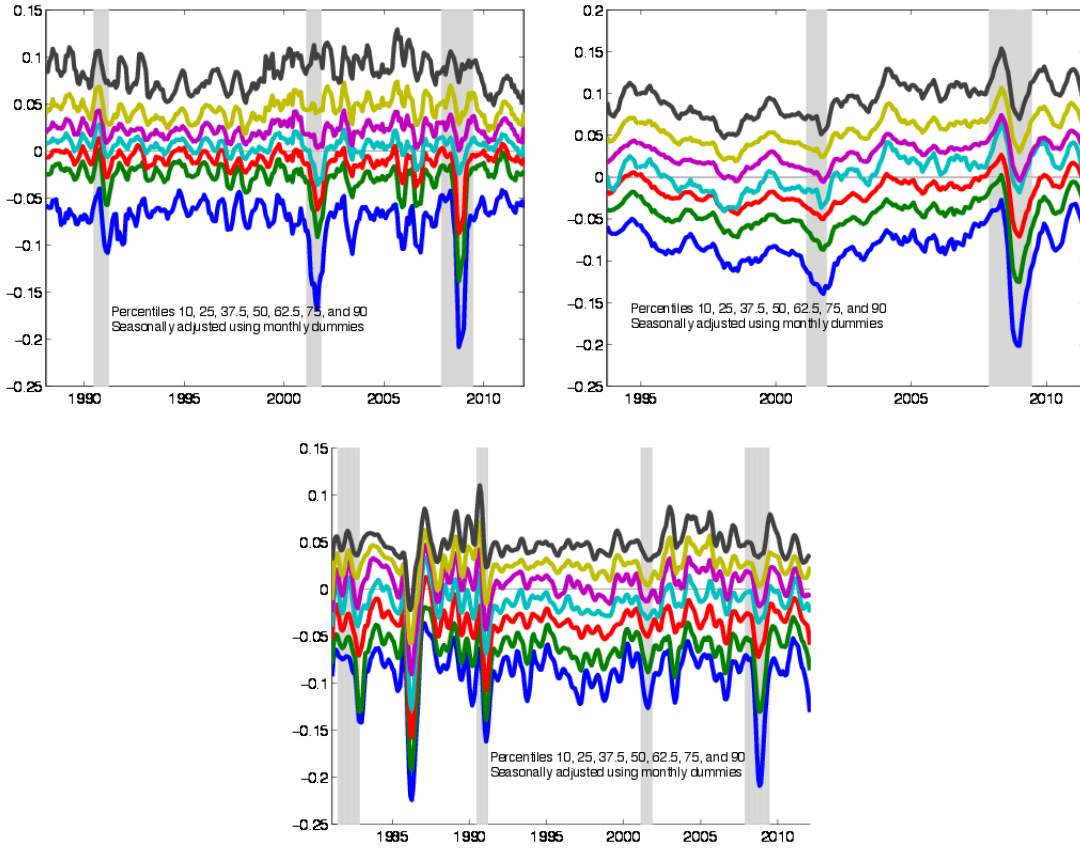


Figure 1: Distribution of price changes across time

variation. Finally, the bottom panel reports results for unfiltered data and shows that all patterns are largely similar.¹⁰ In the Appendix, we also show that similar conclusions obtain when regressing variables on recession indicators.

We document two main facts. The first fact is that the frequency of adjustment is counter-cyclical. Vavra (2014) first documented this fact for the CPI but we see here that it holds at all stages of production. The second fact is that price dispersion is strongly countercyclical. Table 1 presents results for three measures of price change dispersion: the standard deviation (XSD), the interquartile range (IQR) and the difference between the 90th and 10th percentile of the distribution of price changes, and all three measures tell the same story. In almost

¹⁰In the unfiltered specification, we detrended all data with a quadratic trend to eliminate spurious trend correlations, but results are similar with no detrending.

Table 1: Business Cycle Correlations of Pricing Moments

	Freq	XSD	IQR	90-10	Skew	Robust-Skew	Kurt	Robust-Kurt	Obs
BP Filtered									
CPI	-0.53***	-0.59***	-0.65***	-0.52***	0.20	0.14	0.39***	0.40***	76
IPP	-0.36**	-0.66***	-0.61**	-0.68***	0.16	0.55***	0.25***	-0.17	51
PPI	-0.35***	-0.57***	-0.48***	-0.56***	0.04	-0.20	0.17**	-0.08	105
HP + MA Filtered									
CPI	-0.52***	-0.61***	-0.64***	-0.64***	0.15**	0.12	0.38***	0.34***	96
IPP	-0.40**	-0.63***	-0.62***	-0.65***	0.24	0.50***	0.20	0.01	71
PPI	-0.28***	-0.40***	-0.31*	-0.39**	0.09	-0.17	0.06	-0.09	125
Unfiltered									
CPI	-0.35***	-0.46***	-0.45***	-0.44***	0.10	0.14***	0.14*	0.07	96
IPP	-0.26	-0.52***	-0.54***	-0.51**	0.12	0.35**	0.05	0.11	71
PPI	-0.27***	-0.27	-0.22*	-0.24	0.09	-0.10	-0.02	-0.02	125

Each cell displays the correlation of a particular pricing moment in a particular data set with GDP growth. BP uses a baxter king(6,32,10) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 1600 and a 3 quarter moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is quarterly. Robust-Skew = $(P_{90} + P_{10} - 2P_{50}) / ((P_{90} - P_{10}))$. Robust-Kurt = $(P_{90} - P_{62.5} + P_{37.5} - P_{10}) / ((P_{75} - P_{25}))$. Standard errors are computed using a Newey-West correction with optimal lag length. * = 10%, ** = 5%, *** = at least 1% significance.

all of the specifications, the dispersion of price changes is significantly negatively correlated with the business cycle. This fact is consistent with the large body of evidence presented in Bloom et al. (2012) documenting that many variables exhibit countercyclical dispersion and shows that this fact holds in a variety of pricing series.¹¹

The last four columns of Table 1 show that, across datasets, there is a less consistent relationship between the third and fourth moments of the distribution of price changes and the business cycle. The standard moment-based skewness exhibits no notable cyclicity in any data set. The robust, quantile based measure of skewness is strongly procyclical in the IPP but not in the CPI or PPI. The kurtosis of price changes measured using both moments and more robust percentiles is strongly procyclical in the CPI but is not robustly so in

¹¹Vavra (2014) showed it held in the CPI; Berger and Vavra (2015) showed it holds in the IPP. In this paper we show it holds in the PPI as well.

IPP or PPI. This shows the importance of jointly analyzing pricing data at various stages of production, as facts gleaned in one data set may not be representative of more general price-setting patterns. A large recent literature has emerged trying to match features of the kurtosis of price changes in CPI data, but here we show that the time-series behavior of kurtosis in the CPI is somewhat unique.¹²

Table 2: Correlation of Pricing Moments with Frequency of Adjustment

	XSD	IQR	90-10	Skew	Robust-Skew	Kurt	Robust-Kurt	Obs
BP Filtered								
CPI	0.52***	0.55***	0.43***	0.16	0.18	-0.54***	-0.44*	76
IPP	0.44*	0.47**	0.43*	-0.24	-0.41**	-0.05	-0.09	51
PPI	0.41***	0.40***	0.40***	-0.16	0.47*	0.08	0.13	105
HP + MA Filtered								
CPI	0.50***	0.55***	0.41***	-0.10	0.09	-0.54***	-0.52**	96
IPP	0.19	0.23	0.22	-0.20	-0.24*	-0.09	-0.01	71
PPI	0.26**	0.30**	0.30**	-0.04	0.38***	0.01	0.12	125
Unfiltered								
CPI	0.36***	0.43***	0.35***	-0.06	0.01	-0.35***	-0.27*	96
IPP	0.12	0.16	0.16	-0.15	-0.18**	-0.04	0.01	71
PPI	0.18**	0.33**	0.25**	-0.02	0.20**	-0.04	0.03	125

Each cell displays the correlation of the frequency of adjustment in a particular data set with the corresponding moment in the same data set. BP uses a baxter king(6,32,10) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 1600 and a 3 quarter moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is quarterly. Robust-Skew = $(P_{90} + P_{10} - 2P_{50}) / ((P_{90} - P_{10}))$. Robust-Kurt = $(P_{90} - P_{62.5} + P_{37.5} - P_{10}) / ((P_{75} - P_{25}))$. Standard errors are computed using a Newey-West correction with optimal lag length. * = 10%, ** = 5%, *** = at least 1% significance.

Table 2 documents the correlation of pricing moments with the frequency of adjustment. In price-setting models, the frequency of adjustment is typically closely related to the amount of aggregate price flexibility, so it is useful to explore the relationship between the price change distribution and frequency. The first three columns of Table 2 show that the frequency of adjustment is significantly and positively correlated with price dispersion in all specifications

¹²All data sets exhibit excess kurtosis on average, as emphasized by Midrigan (2011).

for the CPI and PPI. The relationship is less consistent for the IPP, however, the point estimates are always positive even when not statistically significant. The next two columns document the relationship between skewness and frequency. Overall, the relationship is idiosyncratic to the specific data set: skewness and frequency are positively correlated in the PPI, negatively correlated in the IPP and there is no time-series relationship in the CPI. Finally, the last two columns of table 2 show that there is a strong negative relationship between kurtosis and frequency in the CPI, but again, this pattern is unique to the CPI: frequency and kurtosis are uncorrelated in the IPP and PPI.

To summarize the more robust patterns in the above tables: we find strong evidence that the frequency and price dispersion are both countercyclical and positively related to each other in all three data sets. Conversely there is no robust relationship between higher moments and the business cycle across data sets: we find that skewness is procyclical only in the IPP and kurtosis is procyclical only in the CPI.

While we find it informative to highlight these particular patterns, it is clear that there are many moments of the price distribution upon which one could focus. In Appendix A, we report additional patterns for ten percentiles of the price change distribution. What should we take away from these empirical facts, and why should we care about matching them? In the next section, we explore the implications of these price facts for the effectiveness of monetary policy, and show that the complicated high-dimensional distribution of price changes at a point in time can be summarized by a useful measure of price flexibility. When viewed through the lens of this price flexibility measure, matching the distribution of price changes across time has important implications for the cyclicity of aggregate price flexibility.

3 Accounting framework

3.1 Basic setup

We use the generalized Ss model developed by Caballero and Engel (2007) to formalize the link between changes in the timing of individual price adjustments and macro price flexibility. The main appeal of this framework is that it flexibly encompasses several pricing mechanisms commonly used in macroeconomic applications in a parsimonious way as well as providing a good fit to the micro data.

First, some preliminaries. There are both aggregate and idiosyncratic shocks, We assume

that shocks to the growth rate of money (or nominal demand) Δm_t are i.i.d with mean μ_A and variance σ_A^2 . Firms also face idiosyncratic (productivity and demand) shocks, v_{it} , which are i.i.d. with potentially time-varying variance σ_I^2 . No assumptions are made regarding the common distribution of idiosyncratic shocks. These shocks are independent across firms and from the aggregate shock. Given these assumptions, the optimal flexible price for firm i (the “desired price”) is:

$$\Delta p_{it}^* = \Delta m_t + v_{it}$$

That is, conditional on adjusting, firm i adjusts to innovations in all the shocks since it last adjusted. Define the price gap as $x \equiv p_{i,t-1} - p_{it}^*$, the difference between firm i 's, current price and the price it would choose if it temporarily faced no adjustment costs. The price gap is the relevant state variable in this pricing model since firms are more likely to adjust the larger the absolute size of the gap.

We further assume that there are i.i.d. idiosyncratic shocks to adjustment costs, ϖ , drawn from a distribution $G(\varpi)$. Integrating over all possible realizations of these adjustment costs, we obtain an adjustment hazard, $\Lambda(x)$, defined as the probability of adjusting—prior to knowing the current adjustment cost draw—by a firm that would adjust by x , if its adjustment cost draw were zero. It is straightforward to prove that $\Lambda(x)$ is decreasing for $x < 0$ and increasing for $x > 0$. This is referred to by Caballero and Engel (2007) as the increasing hazard property: the probability of adjusting is increasing in the absolute size of a firm's price gap.

A nice feature of this generalized Ss framework is that it nests many standard models as special cases. For example, a standard menu cost model is obtained when $G(\varpi)$ has all of its mass at one point. The Calvo model ($\Lambda(x) = \lambda$ for all x) is obtained when $G(\varpi)$ has mass λ at $\varpi = 0$ and mass $1 - \lambda$ at a very large value of ϖ . The model also has empirical relevance: it gives rise to infrequent and lumpy price adjustment, which is a central feature of the price data that we seek to reproduce. It can also well match the observed distribution of price changes, and it is consistent with the evidence in Eichenbaum et al. (2011) that firms are more likely to adjust prices that are out of line with marginal cost.

The model also aggregates nicely. Denote by $f_t(x)$ the cross section of price gaps immediately before any adjustments take place at time t . Inflation is given by the simple formula:

$$\pi_t = - \int x \Lambda_t(x) f_t(x) dx$$

Define $F = \frac{\partial \pi_t}{\partial \Delta m_t}$ as the price flexibility index. It measures the price response upon impact to a nominal shock. When log nominal demand follows a random walk, a common assumption in the literature (Woodford (2003) Nakamura and Steinsson (2010); Vavra (2014)), the flexibility index is a summary measure of monetary non-neutrality because the larger is the (price) flexibility index, the smaller is the output response. Thus knowledge of the flexibility index is a useful proxy for the current efficacy of monetary policy. Fortunately, Caballero and Engel (2007) show how to derive the flexibility index for the generalized Ss model:

$$F = \lim_{\Delta m_t \rightarrow 0} \frac{\partial \pi_t}{\partial \Delta m_t} = \int \Lambda_t(x) f_t(x) dx + \int x \Lambda'_t(x) f_t(x) dx \quad (1)$$

The flexibility index can be decomposed into two components: an intensive margin and an extensive margin. The first term is the intensive margin, which measures the part of inflation coming from firms that would have adjusted even absent the monetary shock. This margin is present in both the *Ss* and Calvo models. The second term is unique to state-dependent models. The extensive margin refers to the amount of inflation coming from firms whose decisions to adjust are altered by the monetary shock. This includes both firms who would have kept their price constant and instead change it, as well as firms who would have changed prices but now choose not to. The extensive margin is only present in *Ss* models since in a Calvo model $\Lambda'_t(x) = 0$.

When will each of these margins be more important? Inspecting the expression for the intensive margin shows that this component is equal to the frequency of adjustment. The more firms that are adjusting absent the aggregate shock, the greater the aggregate price response to that shock through the intensive margin. The extensive margin grows with the number of firms near the margin of adjustment (firms with large $\Lambda'_t(x)$). In addition, the extensive margin is amplified if firms near the margin of adjustment also have large values of $|x|$: if the difference between adjusting and not adjusting grows, then triggering firms to switch their adjustment decisions will have a bigger effect on the overall price level.

The flexibility index is our main object of interest as it tells us how the the price response upon impact to a nominal shock varies over time. Moreover, it is also potentially useful for discriminating between price setting models. Equation (1) shows that if one knew both the hazard function, $\Lambda_t(x)$, and the distribution of price gaps, $f_t(x)$, one could estimate the flexibility index at each moment of time. Of course, both of these objects are unobservable. However, with some minimal structure and data on observed price changes, we are able to

identify this object. First, the product of $\Lambda_t(x)$ and $f_t(x)$ relates unobservable price gaps of size x to the observable distribution of price changes of size x . We put further structure on the problem by assuming that the hazard rate is quadratic (until the point at which firms adjust with probability 1), since this parsimoniously captures the state-dependence of the Ss model while also nesting the Calvo model.

$$\Lambda(x) = \min(a_t + b_t x^2, 1)$$

What determines the distribution of price gaps $f_t(x)$? In traditional structural approaches, one assumes some simple process for v_{it} , and combines this assumption with $\Lambda(x)$ to derive the evolution of $f_t(x)$. For example, Caballero and Engel (2006) assume that v is drawn from a time-invariant normal distribution, while Midrigan (2011) assumes a time-invariant leptokurtic distribution. Permuting these shock processes with the adjustment hazard produces some distribution of price gaps $f(x)$. One then estimates the underlying shock process to match the stationary distribution of price changes. This approach has the advantage of being highly parsimonious since it estimates a limited number of parameters. It is also useful for performing counterfactual exercises in response to changes in the policy environment, under the assumption that the distribution of v is policy invariant. However, it also has an important disadvantage: the imposition of this structure implies strong restrictions on the evolution of price gaps and thus the distribution of price changes across time. Given these tight restrictions and the small number of parameters estimated, this means these models can at best very roughly capture the complicated evolution of the price change distribution documented in the previous section.

In order to try to more directly assess the implications of this complicated price distribution for aggregate price flexibility, we take a different approach that tries to estimate outcomes rather than the underlying shock process. In particular rather than trying to estimate underlying structural parameters of some shock process v , we instead directly estimate a flexible functional form for the distribution of price gaps $f_t(x)$. Given that we have much less theoretical guidance for shape of distribution of price gaps, we leave it largely unrestricted. In our primary specification, we allow $f_t(x)$ to follow a Pearson Type 7 Distribution, which means it has an unrestricted mean, variance, skewness and kurtosis. Given these 4 parameters together with the 2 parameters of the adjustment hazard, equation (1) delivers the price response upon impact at each moment in time. In addition to this functional form, we also

provide additional results following Guvenen, Ozkan and Song (2014) in using a mixture of normals to provide a flexible parameteric form for $f_t(x)$. While it might seem that there is little difference between estimating the distribution of v and that of $f_t(x)$, the key distinction is that the distribution of v is assumed to be time-invariant¹³, while we estimate a separate $f_t(x)$ in each period.¹⁴ That is, the main distinction between the two approaches is on the restrictions placed on parameter variation across time. Our approach estimates $6 \times t$ parameters while a structural approach assuming a time-invariant Pearson distribution for v and a stationary hazard Λ would estimate only 6 parameters.




Figure 2: Effect of parameters on Frequency and Price Flexibility

How does underlying variation in the parameters governing Λ and f affect both the observed distribution of price changes and aggregate price flexibility? We illustrate this by

¹³Or to only vary across time in extremely simple ways.

¹⁴These approaches are exactly equivalent if one allows the distribution of v and Λ to vary across time with equivalent degrees of freedom.

picking some sample parameters,¹⁵ and then varying moments of the price gap distribution and reporting both the frequency of adjustment and the price response upon impact for these various parameter values. The top panel of Figure 2 shows how the frequency and aggregate flexibility vary with the standard deviation of f . It is obvious from the figure that increases in the standard deviation of desired price changes increase both frequency and aggregate price flexibility. Importantly the effect is non-linear: the effect on aggregate price flexibility is highly convex in the standard deviation of the price gap distribution. The logic behind these effects is that an increase in standard deviation of the distribution of price gaps means that there is more mass in the region of the state space where there is a higher probability of adjustment. That is, both the intensive and extensive margins are strengthened.¹⁶

In contrast, the middle panel of Figure 2 shows that there is little relationship between the skewness of $f_t(x)$ and either the frequency of adjustment or price flexibility. Finally, the bottom panel shows that there is a negative relationship between kurtosis and both the frequency of price changes and price flexibility.¹⁷ Why? Higher kurtosis means that the distribution of price gaps has fatter tails relative to a normal distribution. That is, there are both more price gaps near zero and more price gaps at extreme values. Since the hazard of adjustment as a function of the price gap is bounded above by 1, this limits the degree to which the price gaps at the extremes can raise frequency. That is, firms with large price gaps will adjust anyway, while simultaneously pushing more mass towards zero lower the frequency of adjustment. Higher kurtosis also reduces the fraction of intermediate firms who are on the margin of adjustment, which lowers price flexibility through a decline in the extensive margin.

Identification

Thus far, we have shown that variation in the moments of the (unobservable) distribution of price gaps can be mapped through our flexible parametric model into the frequency of adjustment and aggregate flexibility. The next step is to show that there is a mapping from the unobservable distribution of gaps to the observable distribution of price changes.

¹⁵We choose [mean, std. deviation, skewness, kurtosis, a, b]=[.005,.05,0,6,90,.01] for this illustration, as it produces moments in line with the average distribution of price changes.

¹⁶This conforms with the more structural results in Vavra (2014).

¹⁷This is consistent with Midrigan (2011).

Table 3 shows the relationship between parameters of the gap distribution and the (observed) distribution of price changes.

Table 3: Correlation between $f(x)$ parameters and distribution of price changes

Gap Parameter	Observed Price Change Moment				
	Frequency	Inflation	Std. Deviation	Skewness	Kurtosis
Mean	-0.00	0.99***	-0.00	-0.65***	0.00
Std. Deviation	0.99***	-0.02	1.00***	-0.01	-0.83***
Skewness	-0.05	1.00***	-0.01	0.79***	0.02
Kurtosis	-0.70***	-0.04	0.96***	0.05	0.92***

The first column shows that variation in the mean and skewness of the gap distribution does not affect the frequency of adjustment. In contrast, the standard deviation of the gap distribution is strongly positively correlated with frequency while the kurtosis of the distribution is strongly negatively correlated with frequency. This reinforces what we observed in figure 2. The next fact which jumps out is that changes in each moment of the gap distribution are strongly positively correlated with the same moment in the distribution of observed price changes. For example, variation in the mean of the gap distribution implies similar variation in the level of inflation in the distribution of observed price changes. This shows that we can use variation in the moments of the distribution of price changes to identify movements in the unobserved distribution of price gaps. Finally, we see that while variation in each parameter is strongly informative for a particular moment, it also induces variation in other moments of the distribution. For example, variation in the mean of the distribution of price gaps affects skewness, while variation in the standard deviation of the price gap distribution affects kurtosis. In sum, Table 3 shows that there is a tight mapping between moments of the unobserved distribution of price gaps and various moments of the observed distribution of price changes so that the latter is useful for identifying the former.

How restrictive are our identifying assumptions that within each month, the hazard is quadratic and that the distribution of price gaps follows a four parameter Pearson Type 7 distribution? It is clear that our approach is more flexible than more typical structural models, but it is more restrictive than a fully non-parametric approach. However, it is also clear that identification requires *some* parametric assumptions, as a fully non-parametric gap distribution and hazard is unidentified. This is because if one allows for a fully non-

parametric gap distribution, then one can always perfectly replicate the data with a Calvo model by setting $b = 0$, choosing the frequency of adjustment to be equal to a , and picking the gap distribution to correspond to the actual observed distribution of price changes in each period. However, this is not a particularly appealing model of price-setting for several reasons. First, this model would have essentially no actual predictive content. A model which allows for a completely arbitrary distribution of shocks to explain observable data essentially explains nothing. Similarly, it would be extremely difficult to construct any model that generates such complicated distributions of firms' desired price changes. Finally, there is strong empirical evidence of state-dependence in micro pricing data, so that a Calvo model with a complicated gap distribution seems at odds with the data.¹⁸

Conversely, our approach is substantially more flexible than more traditional structural approaches such as those in Vavra (2014) and Alvarez and Lippi (2014). We think there is significant value in exploring the implications of a less restrictive model for aggregate price flexibility. Thus, while our assumption that the gap process is determined by a parametric distribution at each point is restrictive relative to a fully non-parametric gap distribution, it is significantly less restrictive than previous structural frameworks. Within a period, we impose a highly flexible parametric functional form for price gaps¹⁹, but across periods this object can evolve in a fully unrestricted way. Standard structural models impose strong restrictions on the relationship between distributions at a point in time and how they evolve across time.²⁰ We believe it is interesting to take an intermediate approach between fully structural and completely non-parametric approaches and explore its implications for aggregate price flexibility.

¹⁸Midrigan (2010) argues that U.S. manufacturing pricing data are much more consistent with state dependent models than with time dependent ones. Midrigan (2012) and Nakamura and Steinsson (2010) use structural approaches and find that micro price data are consistent with state dependent pricing. Gagnon (2010) and Alvarez and Lippi (2013) use evidence from high inflation episodes from Mexico and Argentina respectively and provide very strong empirical evidence that price setters exhibit state dependence in price. Most directly, Eichenbaum, Jaimovich and Rebelo (2011) show that prices are much more likely to adjust when firms' price gaps (as measured by the deviation in their markup from average) are large.

¹⁹In our baseline we use a Pearson Type 7 distribution, but we also show that results are similar using a time-varying mixture of normals.

²⁰For example, the closest structural analogue is contained in Caballero and Engel (2006). They use the same accounting framework as in Caballero and Engel (2007) and impose the structural assumption that idiosyncratic shocks are normal with mean zero and constant variance, which they try to estimate to match the average distribution of price changes in the CPI. If one instead allows for an arbitrary shock process then these more structural models also have essentially no content.

4 Results and implications

We now use the theoretical framework to explore the level and time-variation in aggregate price flexibility for the CPI, IPP and PPI. We assume that the hazard function is a quadratic in price gaps x which is consistent with both the increasing hazard property of Ss models while also nesting the Calvo model. Since we have less guidance for the structure of the distribution of price gaps, we assume in our baseline results that the distribution follows a Pearson Type 7 distribution, which leaves the first 4 moments unrestricted. This means we have six parameters to estimate: the intercept and slope coefficient of the hazard function and the first four moments of the price gap distribution. We also show below that a version of the model where we assume that the distribution of price gaps is given by the mixture of two normal distributions instead of a Pearson Type 7 distribution delivers similar conclusions.

In our baseline specification, we estimate these six parameters period-by-period using seven moments, M_t , for identification:

$$M_t = [freq_t, mean_t, var_t, skewness_t, kurtosis_t, median_t, IQR_t]$$

The seven moments are the frequency of adjustment as well as the mean, median, variance, skewness, kurtosis and interquartile range of the distribution of price changes.²¹ Each period we minimize a quadratic form of these moments M and find the parameters which provide the best fit.²² That is we find the parameters which minimize $M'WM$, where W is a weight matrix. For our baseline specification we weight each moment equally.²³ Once we have specified the quadratic form, we minimize it period-by-period. Then we compute aggregate price flexibility - the price response upon impact to a nominal shock - using equation (1) and analyze how it co-varies with the business cycle and the frequency of adjustment.

²¹Results are similar if instead on uses 10 rather than 7 moments. The moments we added were the 10th, 25th, 75th and 90th percentile of the distribution of price changes, and the interquartile range was removed because it was redundant.

²²We experimented with multiple quadratic forms but we found the most stable results when the odd moments were specified in percentages, $\left(\frac{m^{sim} - m^{data}}{m^{data}}\right)^2$ and the even moments as $(m^{sim} - m^{data})^2$. The reason is the even moments can be either positive or negative and are often centered around zero. If we specified all the moments in percentage terms, this would lead to us dividing by a number near zero, which lead to unstable estimated. However, most results were robust to specifying all the moments symmetrically, either in percentage terms or as raw quadratic forms.

²³Altonji and Segal (1996) argue that simulated minimum distance estimation often performs better in small samples if an identify weight matrix is used rather than the optimal weight matrix.

We conduct this exercise for all three price series and for all three price filters. Appendix Table 13 reports mean parameter results and goodness of fit for our baseline specification using a Pearson Type 7 distribution to match the seven pricing moments across the three data sets. The fit is slightly worse for PPI data, but is in general quite good for all three data sets. The most important takeaway is that the estimation identifies high values for the quadratic adjustment hazard parameter b in all three data sets. This will imply an important role for extensive margin effects in generating price flexibility.

Table 4: Cyclicalities and Time-Variation in Price Flexibility: Matching Moments

	Cyclicalities	Time-Variation			
		Mean	Std. Dev.	5th	95th
BP Filtered					
CPI	-0.57***	0.11	0.014	0.085	0.140
IPP	-0.27**	0.22	0.020	0.184	0.253
PPI	-0.28*	0.40	0.045	0.317	0.462
HP + MA Filtered					
CPI	-0.56***	0.11	0.013	0.096	0.137
IPP	-0.37***	0.22	0.018	0.188	0.246
PPI	-0.30**	0.40	0.050	0.316	0.475
Unfiltered					
CPI	-0.37***	0.11	0.020	0.084	0.150
IPP	-0.31**	0.22	0.025	0.172	0.258
PPI	-0.30***	0.40	0.087	0.278	0.541

This table shows results for the Pearson Type 7 distribution targeting M moments of the price change distribution. The first column shows the correlation between the price flexibility index and GDP growth. In the first panel, series are filtered using a Baxter King (6,32,10) filter. In the second panel, series are filtered using a Hodrick-Prescott(1600) filter and a 3 quarter moving average. In the third panel, series are unfiltered but are detrended with a quadratic trend. “Mean” shows the mean price flexibility over the entire sample. This is computed prior to filtering, since filtered data is mean zero. “Std. Dev.,” “5th” and “95th” shows the standard deviation, 5th and 95th percentile of price flexibility across time, after filtering. Standard errors for the cyclicalities calculation are adjusted for serial correlation using a Newey-West correction with optimal lag length. *=10% significance, **=5% significance, ***=1% significance.

Given this set of parameters, how does price flexibility vary across time? We have two main results, both of which are shown in Table 4. The first fact is that aggregate price

flexibility varies substantially across time. Column 2 shows the mean level of price flexibility while columns 3-5 show the time-series standard deviation of price flexibility together with the fifth and ninety-fifth percentile. The first two panels filter the price flexibility series, leaving only price flexibility variation at business cycle frequencies. Even with this filtering, there is still substantial time-series variation: the 95th percentile of price flexibility is 30-60% larger than the 5th percentile of price flexibility. The time-series standard deviation of price flexibility is typically around 10% of the mean level of price flexibility. The bottom panel shows that, unsurprisingly, price flexibility variation is even larger when using variation at all time frequencies. The 95th percentile of unfiltered price flexibility is 50-95% higher than the 5th percentile of the same series.

The second important fact shown in Table 4 is that price flexibility is strongly counter-cyclical, so that this variation in price flexibility does not occur at random. Column 1 shows that our flexibility index is higher when the growth rate of gdp is lower in all nine specifications. Since this flexibility index measures how responsive prices are to a nominal shock at a moment in time, this implies that monetary policy which increases nominal output by a given amount will be less effective at stimulating real output during recessions.²⁴ Thus monetary policy is least effective during the times when policy makers most want to conduct expansionary policy.

One concern with our estimation approach thus far is that it might be sensitive to the particular pricing moments we chose to target. Targeting centered moments of the distribution of price changes utilizes all information, but it also makes results more sensitive to outliers. In particular, some of the higher moments we choose to target are difficult to estimate in small samples. Targeting these moments in our estimation could thus lead us astray. In order to explore the robustness of our results to these concerns, we run our estimation targeting nine percentiles of the distribution of price changes rather than centered moments of this distribution. In particular, we target the following ten moments:

$$M_t = [freq_t, p1_t, p5_t, p10_t, p25_t, p50_t, p75_t, p90_t, p95_t, p99_t]$$

These moments robustly summarize the distribution of price changes while at the same

²⁴Statements about recessions vs. general cyclicity can be somewhat more sensitive to the exact timing of recession dates. Nevertheless, regressing price flexibility on recession dummies implies significant increases for all CPI and IPP specifications and for PPI specifications with the HP+MA filter. Point estimates are positive but insignificant for the other two PPI filters.

Table 5: Cyclicalty and Time-Variation in Price Flexibility: Matching Percentiles

	Cyclicalty	Time-Variation		
		Mean	Std. Dev.	5th 95th
BP Filtered				
CPI	-0.36	0.13	0.011	0.110 0.144
IPP	-0.27***	0.20	0.019	0.172 0.232
PPI	-0.35***	0.29	0.031	0.248 0.348
HP + MA Filtered				
CPI	-0.37**	0.13	0.011	0.111 0.147
IPP	-0.29***	0.20	0.018	0.171 0.230
PPI	-0.34**	0.29	0.032	0.237 0.342
Unfiltered				
CPI	-0.26*	0.13	0.017	0.101 0.153
IPP	-0.28***	0.20	0.030	0.151 0.249
PPI	-0.25**	0.29	0.053	0.200 0.384

This table shows results for the Pearson Type 7 distribution targeting 10 percentiles of the price change distribution plus frequency. The first column shows the correlation between the price flexibility index and GDP growth. In the first panel, series are filtered using a Baxter King (6,32,10) filter. In the second panel, series are filtered using a Hodrick-Prescott(1600) filter and a 3 quarter moving average. In the third panel, series are unfiltered but are detrended with a quadratic trend. “Mean” shows the mean price flexibility over the entire sample. This is computed prior to filtering, since filtered data is mean zero. “Std. Dev.,” “5th” and “95th” shows the standard deviation, 5th and 95th percentile of price flexibility across time, after filtering. Standard errors for the cyclicalty calculation are adjusted for serial correlation using a Newey-West correction with optimal lag length. *=10% significance, **=5% significance, ***=1% significance.

time being less susceptible to outliers. The results are shown in Table 5 which has the same structure as Table 4. We see that both our stylized facts hold up to targeting different moments: aggregate price flexibility is highly countercyclical and the magnitude of this variation is economically large, especially from peak to trough.²⁵

How sensitive are our results to the functional form we have chosen for the distribution of price gaps f ? In our baseline results, we assume that this distribution follows a Pearson

²⁵Heteroscedasticity is more substantial for the bandpass filtered CPI specification under this specification so that while point estimates remain strongly negative, p-values are approximately 0.15.

Type 7 distribution which allows us to independently choose values for the mean, variance, skewness and kurtosis. Following Guvenen et al. (2014), we have also experimented with instead using a gap distribution which is a mixture of normals. In this case, we have five parameters to estimate: the mean and variance of each normal distribution together with the weight placed on each of the two normal distributions. This mixed normal specification allows for substantial flexibility. For example, it is straightforward to generate distributions with excess kurtosis or to allow for bimodal distributions. Table 6 shows that again, our conclusions are unchanged by using a different functional form for our price gap distribution.

Table 6: Cyclicalilty and Time-Variation in Price Flexibility: Mixed Normal Gap Distribution

	Cyclicalilty	Time-Variation			
		Mean	Std. Dev.	5th	95th
BP Filtered					
CPI	-0.66***	0.07	0.014	0.055	0.103
IPP	-0.54***	0.17	0.014	0.154	0.199
PPI	-0.50***	0.38	0.056	0.310	0.486
HP + MA Filtered					
CPI	-0.67***	0.07	0.013	0.056	0.102
IPP	-0.54***	0.17	0.015	0.152	0.202
PPI	-0.40***	0.38	0.057	0.303	0.483
Unfiltered					
CPI	-0.43***	0.07	0.018	0.048	0.111
IPP	-0.31**	0.17	0.021	0.142	0.221
PPI	-0.33***	0.38	0.102	0.245	0.564

This table shows results for the mixed normal distribution targeting M moments of the price change distribution. The first column shows the correlation between the price flexibility index and GDP growth. In the first panel, series are filtered using a Baxter King (6,32,10) filter. In the second panel, series are filtered using a Hodrick-Prescott(1600) filter and a 3 quarter moving average. In the third panel, series are unfiltered but are detrended with a quadratic trend. “Mean” shows the mean price flexibility over the entire sample. This is computed prior to filtering, since filtered data is mean zero. “Std. Dev.,” “5th” and “95th” shows the standard deviation, 5th and 95th percentile of price flexibility across time, after filtering. Standard errors for the cyclicalilty calculation are adjusted for serial correlation using a Newey-West correction with optimal lag length. **=5% significance, ***=1% significance.

The results thus far show that price flexibility varies substantially across time and is

countercyclical, but they tell us little about the forces which drive this variation. Recall from Section 3.1 that our price flexibility index can be decomposed into both an intensive margin and an extensive margin component. The intensive margin is given by $IM = \int \Lambda_t(x) f_t(x) dx$, which is just equal to the frequency of adjustment. The extensive margin is given by $EM = \int x \Lambda'_t(x) f_t(x) dx$. The intensive margin gives the response of inflation to monetary shocks driven by firms which would adjust prices independently of the shock. The extensive margin reflects the additional inflation effects which arise from changes in the mix of adjusting firms. In a Calvo model, $\Lambda'_t=0$ so that only the intensive margin is active. Put slightly differently, knowledge of the frequency of adjustment is sufficient for determining aggregate price flexibility. In contrast, standard Ss models such as Golosov and Lucas (2007) imply very important roles for the extensive margin. Recall that our hazard is parameterized as $\Lambda(x) = \min(a_t + b_t x^2, 1)$, so that our framework nests a Calvo model as well as models with strong extensive margin effects, depending on the estimated value of b . Thus, the relative strength of these effects in our framework depends on the particular parameters which are identified to match observed price change behavior.

Is the variation we find in $F = IM + EM$ then driven mainly by movements in IM or movements in EM ? In Table 7, we first show that there is a positive correlation between IM and EM but this correlation is far less than 1. This implies that there are significant movements in the extensive margin which occur independently of movements in frequency of adjustment. We can assess the relative importance of movements in frequency versus higher moments using a formal variance decomposition. By definition, $VAR(F) = VAR(IM) + VAR(EM) + 2COV(IM, EM)$. If movements in price flexibility are driven by movements in the intensive margin, as in a Calvo model, then $VAR(IM)/VAR(F)$ should be equal to 1. Table 7 shows that this is clearly not the case. Pure movements in the intensive margin explain only around 20-30% of overall movements in price flexibility. If we also include covariance terms, in which movements in the intensive margin and extensive margin amplify each other, then the share of explained variance rises. However, it remains the case that 30-40% of movements in price flexibility are explained by independent movements in the extensive margin.

Thus, extensive margin effects driven by changes in the distribution of price changes are crucial for understanding price flexibility in our data. It is worth stressing again that this did not have to be the case. For certain parameters, our model nests a Calvo model where all flexibility is driven by time-variation in the frequency of adjustment. However, our model

Table 7: Intensive Margin and Aggregate Price Flexibility

	$corr(IM, EM)$	Match Moments		$corr(IM, EM)$	Match Percentiles	
		$\frac{VAR(IM)}{VAR(F)}$	$\frac{VAR(IM)+2COV(IM,EM)}{VAR(F)}$		$\frac{VAR(IM)}{VAR(F)}$	$\frac{VAR(IM)+2COV(IM,EM)}{VAR(F)}$
BP Filtered						
CPI	0.68***	0.22	0.61	0.69***	0.19	0.58
IPP	0.80***	0.15	0.52	0.83***	0.29	0.74
PPI	0.46***	0.17	0.61	0.80***	0.28	0.68
HP + MA Filtered						
CPI	0.69***	0.21	0.60	0.67***	0.20	0.59
IPP	0.61***	0.16	0.50	0.84***	0.30	0.76
PPI	0.47***	0.26	0.60	0.67***	0.32	0.71
Unfiltered						
CPI	0.70***	0.17	0.56	0.49***	0.24	0.57
IPP	0.45***	0.22	0.51	0.83***	0.31	0.77
PPI	0.49***	0.31	0.61	0.68***	0.35	0.76

This table assesses the contribution of the extensive margin for price flexibility. The first two columns “Match Moments” do this for the Pearson version of the model which matches centered moments while the “Match Percentiles” columns do this for the version of the model which matches percentiles of the observed price distribution. In the first panel, series are filtered using a Baxter King (6,32,10) filter. In the second panel, series are filtered using a Hodrick-Prescott(1600) filter and a 3 quarter moving average. In the third panel, series are unfiltered but are detrended with a quadratic trend. Correlations are computed using standard errors which are adjusted for autocorrelation using a Newey-West correction with optimal lag length. ***=1% significance.

prefers alternative regions of the parameter space in order to fit the observed distribution of price changes, and these alternative parameters imply a crucial role for extensive margin effects. Thus, our quantitative results imply that while one could exogenously calibrate a Calvo model to match the time-series for frequency, this exercise would grossly understate the extent of time-variation in price flexibility. Higher moments of the distribution of price changes vary across time and play a crucial role in determining flexibility.

5 Conclusion

In this paper, we synthesized many new and existing facts on the evolution of price changes across time using a consistent empirical methodology to simultaneously study pricing at three stages of production. Our empirical results are unique in the variety of moments collected as well as in the unification of moments across data sets. This simultaneous study of data from the CPI, IPP and PPI is important: we show that several price-setting statistics are idiosyncratic to particular stages of production. For example, patterns related to skewness and kurtosis are quite sensitive to particular data sets. However, there are also several robust patterns which hold across all three data sets: the frequency of adjustment and all measures of price change dispersion are always positively correlated with each other and negatively correlated with output. More generally, there are large movements in the distribution of price changes across time in each of these data sets.

While our paper is largely empirical, we also use a flexible version of the Caballero and Engel (2007) generalized Ss model to summarize the implications of these complicated pricing moments for aggregate price flexibility and the effectiveness of monetary policy. While there is a large and growing literature using structural models to explore the implications of various pricing moments for aggregate price flexibility, our paper is unique in exploring these implications using a much broader set of moments with more minimal identifying assumptions. Using this generalized Ss framework, we show that in all three data sets: 1) Aggregate price flexibility is countercyclical and 2) The level of time-variation is significant. The 95th percentile of the price flexibility index is often 50% or more larger than the 5th percentile of the price flexibility index. This suggests that the efficacy of monetary policy varies significantly over time with policy being much less effective in recessions than in booms.

References

- Altonji, Joseph G and Lewis M Segal**, “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business & Economic Statistics*, 1996, *14* (3), 353–366.
- Alvarez, Fernando and Francesco Lippi**, “Price Setting With Menu Cost for Multiproduct Firms,” *Econometrica*, 2014, *82* (1), 89–135.
- , **Hervé Le Bihan, and Francesco Lippi**, “Small and Large Price Changes and the Propagation of Monetary Shocks,” *National Bureau of Economic Research Working Paper Series*, 2014, No. 20155.
- Berger, David and Joseph S Vavra**, “Volatility and Pass-through,” *National Bureau of Economic Research Working Paper Series*, 2015, No. 19651.
- , **Jon Faust, John H. Rogers, and Kai Steverson**, “Border prices and retail prices,” *Journal of International Economics*, sep 2012, *88* (1), 62–73.
- Bils, Mark**, “Do Higher Prices for New Goods Reflect Quality Growth or Inflation?,” *The Quarterly Journal of Economics*, may 2009, *124* (2), 637–675.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-eksten, and Stephen J Terry**, “Really Uncertain Business Cycles,” *NBER Working Paper #18245*, 2012.
- Caballero, Ricardo J. and E. M R Engel**, “Price Stickiness in Ss Models: Basic Properties,” *unpublished, MIT*, 2006.
- **and** – , “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics*, 2007, *54* (SUPPL.), 100–121.
- Caplin, Andrew S and Daniel F Spulber**, “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, nov 1987, *102* (4), 703–725.
- Chu, Ba H., Kim P Huynh, David Jacho-Chavez, and Oleksiy Kryvtsov**, “Functional Principal Component Analysis of Density Families with Complex Survey Data on UK Prices,” *Working Paper*, 2015.

- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo**, “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, 2011, *101* (1), 234–262.
- Golosov, M and Robert E. Lucas**, “Menu Costs and Phillips Curves,” *Journal of Political Economy*, apr 2007, *115* (2), 171–199.
- Gopinath, Gita and Oleg Itskhoki**, “Frequency of Price Adjustment and Pass-Through,” *The Quarterly Journal of Economics*, may 2010, *125* (2), 675–727.
- **and Roberto Rigobon**, “Sticky Borders,” *The Quarterly Journal of Economics*, may 2008, *123* (2), 531–575.
- Guvenen, Fatih, Ozkan Serdar, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, jun 2014, *122* (3), 621–660.
- Kehoe, Patrick and Virgiliu Midrigan**, “Prices are sticky after all,” *Journal of Monetary Economics*, oct 2015, *75*, 35–53.
- Klenow, Peter J. and Benjamin A. Malin**, “Microeconomic Evidence on Price-Setting,” in “Handbook of Monetary Economics,” Vol. 3 2010, pp. 231–284.
- Midrigan, Virgiliu**, “Is Firm Pricing State or Time Dependent? Evidence from U.S. Manufacturing,” *Review of Economics and Statistics*, 2010.
- , “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 2011, *79* (4), 1139–1180.
- Nakamura, Emi and Jón Steinsson**, “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, nov 2008, *123* (4), 1415–1464.
- **and** – , “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, aug 2010, *125* (3), 961–1013.
- Neiman, B**, “Stickiness, synchronization, and passthrough in intrafirm trade prices,” *Journal of Monetary Economics*, 2010, *57* (3), 295–308.
- Vavra, Joseph**, “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation*,” *The Quarterly Journal of Economics*, sep 2014, *129* (1), 215–258.
- Woodford, Michael**, *Interest and Prices*, Princeton University Press, 2003.

APPENDIX

This appendix provides a number of robustness checks and results which are suppressed from the body of the text. In Table 1 in the body of the text, we report a number of statistics on the business cycle properties of price moments. We prefer these measures both because they are less sensitive to slight changes in the timing of recessions, and more importantly, because we do not have a large number of recessions in our data set. Using more general correlations with output allows us to exploit much more variation. Nevertheless, Table 8 shows that similar conclusions are obtained when regressing price-setting statistics on recession dummies. For these specifications, we do not perform any filtering beyond a simple quadratic detrending, so they are most comparable to “Unfiltered” panel in Table 1. The coefficients in these tables are reported in units which reflect the proportionate change in statistics during recessions relative to average. For example, in CPI data, during recessions the frequency of adjustment is 20% above the average over the whole sample.

Table 8: Regression of pricing moments on recession dummies

	Freq	XSD	IQR	90-10	Skew	Q-Skew	Kurt	Q-Kurt	Obs
CPI	0.20***	0.19***	0.33***	0.24***	-0.58	-0.06	-0.28**	-0.08**	96
IPP	0.05	0.10***	0.12**	0.12***	1.62***	9.32**	-0.013	0.004	71
PPI	0.10	0.11**	0.11	0.118**	0.23	-0.40	-0.07	-0.005	125

Data is quadratically detrended but otherwise unfiltered. Coefficients represent the relative change in recessions: (Recession Value / Average Value)-1. *=10%, **=5%, ***1% significance.

Tables 9 and 10 extend the analysis in Tables 1 and 2 to a number of additional pricing percentiles. (See also Figure 1 in the text). We concentrate on centered moments in the text since they are more easily interpretable, but the various percentiles in Table 9 and Table 10 are less sensitive to outliers and could potentially be useful targets for business cycle models. In Table 5 in the text, we present results for a version of our pricing model which targets these percentiles instead of the pricing moments in Table 1 and show that we obtain similar conclusions.

Table 9: Business Cycle Correlations of Percentiles of the Price Change Distribution

	<i>p</i> 1	<i>p</i> 5	<i>p</i> 10	<i>p</i> 25	<i>p</i> 375	<i>p</i> 50	<i>p</i> 625	<i>p</i> 75	<i>p</i> 90	<i>p</i> 95	<i>p</i> 99	Obs
BP Filtered												
CPI	0.36**	0.65***	0.58***	0.63***	0.58***	0.34***	0.25**	-0.00	-0.15	-0.38**	0.15	76
IPP	0.38*	0.64***	0.65***	0.58***	0.49***	0.42***	0.33***	0.28**	0.17	0.01	-0.22	51
PPI	0.36**	0.54***	0.55***	0.46***	0.39***	0.43***	0.34***	0.17	-0.09	-0.13	-0.01	105
HP + MA Filtered												
CPI	0.33***	0.60***	0.57***	0.56***	0.52***	0.28**	0.11	-0.12	-0.21***	-0.37***	0.06	96
IPP	0.48**	0.62***	0.64***	0.57***	0.50***	0.44***	0.35***	0.27*	0.16	0.05	-0.08	71
PPI	0.30*	0.36**	0.39**	0.35***	0.28***	0.34***	0.33***	0.17*	-0.03	-0.07	0.04	125
Unfiltered												
CPI	0.22***	0.46***	0.46**	0.38**	0.36**	0.17*	0.05	-0.06	-0.13	-0.18*	0.00	96
IPP	0.37*	0.48**	0.52**	0.48**	0.40*	0.34*	0.28	0.21	0.16	0.02	-0.05	71
PPI	0.22	0.20	0.21	0.16*	0.11	0.14*	0.14	0.00	-0.07	-0.06	-0.01	125

Each cell displays the correlation of a particular pricing moment in a particular data set with GDP growth. BP uses a baxter king(6,32,10) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 1600 and a 3 quarter moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is quarterly. Standard errors are computed using a Newey-West correction with optimal lag length. *=10%, **=5%, ***=at least 1% significance.

Table 10: Correlations of Percentiles of the Price Change Distribution with the Frequency of Adjustment

	<i>p</i> 1	<i>p</i> 5	<i>p</i> 10	<i>p</i> 25	<i>p</i> 375	<i>p</i> 50	<i>p</i> 625	<i>p</i> 75	<i>p</i> 90	<i>p</i> 95	<i>p</i> 99	Obs
BP Filtered												
CPI	-0.13	-0.51***	-0.37**	-0.41*	-0.35	-0.20	-0.04	0.30*	0.35***	0.56***	0.03	76
IPP	-0.43**	-0.40**	-0.40**	-0.33*	-0.23	-0.21	-0.14	-0.04	-0.08	0.00	0.42***	51
PPI	-0.28*	-0.18*	-0.27**	-0.30**	-0.26*	-0.23	-0.19	-0.02	0.24*	0.14	-0.23*	105
HP + MA Filtered												
CPI	-0.19	-0.48***	-0.35**	-0.34*	-0.30*	-0.15	0.08	0.38***	0.30**	0.42***	-0.13	96
IPP	-0.21	-0.22	-0.21	-0.16	-0.10	-0.09	-0.06	-0.02	-0.05	-0.06	0.04	71
PPI	-0.11	-0.03	-0.14	-0.17*	-0.14	-0.15	-0.11	0.08	0.27**	0.20**	-0.13	125
Unfiltered												
CPI	-0.02	-0.34***	-0.28**	-0.23*	-0.17	-0.05	0.08	0.25**	0.22***	0.24***	-0.10	96
IPP	-0.16	-0.14	-0.12	-0.06	0.00	0.03	0.04	0.04	0.02	-0.03	0.01	71
PPI	-0.00	0.05	-0.05	-0.06	0.01	0.02	0.03	0.22*	0.29***	0.19**	-0.07	125

Each cell displays the correlation of a particular pricing moment in a particular data set with GDP growth. BP uses a baxter king(6,32,10) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 1600 and a 3 quarter moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is quarterly. Standard errors are computed using a Newey-West correction with optimal lag length. *=10%, **=5%, ***=at least 1% significance.

All results in the body of the paper are reported using quarterly data since this is the frequency with which GDP is measured and is less sensitive to high frequency noise. However, the raw pricing data is available monthly, so Tables 11 and 12 repeat our analysis at the monthly frequency, using industrial production instead of GDP as our measure of the business cycle.

Table 11: Business Cycle Correlations of Pricing Moments: Monthly

	Freq	XSD	IQR	90-10	Skew	Robust-Skew	Kurt	Robust-Kurt	Obs
BP Filtered									
CPI	-0.57***	-0.71***	-0.76***	-0.63***	0.18**	0.15	0.37***	0.37**	222
IPP	-0.20	-0.44**	-0.42*	-0.48***	0.26**	0.55***	0.07	-0.22	147
PPI	-0.23**	-0.57***	-0.37**	-0.57***	0.08	0.07	0.26	-0.21	306
HP + MA Filtered									
CPI	-0.47***	-0.60***	-0.62***	-0.56***	0.15**	0.16**	0.28***	0.25**	288
IPP	-0.35***	-0.52***	-0.59***	-0.56***	0.32***	0.49***	0.10	0.13	213
PPI	-0.15*	-0.36**	-0.27**	-0.39***	0.10	0.01	0.12	-0.13	372
Unfiltered									
CPI	-0.12**	-0.33***	-0.33**	-0.34***	0.07	0.06	0.14***	0.06	288
IPP	-0.01	-0.33***	-0.38**	-0.37***	0.20***	0.19**	0.08	-0.00	213
PPI	-0.13**	-0.19**	-0.17***	-0.18**	0.06	-0.03	0.01	-0.00	372

Each cell displays the correlation of a particular pricing moment in a particular data set with industrial production growth. BP uses a baxter king(18,96,33) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 129,600 and a 5 month moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is monthly. Robust-Skew = $(P_{90} + P_{10} - 2P_{50}) / ((P_{90} - P_{10}))$. Robust-Kurt = $(P_{90} - P_{62.5} + P_{37.5} - P_{10}) / ((P_{75} - P_{25}))$. Standard errors are computed using a Newey-West correction with optimal lag length. *=10%, **=5%, ***=at least 1% significance.

Table 12: Correlation of Pricing Moments with Frequency of Adjustment: Monthly

	XSD	IQR	90-10	Skew	Robust-Skew	Kurt	Robust-Kurt	Obs
BP Filtered								
CPI	0.54***	0.56***	0.47**	0.14	0.21	-0.52***	-0.38	222
IPP	0.32**	0.37**	0.31**	-0.12	-0.30**	-0.04	-0.04	147
PPI	0.38***	0.41**	0.38**	-0.12	0.39**	0.03	0.07	307
HP + MA Filtered								
CPI	0.41***	0.43***	0.34***	-0.12	0.05	-0.35***	-0.33**	288
IPP	0.16	0.23	0.24	-0.19	-0.21	-0.25*	0.06	213
PPI	0.18**	0.26*	0.26**	-0.00	0.17**	-0.06	0.15*	373
Unfiltered								
CPI	0.24***	0.28***	0.25***	0.01	-0.03	-0.23***	-0.12	288
IPP	0.03	0.10	0.13*	0.02	-0.05	-0.20**	0.07	213
PPI	0.04	0.14	0.11*	0.08	0.10**	-0.07	0.06	373

Each cell displays the correlation of the frequency of adjustment in a particular data set with the corresponding moment in the same data set. BP uses a baxter king(6,32,10) filter. BP uses a baxter king(18,96,33) filter. HP+MA uses a hodrick-prescott filter with smoothing parameter 129,600 and a 5 month moving average. Unfiltered data uses no filters but detrends series using a quadratic trend. All data is monthly. Robust-Skew = $(P_{90} + P_{10} - 2P_{50}) / ((P_{90} - P_{10}))$. Robust-Kurt = $(P_{90} - P_{62.5} + P_{37.5} - P_{10}) / ((P_{75} - P_{25}))$. Standard errors are computed using a Newey-West correction with optimal lag length. *=10%, **=5%, ***=at least 1% significance.

Table 13 displays the mean parameters and goodness of fit for our primary pricing model whose results are shown in Table 4.

Table 13: Mean Parameter Values and Goodness of Fit

	Parameters						Fit	
	mean	xsd	skew	kurtosis	slope	intercept	mean squared error	std. dev errors
CPI	0.0008	0.042	0.016	6.74	16.01	0.026	0.014	0.01
IPP	0.0043	0.053	0.081	9.78	22.25	0.043	0.0085	0.005
PPI	0.0047	0.050	-0.038	6.87	46.44	0.093	0.037	0.19