Multiproduct Firms and Price-Setting: Theory and Evidence from U.S. Producer Prices

by

Saroj Bhattarai, Pennsylvania State University
Raphael Schoenle, Brandeis University

CEPS Working Paper No. 211
July 2010

Multiproduct Firms and Price-Setting: Theory and Evidence from U.S. Producer Prices*

Saroj Bhattarai and Raphael Schoenle†
Pennsylvania State University and Brandeis University

July 29, 2010

Abstract

In this paper, we establish three new facts about price-setting by multi-product firms and contribute a model that can match our findings. On the empirical side, using micro-data on U.S. producer prices, we first show that firms selling more goods adjust their prices more frequently but on average by smaller amounts. Moreover, the higher the number of goods, the lower is the fraction of positive price changes and the more dispersed the distribution of price changes. Second, we document substantial synchronization of price changes within firms across products and show that synchronization plays a dominant role in explaining pricing dynamics. Third, we find that within-firm synchronization of price changes increases as the number of goods increases. On the theoretical side, we present a state-dependent pricing model where multi-product firms face both aggregate and idiosyncratic shocks. When we allow for firm-specific menu costs and trend inflation, the model matches the empirical findings.

JEL Classification: E30; E31; L11.

Keywords: Multi-product firms; Number of Goods; State-dependent pricing; U.S. Producer prices.

*We thank, without implicating, Jose Azar, Alan Blinder, Thomas Chaney, Gauti Eggertsson, Penny Goldberg, Oleg Itskhoki, Nobu Kiyotaki, Kalina Manova, Marc Melitz, Virgiliu Midrigan, Emi Nakamura, Woong Yong Park, Sam Schulhofer-Wohl, Chris Sims, Jon Steinsson, Mu-Jeung Yang and workshop and conference participants at the Midwest Macro Meetings, Recent Developments in Macroeconomics at Zentrum für Europäische Wirtschaftsforschung (ZEW) and Mannheim University, and Princeton University for helpful suggestions and comments. 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 project coordinators, Ryan Ogden, and especially Kristen Reed, for substantial help and effort, as well as Greg Kelly and Rosi Ulicz for their help. We gratefully acknowledge financial support from the Center for Economic Policy Studies at Princeton University.

†Contact: Saroj Bhattarai, The Pennsylvania State University, 610 Kern Building, University Park, PA 16802-3306. Phone: +1-814-863-3794, email: sub31@psu.edu. Raphael Schoenle, Mail Stop 021, Brandeis University, P.O. Box 9110, 415 South Street, Waltham, MA 02454. Phone: +1-617-680-0114, email: schoenle@brandeis.edu.

1 Introduction

In this paper, we analyze price-setting behavior from a new angle, using the firm as the unit of analysis. We examine the micro-data underlying the U.S. Producer Price Index (PPI) and establish three new empirical facts, showing that the number of goods produced by a firm, synchronization of price changes within firms, and the interaction of these two are key variables for explaining price adjustment decisions. On the theoretical side, we find that in order to match our empirical findings, we need to include firm-specific menu costs and trend inflation into a state-dependent model of multi-product firms.

Our results highlight that heterogeneity among firms, here captured by the number of goods produced by them, plays a critical role in explaining pricing dynamics. Our findings directly suggest that it is necessary to model multi-product firms as distinct from an aggregate of many single-product firms. Moreover, we know from recent work by Midrigan (2008) that allowing for the features which we document empirically has large aggregate implications in terms of monetary policy.[1] Finally, given that approximately 98.55% of all prices in the PPI are set by firms with more than one good,[2] analyzing how the number of goods in a firm relates to pricing decisions appears to be of independent interest since it contrasts with the standard macro-economic assumption of price-setting by single-product firms.

We analyze the PPI micro data and establish the following: first, pricing behavior is systematically related to the number of goods produced by firms.[3] We find that as the number of goods increases, the frequency of price adjustment increases. At the same time, the average size of price changes, conditional on adjustment, decreases. This result holds for both upwards and downwards price changes. In addition, we find that small price changes

[1] In particular, when there are complementarities in the cost of adjusting the prices of goods, price changes will always be dispersed with some very large and some very small price changes. Therefore, monetary shocks will have large real effects similar to models of time-dependent adjustment.

[2] We define a good as a particular brand of product which is moreover identified according to certain characteristics that do not change over time, such as having the same buyer over time. The data section contains further details on the good definition.

[3] The number of goods is a new dimension for the analysis of price-setting behavior relative to previous studies. We can analyze this dimension because there is substantial variation in the number of goods across the more than 20,000 firms in the PPI. Previous studies closest to our analysis such as Lach and Tsiddon (1996), Lach and Tsiddon (2007) or Midrigan (2008) have focused on price-setting not across but within single large retailers such as Dominick's Supermarket.

are highly prevalent in the data and become more prevalent when the number of goods increases. Finally, there is substantial dispersion in the size of price changes and it increases with the number of goods. For example, the distribution of price changes conditional on adjustment is highly leptokurtic, and becomes more leptokurtic as the number of goods per firm increases. Moreover, both the first and the 99^{th} percentiles of price changes take more extreme values.

Second, we find evidence for substantial within-firm synchronization of individual price adjustment decisions: when the price of one good in a firm changes, the probability increases that the price of another good in the firm changes in the same direction. We estimate a multinomial logit model to relate individual, good-level price adjustment decisions to the fraction of price changes of the same sign within a firm or within the same industry, while controlling for economic fundamentals such as inflation. We find that price changes of individual goods are synchronized within firms. This result holds for both upwards and downwards adjustment decisions. Moreover, our results show that such synchronization within the firm is stronger than within the industry. In fact, the fraction of price changes of the same sign within a firm is economically more important for individual adjustment decisions than the fraction of price changes within the same industry or economic fundamentals such as inflation. The within-firm synchronization is also stronger for positive than for negative price adjustment.

Third, we document that the number of goods and the degree of within-firm synchronization strongly interact in determining individual price adjustment decisions. When we group firms by the number of goods and estimate the same multinomial logit model for each category, we find that the strength of within-firm synchronization increases monotonically as we move to groups with more goods. Again, this result holds both for upwards and downwards adjustment decisions. At the same time, we find that the strength of synchronization within the same industry decreases monotonically as the number of goods increases.

These findings thus contribute to the literature a first account of price-setting dynamics from the perspective of the firm, using the micro data underlying the U.S. PPI. Moreover, our findings are significantly wider and more general compared to typical studies that use retail or grocery store data. Importantly, given the variation in the number of goods produced by firms in our dataset, an important source of firm heterogeneity, we are able to systematically uncover patterns across firms as we vary this dimension.

Next, on the theoretical side, we develop a state-dependent pricing model which is consistent with these empirical findings. In particular, we show how these trends are critical in validating different features of the model. In our model, the firm is subject to both idiosyncratic productivity shocks and aggregate inflation shock. In addition, there is a menu cost of changing prices which is firm-specific. We understand this menu cost very broadly as a cost of price adjustment as Blinder et al. (1998) or Zbaracki et al. (2004) argue. In our setup, firm-specific menu costs imply that when a firm produces more than one good, the menu cost is shared among the goods. Moreover, we assume that there are economies of scope in the cost of adjusting prices.

To build intuition for the predictions of the model under these assumptions, consider a 1-good and a 2-good firm. Given economies of scope, the per-good menu cost is lower for a 2-good firm compared to the menu cost of a 1-good firm. Therefore, this implies that a 2-good firm essentially gets to change the price of a second good for free when it decides to change a particular price. This leads, on average, to a higher frequency of price changes. At the same time, since a lot of price changes occur even when a desired price is not very different from the current price of the item, the mean absolute, positive, and negative size of price changes for the 2-good firm is lower. This also implies that the fraction of small price changes is higher for the 2-good firm. This mechanism of the model is reinforced when going from two to three goods.

How can the model explain the increase in the fraction of negative price changes as the number of goods increases? With trend inflation, firms adjust downwards only when they receive substantial negative productivity shocks. With firm-specific menu costs, since the firm adjusts both prices when the desired price of one item is very far from its current price, a higher fraction of downward price changes becomes sustainable. The model also predicts that kurtosis increases as the number of goods goes up because of a higher fraction of price changes in the middle of the distribution. Moreover, both positive and negative adjustment decisions become more synchronized within the firm. At the same time, as we find in the data, positive adjustment decisions are more synchronized than negative adjustment decisions. This is due to common positive shocks from upward trend inflation.

Our model is related to work by Sheshinski and Weiss (1992) and Midrigan (2008). We make two theoretical contributions to the literature as we solve our state-dependent model of price-setting. First, we vary the number of goods produced by a firm from 1 to 3 goods. This allows us to investigate what features of such a state-dependent model are necessary to match the empirical trends we observe as the number of goods per firm increases. Second, we check whether theoretical synchronization results when comparing the pricing behavior of a 2-good and a 3-good firm are consistent with our empirical findings. Such an analysis cannot be done by comparing a 1-good with a 2-good firm since there is no synchronization metric for a single product firm. We find that our model qualitatively matches the stylized facts on pricing behavior and synchronization of adjustment decisions which we establish. Most importantly, we only need to adjust the menu cost parameter in order to match these multiple moments, allowing for economies of scope in the cost of price adjustment as the number of goods produced by a firm increases.

Our paper is also related to two strands of the literature on multi-product firms. Our focus on the effect of variations in the number of goods connects our analysis to the recent literature in industrial organization and international trade on multiproduct firms. This literature examines how firms adjust both their product mix and the number of goods produced in response to changes in the economic environment. For example, Bernard et al. (2009, 2010) analyze the key role of such adjustment for patterns of international trade, aggregate productivity dynamics and product turnover for U.S. firms while Goldberg et al. (2008) do so for Indian firms. These papers do not consider price-setting behavior of firms, however, which is the focus of our paper.

Our analysis is also closely related to work in monetary economics that explores both theoretically and empirically price-setting behavior of multi-product firms. In a seminal paper, Sheshinski and Weiss (1992) show the conditions under which price-setting by multi-product firms is likely to be synchronized or staggered. Lach and Tsiddon (1996) use retail store data and Fisher and Konieczny (2000) use data from a newspaper chain to show that price changes are synchronized within an firm while staggered across firms. Moreover, Lach and Tsiddon (2007) show how small price changes are prevalent in retail store data and argue that this feature can be consistent with a model of multi-product firms where part of the cost of price adjustment is firm-specific, an idea we connect to in our modeling part. Finally, Midrigan (2008) extends the insights in these papers substantially. On the theoretical side, he presents a general equilibrium model where two-product firms face economies of scope in the technology of adjusting prices. The striking quantitative result of his paper is that aggregate fluctuations from monetary shocks are substantially larger than in traditional state-dependent models and almost as large as time-dependent models. On the empirical side, he uses grocery store data to show that a large fraction of price changes are small in absolute values and that the distribution of price changes, conditional on adjustment, is leptokurtic. Both his theoretical and empirical results however, do not contain an analysis of price setting behavior as the number of goods produced varies.

2 Empirics

2.1 Data

We use monthly producer price micro-data from the Bureau of Labor Statistics (BLS). We draw producer prices from the dataset that is normally used to compute the Producer Price Index (PPI). Using producer prices makes our results comparable to other studies of price-setting behavior as well as consistent with a model where firms, and not retailers, set prices. In the discrete choice regressions, we supplement the PPI data with the monthly inflation rate which we gather data from the OECD "Main Economic Indicators (MEI)." We describe the PPI dataset, the OECD data, and our method of categorizing firms by the number of goods produced in the following sections.

2.1.1 Producer Price Data

The PPI contains a large number of monthly price quotes for individual "items", that is, particular brands of products with certain time-persistent characteristics. These items which we henceforth refer to as goods are selected to represent the entire set of goods produced in the US and are sampled according to a multi-stage design.[4] This sampling procedure takes three main steps: in a first step, the BLS compiles a sampling universe of all firms producing in the US using lists from the Unemployment Insurance System. Most firms are required to participate in this system and the BLS verifies and completes the sampling frame using additional publicly available lists, for example in the service sector. In a second step, "price-forming units" which are usually defined to be "production entities in a single location" are selected for the sample according to the total value of shipment of these units or according to their total employment. In a final series of steps called "disaggregation," a BLS agent conducts a field visit and selects the actual goods to be selected into the sample. Again, total values of shipment are used for selection.

In this last step, the BLS takes great care to obtain actual transaction prices. This emphasis on transaction prices goes back to a critique by Stigler and Kindahl (1970) when the data was based on list and not transaction prices. In addition, the BLS also uniquely identifies a good according to its "price-determining" characteristics such as 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. Moreover, the BLS collects information on price discounts and special surcharges. Once a good has been sampled and uniquely identified according to its price-determining characteristics, the BLS collects monthly prices for that very same good and the same customer through a re-pricing form. Moreover, neither order prices nor futures prices are included in the dataset.[5]

Despite this emphasis on transaction prices, there might be some concern about the quality of the price data: respondents have the option to report on the re-pricing form that a price has not changed. This might induce a bias in the price data towards higher price stickiness if respondents are lazy. Using the episode of the 2001 anthrax scare when the BLS exclusively collected prices by phone, Nakamura and Steinsson (2008)[6] show however, that the frequency of price changes, controlling for inflation and seasonality, was the same

[4]For a detailed description of the sampling procedures, see Chapter 14 of the BLS Handbook of Methods (US Department of Labor, 2008).

[5]The PPI price then is defined as "the net revenue accruing to a specified producing establishment from a specified kind of buyer for a specified product shipped, or service provided, under specified transaction terms on a specified day of the month" (BLS, 2008).

[6]See footnote 12 in Nakamura and Steinsson (2008). This idea was first used in Gopinath and Rigobon (2008) where it is applied to export and import prices.

in months when data were collected using the standard mail form as when the collection was done through personal phone calls.

Moreover, since the same product is priced every month, the BLS accounts for instances of product change and quality adjustments. When there is a physical change in a product, one of several quality adjustment methods are used. These include the direct adjustment method for minor physical specification changes, and either the explicit quality adjustment method or the overlap method for major changes. Hedonic regressions have also now been introduced by the BLS into these adjustment processes.

Using these data, we choose our sample for analysis after undertaking the following steps. First, we focus on prices for market transactions, eliminating all intra-firm trade prices. This eliminates 2.83% of the data. Second, we choose the years from 1998 through 2005 as a time-frame. This has the advantage that consistent sampling methods were applied during that time period. Third, we drop all time series where the buyer type is classified as "foreign buyer" during the entire time series. This removes 1.06% of the remaining prices. Fourth, we follow Neiman (2010) and Gopinath and Rigobon (2008) by dropping time series with fewer than six data points. This affects only 0.37% of the data in the previous step. Fifth, we drop prices where the associated price changes are larger in magnitude than two log points. This applies to 0.006% of the data from the previous step. Finally, we concord four-digit SIC industry coding of the data to six-digit NAICS coding.

2.1.2 Macro-Economic Data

We supplement the BLS PPI data with data on CPI inflation. We obtain the series at monthly frequency from the OECD "Main Economic Indicators (MEI)." We use both CPI inflation including food and energy prices as well as excluding food and energy prices. Since we find no qualitative difference in our results, we only report results from the inclusive CPI measure in the paper.

2.1.3 Identifying and Grouping Firms

The PPI data allow us to identify firms according to the number of goods produced by them. This distinction uses the firm identifiers and then counts the number of goods in the data for each firm and at any point in time. We define firms at the establishment level (for example, "Company XYZ").[7] We then group the firms into the following four good bins:

a) Bin 1: firms with 1 to 3 goods, on average.

b) Bin 2: firms with 3 to 5 goods, on average.

c) Bin 3: firms with 5 to 7 goods, on average.

d) Bin 4: firms with 7 goods, on average.

Thus, firms in higher bins sell a greater number of goods than firms in lower bins.

Importantly, this way of grouping the sampled data ensures, beyond the sampling scheme, that the sampled data monotonically map the number of actual goods per firm.[8] On the one hand, the BLS sampling design in the "disaggregation" stage is such that all the economically important products tend to be sampled, with probability proportional to their sales.[9] In addition, the BLS pays special attention to cover all distinct product categories if they exist in a firm and allows some discretion in sampling when there are many products in a firm. Thus, if a firm has more products, more products will be sampled on average. Importantly, on the other hand, our strategy of binning goods into the ranges given above leaves some room for potential errors in the sampling and allows us to average out these potential errors when we calculate our statistics of interest. Finally, as the results show, our choice of binning leads to results that our model in all cases predicts would be indeed identified with an increasing number of goods per firm.

We present in Table 1 some descriptive statistics on firms according to the groups that we construct. The average number of goods per firm across these bins is 2.2, 4.0, 6.1, and 10.3 respectively. The table shows that while the majority of firms, around 80%, fall in bins 1 and 2, there are a substantial number of firms in bins 3 and 4 as well. In fact, since firms in bins 3 and 4 produce more goods, they account for a much larger share of prices than of firms. Firms in bins 3 and 4 set around 40% of all prices in our data.

Regarding firm size, the table reports two statistics which we compute as follows. First,

[7] Therefore, we will use the terms "firms" and "establishments" interchangeably in this paper. In the PPI dataset, they correspond to, as we have described above, what are called "price forming units."

[8] We thank Alan Blinder and Chris Sims for pointing us to this potential sampling issue.

[9] We know from Bernard et al. (2010) and Goldberg et al. (2008)'s Table 4 that large firms are multi-product firms with substantial value of sales concentrated in a few goods. We present an analogous table for our dataset in the appendix. Results suggest that sampling is likely to monotonically capture the actual number of economically important goods. Please see the appendix for details.

after placing firms in different bins, we compute mean employment at the firm level, which is defined as employment per average number of goods per firm. Then, we take the median across different industries (NAICS 3 digit). Finally, we report in Table 1 as mean employment, the average of these medians across all industries in a bin, and as median employment, the median of these medians. The table shows that there is no clear trend in terms of employment per good across these different bins.

Tables 2 and 3 present the distribution of firms across bins and industries (NAICS 2 digit). Table 2 shows that no particular industry substantially dominates a particular good bin and that in fact, NAICS 31, 32, and 33 (durable and non-durable manufacturing) are the dominant industries for all bins, accounting for around 45-70% of all firms. Table 3 shows that for a particular industry, typically good bins 1 and 2 contain the vast majority of firms. Notable exceptions are NAICS 22 (utilities) which contains a very high proportion of firms that fall in bin 4, and NAICS 62 (health care and social assistance) where almost half the firms are in bins 3 and 4.

2.2 Results

We report the results from our empirical analysis in two parts below. First, we document important aggregate statistics on price changes by firms according to the good bins. Second, we show the role played by economic fundamentals in pricing decisions at the good level.

In relation to recent studies of price adjustment using micro data from the BLS and retail stores, it is worth noting at the outset the following aspects of the PPI data. As documented by Nakamura and Steinsson (2008), while sale prices are important in the CPI data, they are not prevalent in the PPI data. Therefore, we do not distinguish between sale and non-sale prices, for example, by using a sales filter. Nakamura and Steinsson (2008) also show that for aggregate statistics on price changes, accounting for product substitutions can make a difference, especially in the CPI. All the baseline results we report below are excluding product substitutions.[10] The results nevertheless remain the same while including product substitutions. Finally, using grocery store data, Eichenbaum et al. (2009) have documented

[10]We identify product replacement by changes in the so-called "base price" which contains the price at each resampling of a good in the PPI. When this base price changes within a price time series, but the data show no change in the actual price series, we set our product substitution dummy to one.

"reference pricing" and non-permanent price changes. In the PPI, in contrast, these features are not prevalent.

2.2.1 Basic Statistics

We present below aggregate statistics on the frequency, size, direction, and distribution of price changes according to the four good bins that we construct.[11]

Frequency of price changes We compute the frequency as the mean fraction of price changes during the life of a good. We do not count the first observation as a price change and assume that a price has not changed if a value is missing. Also, we do not explicitly take into account issues of left-censoring of price-spells. However, we verify that taking into account left-censoring leaves the resulting distribution of frequencies in the PPI essentially unchanged. For our purpose, it is most relevant that we apply our method consistently across all firms. After computing the frequency of price changes at the good level, we calculate the median frequency for all goods within the firm. Then, we report the mean, median, and standard error of frequencies across firms in a given good bin. We use the standard error to compute 95% confidence intervals through out the paper.

Figures 6 and 6 show that the monthly mean and median frequency of price changes increase with the number of goods produced by firms. The mean frequency increases from 20% in bin 1 to 29% in bin 4 while the median frequency increases from 15% to 23%. The relationship is monotonic across bins except for the mean frequency of price changes for bins 1 and 2. Inverting these frequency values, this implies that the mean duration of a price spell decreases from 5 months in bin 1 to 3.4 months in bin 4 while the median duration decreases from 6.7 months to 4.3 months. Therefore, in general, firms that produce a greater number of goods change prices more frequently.[12]

[11] While we control for a large set of potential explanators of price adjustment in the next section, we have also separately estimated "aggregate" statistics specifically controlling for size only. This is done by regressing firm-level statistics on bin dummies and measures of firm size. "Aggregate" results as discussed in this section remain essentially unchanged and are available on request. This finding is consistent with the insight in Goldberg and Hellerstein (2009) that large firms change prices more frequently and by smaller amounts if one takes into account that firms which sell a higher number of goods have bigger total sales.

[12] Using our dataset to compute an aggregate measure of frequency and duration of price changes in the PPI, we get estimates of 0.21 and 0.16 for the mean and median frequency and 6.91 and 5.74 months for the mean and median duration. This is calculated by first computing the frequency at the good level, second

Direction and size of price changes We define the fraction of positive price changes as the number of strictly positive price changes over all zero and non-zero price changes. We compute this at the firm level and then report the mean across firms in a given good bin.[13] Figure 6 shows that the mean fraction of positive price changes decreases with the number of goods produced by firms as it goes down from 0.64 in bin 1 to 0.61 in bin 4. Firms with many goods therefore adjust prices upwards less frequently.

We next compute the size of price changes as the percentage change to last observed price. Again, we compute this at the good level, take the median across goods in an firm, and then report the mean across firms in a good bin. Figure 6 shows that the mean absolute size of price changes decreases with the number of goods produced by firms as it goes down from 8.5% in bin 1 to 6.6% in bin 4.[14]

Moreover, this relationship holds even when we separate out the price changes into positive and negative price changes. Figure 6 shows that the mean size of positive price changes decreases with the number of goods while the mean size of negative price changes increases with the number of goods. Thus, in general, firms that produce a greater number of goods adjust their prices by a smaller amount, both upwards and downwards.

An interesting statistic in this context is the fraction of small price changes where we define a small price change as:

$$|\Delta p_{i,t}| \leq \kappa \overline{|\Delta p_{i,t}|} \tag{1}$$

where i is a firm and $\kappa = 0.5$. That is, a price change is small if it is less in absolute terms than a specified fraction (here 0.5) of the mean absolute price change in a firm. After computing this at the firm level, we then report the mean in a good bin. Figure 6 shows that the mean fraction of small price changes increases from 0.38 in bin 1 to 0.55 in bin 4.

by taking the median across goods in a classification group, third, by taking the median across classification groups within six-digit categories and fourth, by taking means and medians across six-digit categories.

[13] While we report only the mean across firms in a given bin for all statistics other than frequency, all our results are completely robust to whether we compute the mean or the median across firms.

[14] Using our dataset to compute an aggregate measure of the absolute size of price changes in the PPI, we find estimates of 6.96% and 5.34% for the mean and median size. This is calculated by first computing the mean absolute price change at the good level, second by taking the median across goods in a classification group, third, by taking the median across classification groups within six-digit categories and fourth, by taking means and medians across six-digit categories.

Therefore, small price changes are more prevalent when firms produce many goods and in fact, for bin 4, more than half the price changes are small.

Dispersion of price changes We define the kurtosis of price changes as the ratio of the fourth moment about the mean and the variance squared:

$$K = \frac{\mu_4}{\sigma^4} \tag{2}$$

where

$$\mu_4 = \frac{1}{T-1} \sum_{i}^{n} \sum_{t=1}^{T_i} (\Delta p_{i,t} - \overline{\Delta p})^4 \tag{3}$$

and where in a given firm with n goods, $\Delta p_{i,t}$ denotes the price change of good i, $\overline{\Delta p}$ the mean price change and σ^4 is the square of the usual variance estimate.

Figure 6 shows that the mean kurtosis of price changes increases with the number of goods produced by firms as it goes up from 5.3 in bin 1 to 16.8 in bin 4. Thus, even for bin 1, the distribution is leptokurtic. We also document in Figure 6 that both the 1^{st} and the 99^{th} percentiles of price change take more extreme values as the number of goods increases. Thus, the dispersion in price changes increases with the number of goods produced by firms.

2.2.2 Regression Analysis

Here, we go beyond providing aggregate statistics and estimate a discrete choice model to analyze what economic fundamentals determine pricing dynamics at the good level. In particular, we estimate a multinomial logit model for the decision to change prices. This allows us to separately examine the relationship of upwards and downwards adjustment decisions with the explanators.[15]

We impose a multinomial logit link function with 3 categories: $m = 0$ for no price

[15] Instead of the multinomial logit model, we could also estimate an ordered probit model. This would assume that there is a “ranking” of outcomes in terms of how the latent underlying variable cutoffs relate to the right-hand side variables. The latent variable in our case could be interpreted as deviation from the desired optimal price. This might indeed result in some ordering - for example, high inflation means one is likely below the desired optimal price and hence, 1 is the adjustment decision preferred to 0 and -1. For other right-hand side variables, however, this relationship is unclear, for example for the fraction of price changes within the firm or even month dummies. Hence, we estimate conservatively using the multinomial logit model.

change, $m = 1$ for a price increase, and $m = -1$ for a price decrease. The multinomial logit model is described in detail for example, by Agresti (2007). Denoting by $\Pi_{i,m}$ the probability that decision m is taken for good i, the probability under the multinomial logit link is given by:

$$\Pi_{i,m,t} = \frac{e^{X_{i,t}\alpha_m}}{\sum_m e^{X_{i,t}\alpha_m}}. \tag{4}$$

Since $\sum_m \Pi_{m,t} = 1$, the three sets of parameters are not unique. Therefore, we follow standard practice and choose category $m = 0$ as a baseline category:

$$\Pi_{i,0,t} = \frac{1}{1 + \sum_{m \in (-1,1)} e^{X_{i,t}\alpha_m}}. \tag{5}$$

We estimate the two remaining logit equations simultaneously. The logit model has the convenient property that the estimated coefficients take on the natural interpretation of the effect of the explanators on the probability of adjusting prices up or down over taking no action.

We include as controls in X the fraction of price changes at the same firm and the same six-digit NAICS sector, excluding the price change of the good we are trying to explain. These variables are meant to capture the extent of synchronization in price setting at the firm and the sectoral level. These variables are in the spirit of the probit analysis of Midrigan (2008). Moreover, to control for a measure of marginal costs, we also include in X the average price change of goods in the same firm and six-digit NAICS sector. We also include a dummy for product replacement where we can identify it: so-called "base prices" in the PPI contain the first price at each resampling. When this base price changes within a price time series but the data show no change in the actual price series, we set the product replacement dummy to one. As an important fundamental factor, we include energy and food CPIs in X. Finally, we control for the total number of employees in the firm, industry fixed effects, month fixed effects, and time trends in the data.

We are not aware of any other similar broad-based analysis of U.S. producer micro prices. Since the results for the PPI as a whole are likely to be of independent interest, we first estimate the model on pooled data across all good bins. Then, we focus on estimation separately by good bins.[16]

Pooled Data Table 4 shows the results from this multinomial logit model for the pooled data across all good bins. We report what are called the relative risk ratios, equivalently the odds ratios, for the different independent variables. Therefore, a coefficient value greater than 1 indicates that a change in the independent variable increases the odds of the dependent category compared with the base category. We also report the marginal effect at the mean and the effect when the dependent variable changes from mean $-$ 1/2 standard deviation to mean + 1/2 standard deviation. Our main findings are as follows:

Synchronization of price changes We find robust evidence for synchronization of price setting both within the industry and the firm.

First, we find that adjustment decisions are synchronized within the industry: The probability of adjusting the price of a good in a firm is higher when the fraction of price changes of the same sign in the industry, excluding that good, increases. This holds for both negative and positive price changes. The effects are both statistically and economically significant. When evaluated at the mean, a 1% increase in the fraction of negative price changes of other goods in the industry leads to a 0.06% increase in probability of a negative price change of a good. Similarly, a 1% increase in the fraction of positive price changes of other goods in the industry leads to a 0.1% increase in probability of a positive price change of a good. The economic significance can also be discerned from the effects when at the mean the fraction of price changes of the same sign changes by one standard deviation: for negative price changes, the probability of a downward price change of a good increases by 1.32%, while for positive price changes, the probability of an upward price change of a good increases by 2.27%.

Second, the results also show that there is substantial synchronization of adjustment decisions within the firm. When the fraction of price changes of the same sign of other goods within the firm increases, then the likelihood of a price change of a given good increases.

[16] As a robustness test, we also control for statistical importance using weights, non-linear effects of size and potential clustering at the six-digit industry level in both our aggregate and bin-by-bin specifications. Since this leaves results essentially unchanged, we skip their discussion in the following. Results are available upon request.

Again, this holds for both negative and positive price changes. When evaluated at the mean, a 1% increase in the fraction of negative price changes of other goods in the industry leads to a 0.3% increase in probability of a negative price change of a good. Similarly, a 1% increase in the fraction of positive price changes of other goods in the industry leads to a 0.5% increase in probability of a positive price change of a good. These effects are therefore not only statistically, but also highly economically, significant. The economic significance can also be discerned from the effects when at the mean, the fraction of price changes of the same sign changes by one standard deviation: for negative price changes, the probability of a downward price change of a good increases by 8.81%, while for positive price changes, the probability of an upward price change of a good increases by 14.7%. Finally, the coefficient on the fraction of same-signed price changes is larger for positive adjustment decisions than for negative adjustment decisions as results in Table 4 show. In fact, the marginal effect is about twice as large for positive price adjustment decisions. Table 5 summarizes this finding.

State-dependent response to inflation The results also show evidence for the fundamental role played by inflation, an important aggregate shock, in pricing decisions. In particular, the likelihood of a price decrease decreases with higher CPI inflation while the likelihood of a price increase increases. This is as one would expect from a model where firms adjust prices in a state-dependent fashion. The effects are both statistically and economically significant. When evaluated at the mean, a 1% increase in CPI inflation decreases the probability of a negative price change of a good by 0.6%, while increasing the probability of a positive price change by 0.5%. Similarly, when at the mean, CPI inflation changes by one standard deviation, the probability of a negative price change of a good decreases by 0.21%, while the probability of a positive price change increases by 0.17%.

Seasonality in price changes We uncover strong seasonality in price changes in the data. In particular, there are strong month effects for both negative and positive price changes. These effects we report are relative to December, which is the excluded month. The results show that the odds of a negative price changes are extremely high in March and relatively high in August, September, October, and November, while extremely low in February and July. The odds of a positive price changes are extremely high in February and July and relatively high in January, April, September, October, and November, while extremely low in March. The results thus show strong seasonal effects in the first quarter and also imply that price changes, both positive and negative, are likely in September, October, and November. There is also evidence that price changes, both positive and negative, are less likely in May and June. Our results on seasonality in price changes are consistent with similar findings in Nakamura and Steinsson (2008), in particular, the strong effects in the first quarter.

Critical role of within-firm variables The inclusion of within-firm variables, which is unique to our analysis, plays a critical role in explaining price setting behavior. Notice from above that the within-firm effects are much stronger than the within industry effects. To isolate the key role of the within-firm variables more clearly, we run the same regressions as above but exclude from X the fraction of price changes in the same firm and the average price change of goods in the same firm. Table 6 shows that while the evidence on synchronization of price changes within the industry, the fundamental role of inflation, and strong month effects remain intact, the explanatory power of the model drops significantly. The R^2 goes down from 48% to 30%.[17] Finally, we also find that among within-firm variables, product substitution is highly significantly associated with a price change. This again holds for both upwards and downwards adjustments.

By Bins Next, we run the mulitnominal logit regression for the four good bins separately to investigate differences due to the number of goods produced by firms. Tables 7, 8, and 9 show the results for the relative risk ratios, the marginal effect at the mean, and the effect when the dependent variable changes from mean $-$ 1/2 standard deviation to mean + 1/2 standard deviation. To avoid cluttering, we do not report the p-values, but all the results we report below are statistically significant. Our main results, as we move from bin 1 to bin 4, are as follows:

[17] The R^2 measure we report denotes the usual pseudo-R^2. This statistic is based on the likelihood and measures improvements of the model fit.

Decreasing within-industry synchronization It is clear that the coefficient on the fraction of price changes in the industry decreases as the number of goods produced by firms increases. When evaluated at the mean, the effect of a 1% increase in the fraction of negative price changes of other goods in the industry on the probability of a negative price change of the good goes down from 0.11% in bin 1 to -0.06% in bin 4. Similarly, the effect of a 1% increase in the fraction of positive price changes of other goods in the industry on the probability of a positive price change of the good goes down from 0.16% in bin 1 to 0% in bin 4. Finally, the effect when at the mean, the fraction of price changes of the same sign changes by one standard deviation goes down from 2.1% to -1.51% for negative price changes and from 3.1% to 0.2% for positive price changes.

Increasing within-firm synchronization The tables also show that the coefficients on the fraction of price changes on goods within the same firm increases as the number of goods produced by firms increases. When evaluated at the mean, the effect of a 1% increase in the fraction of negative price changes of other goods in the firm on the probability of a negative price change of the good goes up substantially from 0.25% in bin 1 to 0.51% in bin 4. Similarly, the effect of a 1% increase in the fraction of positive price changes of other goods in the firm on the probability of a positive price change of the good goes up from 0.44% in bin 1 to 0.78% in bin 4. Finally, the effect when at the mean, the fraction of price changes of the same sign changes by one standard deviation, increases from 5.38% to 15.8% for negative price changes and from 9.61% to 24.24% for positive price changes. Importantly, the marginal effects of a change in the fraction of price adjustment within the firm on individual adjustment decisions is systematically larger across all bins for positive adjustment decisions compared to negative adjustment decisions. Thus, we find that there is evidence for greater synchronization of price changes within the firm as the number of goods produced by the firm increases. Moreover, synchronization is always stronger for positive than for negative price changes.

Lower probability of downwards changes with inflation Finally, while the relationship between number of goods and the coefficient on inflation for upward changes do not show a clear trend, there is a clear decrease in odds of a downward change in prices in response to inflation when the number of goods increases. When evaluated at the mean, the effect of a 1% increase in CPI inflation on the probability of a negative price change decreases from -0.24% in bin 1 to -1.55% in bin 4.

2.2.3 Implications for Theory

Our empirical results have potentially important implications for the way we model price setting by firms. This includes the way we think of the role of nominal rigidities in monetary models since we have shown that the number of goods plays a key role in explaining why firms change price. In the international context, the results should equally have consequences for modeling pass through and real exchange rate and trade dynamics. In fact in ongoing work, Bhattarai and Schoenle (2009), we find that in the micro-data underlying the U.S. import and export price indices, frequency of price adjustment, and exchange rate pass-through are systematically related to the number of goods. In the following, we outline some implications for the way we model price-setting and how they relate to other recent empirical findings.

First, in terms of aggregate stylized facts, our results complement several arguments that have been made in the literature. The significant fraction of negative price changes that we document imply, as also argued by Golosov and Lucas Jr. (2007) and Nakamura and Steinsson (2008), that models that rely on only aggregate shocks, and hence predict predominantly positive price changes with modest inflation, are inconsistent with micro data. We have also shown that the absolute size of price changes are large, which again suggests the need for idiosyncratic firm level shocks, as emphasized by Golosov and Lucas Jr. (2007) and Klenow and Kryvtsov (2008). At the same time however, there is a substantial fraction of small price changes in the data and the distribution of price changes is highly leptokurtic. This observation, as argued by Klenow and Kryvtsov (2008) and Midrigan (2008), implies that simple menu cost models are inconsistent with micro data since they do not predict enough small price changes.

Second, using a discrete choice model for changes of producer prices, we have uncovered additional relationships that have implications for theory. We have shown broad based evidence for synchronization of producer price changes within industries and firms. This suggests a model with industry and firm level strategic complementaries. In particular, firm level synchronization is quantitatively important, as it explains a significant part of the variation in the probability of price changes. We also show that inflation plays a fundamental role in pricing decisions by increasing the likelihood of a price increase while decreasing the likelihood of a price decrease. This suggests a model where pricing decisions are state-dependent.

Third, the central empirical focus of this paper, the analysis of price dynamics according to the number of goods produced by firms, has additional implications for theory. We have shown that price setting dynamics change substantially with the number of goods produced by firms. The increasing frequency, decreasing size, and a higher fraction of small price changes with an increasing number of goods suggest a model where part of the menu cost of adjusting prices is shared among goods produced by a firm.

Compared to the literature, our main challenge here is to explain the trends we observe in price setting as we vary the number of goods produced by firms, an analysis that has not been undertaken before. We turn to this task next.

3 Theory

3.1 Model

We use a partial equilibrium setting of a firm that decides each period whether to update the prices of its n goods indexed by $i \in (1, 2, 3)$, and what prices to charge if it updates. The firm's decision problem arises for three reasons: First, the firm is subject to an aggregate inflationary shock. Second, each good is subject to an idiosyncratic productivity shock. Third, changing the price costs a total fixed "menu cost" each time when the firm decides to adjust prices. The menu cost is thus firm-specific and not good-specific,[18] and firms

[18]This assumption is motivated by the observed synchronization patterns in the data. Conditional on observing at least one adjustment, the median monthly fraction of price changes in a firm is 75% as Table 13 in the appendix shows. Thus, assuming full joint adjustment should be a good first approximation. Moreover, it is important to note that the high fraction of joint price adjustment decisions is not "mechanical" due to the way prices are reported: one could presume that each firm reports prices on only one reply form and whenever there is one price change, the reporter indicates price changes for all goods. However, prices are reported on separate reply forms and frequently by separate agents, especially as the number of goods

trade off this cost with changes in the expected stream of profit.

Our model is similar to the ones in Sheshinski and Weiss (1992) and Midrigan (2008). The main difference is that compared to Sheshinski and Weiss (1992) we allow for a stochastic aggregate shock while compared to Midrigan (2008) we solve for equilibrium as we vary n, the number of goods produced by firms. In particular, the latter variation allows us to make two contributions. First, we can solve for trends in price-setting behavior with respect to how many goods firms produce. Second, we can compute trends in synchronization of price-setting by comparing 2-good and 3-good firms. Since no measures of synchronization can be computed in the one-good case, the comparison of 2-good and 3-good firms is necessary to model trends in synchronization of adjustment decisions.

In our model, the firm produces output of good i using a technology that is linear in labor:

$$c_{i,t} = A_{i,t}\ l_{i,t}$$

where $A_{i,t}$ is a good-specific for good i productivity shock that follows an exogenous process:

$$\ln A_{i,t} = \rho_A^i \ln A_{i,t-1} + \epsilon_{A,t}^i$$

where $\mathbf{E}[\epsilon_{A,t}^i] = 0$ and $var(A_{i,t}) = (\sigma_A^i)^2$. We assume that there is no correlation between good-specific shocks. We do this in order to isolate the effect of multi-product firms on the decision to synchronize price adjustment when the underlying idiosyncratic shocks are uncorrelated, while controlling for the common inflationary shock.

The firm's product i is subject to the following demand:

$$c_{i,t} = \left(\frac{p_{i,t}}{P_t}\right)^{-\theta} C_t \qquad i = 1, 2, ..., n$$

where C_t is aggregate consumption, P_t is the aggregate price level, $p_{i,t}$ is the price of good i, and θ is the elasticity of substitution across goods.

In this partial equilibrium setting, we normalize $C_t = \bar{C}$. We also assume that the price level P_t exogenously follows a random walk with a drift:

increases for which a firm reports prices.

$$\ln P_t = \mu_P + \ln P_{t-1} + \epsilon_{P,t}$$

where $\mathbf{E}[\epsilon_{P,t}] = 0$ and $var(\epsilon_{P,t}) = (\sigma_P)^2$. Given our assumption about technology, the real marginal cost of the firm for good i, $MC_{i,t}$, is therefore given by:

$$MC_{i,t} = \frac{W_t}{A_{i,t} P_t}$$

where W_t is the nominal wage. We normalize $\frac{W_t}{P_t} = \bar{w}$.

Whenever the firm adjusts one or more than one of its prices, it has to pay a constant total "menu cost", $K(n)$, whenever it chooses to adjust prices. We understand this "menu cost" very broadly as a general cost of price adjustment, not the literal cost of relabeling the price tags of goods. Blinder et al. (1998) and Zbaracki et al. (2004) provide some evidence for such a broader interpretation of "menu costs." Moreover, this cost in our model is not a good-specific cost, but firm-specific:

$$K(n) > 0. \tag{6}$$

The cost of changing prices may depend on the number of goods produced by the firm and in particular, we assume that:

$$\frac{\partial K(n)}{\partial n} > 0 \tag{7}$$

and that:

$$\frac{K(n+1)}{K(n)} < \frac{n+1}{n}. \tag{8}$$

These assumptions mean that the cost of changing prices increases monotonically with the number of goods produced, and that there are increasing cost savings as more and more goods are subject to price adjustment. Therefore, there are economies of scope in the cost of changing prices. While we do not model this adjustment process in detail, one can for example think about the adjustment technology as a fixed cost of hiring a manager to change prices: it is costly to hire him in the first place, but much less costly to have him adjust the price of each additional good.

Given this setup, the firm maximizes the expected discounted sum of profits from selling all of its goods. Total period gross profits, before paying the menu costs, are given by:

$$\pi_t = \sum_i^n \left(\frac{p_{i,t}}{P_t} - \frac{\overline{w}}{A_{i,t}} \right) \left(\frac{p_{i,t}}{P_t} \right)^{-\theta} \bar{C}.$$

The problem of the firm is to choose whether to update all prices in a given period, and if so, by how much. Whenever it updates prices, it has to pay the menu cost K.

The problem of the firm can most easily be described and solved recursively: the state variables in this problem are last period's real prices, $\frac{p_{i,t-1}}{P_t}$, and the current productivity shocks, that is, $\mathbf{p}_{-1} = \left(\frac{p_{1,t-1}}{P_t}, \frac{p_{2,t-1}}{P_t}, ..., \frac{p_{n,t-1}}{P_t} \right)$ and $\mathbf{A} = (A_{1,t}, A_{2,t}, ..., A_{n,t})$. Let $V^a(\mathbf{A})$ denote the firm's value of adjusting all prices:

$$V^a(\mathbf{A}) = \max_{\mathbf{p}} \left[\pi(\mathbf{p}; \mathbf{A}) - K + \beta \int \int V \left(\mathbf{p}_{-1}^{'}, \mathbf{A}^{'} \right) dF \left(\epsilon_A^1, \epsilon_A^2, ..., \epsilon_A^n \right) dF(\epsilon_P) \right]$$

where we have omitted time subscripts and $'$ denotes the subsequent period.

Let $V^n(\mathbf{p}_{-1}, \mathbf{A})$ denote the firm's value of not adjusting prices:

$$V^n(\mathbf{p}_{-1}, \mathbf{A}) = \pi(\mathbf{p}_{-1}; \mathbf{A}) + \beta \int \int V \left(\mathbf{p}_{-1}^{'}, \mathbf{A}^{'} \right) dF \left(\epsilon_A^1, \epsilon_A^2, ..., \epsilon_A^n \right) dF(\epsilon_P)$$

where

$$V = \max(V^a, V^n).$$

3.2 Computation

We solve this problem for a firm that produces $n = 1$, $n = 2$ and $n = 3$ goods. First, we employ collocation methods to find the policy functions of the firm. Second, given the policy functions, we simulate time series of shocks and corresponding adjustment decisions for many periods. Finally, we compute statistics of interest for each simulation and good, and across simulations. We also estimate a multi-nomial logit model of adjustment decisions using the simulated data to compare our theoretical results with the empirical findings. The appendix provides further details about computation and analysis.

We present our choice of parameters in Table 10. Since our model is monthly, we choose a discount rate β of $(0.96)^{\frac{1}{12}}$. We use a value of K such that menu costs are a 0.35% of steady-state revenues for the 1-good firm, 0.65% for a 2-good firm, and 0.75% for a 3-good firm. We choose θ to be 4, which implies a markup of 33%.

To parametrize the exogenous processes, we set the trend in aggregate inflation to be a monthly increase of 0.21%. We use persistent idiosyncratic productivity shocks, where the $AR(1)$ parameter is 0.96. We choose 0.37% and 2% respectively for the standard deviations of the aggregate inflation and the idiosyncratic productivity shock. For the firms with 2 and 3 goods, we use the same values for the persistence and variance of the all idiosyncratic productivity shocks.[19] All of our parameters are standard in the literature.

3.3 Results

We present the main results from the simulations in Table 11 and illustrate them graphically in Figures 9-15. As we increase the number of goods from 1 to 3, the model predicts clear and systematic trends in the key price-setting statistics which align with our empirical findings.

First, we find that the frequency of price changes goes up from 15.22% to 19.72% while the mean absolute size of price changes goes down from 5.21% to 3.96% as we increase the number of goods. Table 11 summarizes these trends. Second, the decrease in the absolute size of price changes also holds for both positive and negative prices changes: they go down from 5.34% to 4.23% and from -5.02% to -3.57% respectively. Third, the fraction of small price changes increases from 1.33%, barely none, to 23.59%. Thus, while firms with more goods change prices more frequently, they does so by smaller amounts on average. Fourth, we also see that the fraction of positive price changes decreases from 61.68% to 59.38%. Thus, as in the data, firms with more goods adjust downwards more frequently. Finally, the model predicts that kurtosis increases from 1.38 to 1.97, again consistent with

[19] As emphasized before, we assume that there is no correlation among the idiosyncratic productivity shocks within the firm.

our empirical findings.

What is the mechanism behind our results? For simplicity, compare a 1-good firm with a 2-good firm. For the case of a 2-good firm, when the firm decides to pay the firm-specific menu cost to adjust one of the prices, it also changes the price of the other good because it gets to change it basically for free. This leads, on average, to a higher frequency of price changes. At the same time, for the 2-good firm, since a lot of price changes happen even when the desired price is not very different from the current price of the good, the mean absolute size of price changes is lower. This smaller mean also implies that the fraction of "small" price changes is much higher for the 2-good firm. In fact, for the 1-good firm, which is the standard menu cost model, the fraction of small price changes is negligible because in that case, the firm adjusts prices only when the desired price is very different from the current price.

What causes the decrease in the fraction of positive price changes? With trend inflation, firms adjust downwards only when they receive very big negative productivity shocks. With firm-specific menu costs, since the firm adjusts both prices when the desired price of one good is very far from its current price, it is now more sustainable to have a higher fraction of downward price changes. Finally, kurtosis increases as we go from one good to two goods because of a higher fraction of price changes in the middle of the distribution.

Next, we address trend in synchronization of individual good-level price changes. Using simulated data, we run the same multinominal logit regression as in the empirical section to investigate if price changes become more synchronized as the number of goods produced by firms increases. We thus estimate the following equation, with no price changes as the base category:

$$\{\Delta p_t \neq 0\}_{-1,0,+1} = \beta_0 + \beta_1 f_t + \beta_2 \Pi_t + \epsilon_t \tag{9}$$

where f_t is the fraction of same-signed adjustment decisions at time t within the firm and Π_t is the inflation rate at time t.

It is important to emphasize here the need to go beyond a 2-good case and consider a 3-good case, because we otherwise cannot check if the model predicts trends in synchronization of price changes that are consistent with the empirical findings. Table 11 shows that the strength of synchronization, that is the coefficient estimate in the multinominal logit regressions for the fraction of other goods of the firm changing in the same direction, increases as we go from a 2-good firm to a 3-good firm. Importantly, as is the case empirically, this is the case with both upwards and downwards price changes.

In addition, Table 11 also makes clear that the simulations predict that positive price adjustment decisions are more synchronized than negative adjustment decisions. That is, the synchronization coefficient for upwards price adjustment decisions is always higher than the coefficient for downwards adjustment decisions. This difference in synchronization probabilities is due to positive trend inflation. Without positive trend inflation, the difference disappears. The model thus matches our findings on synchronization established in the empirical section.

3.4 Robustness

In this section, we discuss several extensions of our baseline model and investigate alternative ways to generate the trends which we find in the data. Appendix 3 contains the tables summarizing results from these simulations. First, we investigate the possibility that perhaps our empirical results are driven by the fact that firms that produce more goods also produce more substitutable goods. We simulate a specification of 2 and 3 good firms with no economies of scope in the cost of adjusting prices but with an elasticity of substitution among goods that is higher than that compared to the 1 good firm. As Table 14 shows, while matching the trends in frequency, size, and fraction of positive price changes, this specification fails to match trends in the fraction of small price changes, kurtosis, and synchronization as we increase the number of goods.

Second, we allow for correlation among the productivity shocks. In our main setup, we had not allowed for correlation of productivity shocks at the firm-level in order to isolate the predictions of the model arising solely due to firm-specific menu costs. Table 15 shows that allowing for such correlation does yield a higher synchronization of price changes within the firm, but at the same time, fails to match trends in frequency and size of price changes.

Third, there might be concern that our synchronization results could be due to purely mechanical, statistical reasons. To investigate this possibility, we perform two tests. In the first, we run a Monte-Carlo exercise based on a simple statistical model of price changes. We model price changes as i.i.d. Bernoulli trials with a fixed probability of success. Using simulated time series for an arbitrary number of goods, we estimate our synchronization equation. We find no mechanical trend in synchronization as the number of goods increases.

In the second test, we use the same parametrized model of a single-good firm from our simulation exercise and model a multi-product firm as a collection of multiple, independent single-product firms. That is, a 2-good firm now is simply a collection of two single-good firms, with no economies of scope in cost of adjusting prices. Similarly, a 3-good firm is a collection of three single-good firms, with no economies of scope in cost of adjusting prices. When we run our synchronization test with simulated data from these specifications, we find that the model cannot produce synchronization results that are consistent with the empirical findings. For example, the synchronization coefficient estimated using simulated data on negative price changes is negative, as opposed to positive in the data. Incidentally, this specification of multiproduct firms as multiple single-good firms fails to match any of our aggregate trends in price-setting as shown in Table 16. These results therefore highlight the need to model a multi-product firm as distinctly different from an aggregate of multiple single-product firms.

Fourth, we allow for a menu of menu costs instead of the firm-specific menu cost structure that we use. For example, in the 2-good case, the firm now has a choice of adjusting 0, 1, or 2 goods. There are different menu costs for adjusting the price of 1 good and 2 goods, but some savings in the cost when adjusting the prices of both goods. Our results are robust to this possible extension. In fact, as shown in Table 17, this specification cannot account for the extent of synchronization in price changes observed in the data.

Fifth, we consider two alternative demand specifications. In the first specification, we allow for non-zero cross-elasticities of demand among the goods produced by the firm, which is precluded in our baseline model with CES demand. As Table 18 shows, while this case generates a higher fraction of small price changes and greater kurtosis, it cannot account for the trend in the size of price changes. In the second specification, we allow for idiosyncratic demand shocks as an alternative to productivity shocks. We present the results in Table 19. We find that demand shocks are unable to generate the fraction of negative price changes that we observe in the data.

4 Conclusion

In this paper, we have established three new facts regarding multi-product price-setting in the U.S. Producer Price Index. First, we show that as the number of goods produced increases, price changes are more frequent, the size price changes is lower, the fraction of positive price changes decreases, and price changes become more dispersed. Second, we find evidence for substantial synchronization of price adjustment decisions within the firm. Third, we find that the number of goods and the degree of synchronization within firms strongly interact in determining price adjustment decisions: as the number of goods increases, synchronization within firms increases.

Motivated by these findings, we present a model with firm-specific menu costs where firms are subject to both idiosyncratic and aggregate shocks. We show that as we change the number of goods produced by the firms, the patterns predicted by the model regarding frequency, size, direction, dispersion, and synchronization of price changes are consistent with the empirical findings.

In future work, it will be fruitful to explore the generality of our empirical results. In fact, in ongoing work, Bhattarai and Schoenle (2009), we find that in the micro-data underlying the U.S. import and export price indices, frequency of price adjustment, and exchange rate pass-through are indeed systematically related to the number of goods produced. We also hope to investigate the implications of our findings for business cycle dynamics, real exchange rate behavior, and monetary policy.

References

Agresti, A. (2007): An Introduction to Categorical Data Analysis, John Wiley and Sons, New York, N.Y.

Bernard, A., S. Redding, and P. Schott (2010): "Multi-Product Firms and Product Switching," American Economic Review, 100, 70–97.

Bernard, A. B., S. Redding, and P. K. Schott (2009): "Multi-Product Firms and Trade Liberalization," CEP Discussion Papers dp0769, Centre for Economic Performance, LSE.

Bhattarai, S. and R. Schoenle (2009): "Frequency of Price Adjustment, Exchange Rate Pass-Through, and Multi-Product Firms," Work in progress, Princeton University.

Blinder, A. S., E. R. D. Canetti, D. E. Lebow, and J. B. Rudd (1998): Asking About Prices: A New Approach to Understanding Price Stickiness, Russell Sage Foundation, New York, N.Y.

Eichenbaum, M., N. Jaimovic, and S. Rebelo (2009): "Reference Prices and Nominal Rigidities," Forthcoming American Economic Review.

Fisher, T. C. G. and J. D. Konieczny (2000): "Synchronization of Price Changes by Multiproduct Firms: Evidence from Canadian Newspaper Prices," Economics Letters, 68, 271–277.

Goldberg, P. and R. Hellerstein (2009): "How Rigid Are Producer Prices?" Federal Reserve Bank of New York Staff Reports.

Goldberg, P., A. Khandelwal, N. Pavcnik, and P. Topalova (2008): "Multi-product Firms and Product Turnover in the Developing World: Evidence from India," Review of Economics and Statistics forthcoming.

Golosov, M. and R. E. Lucas Jr. (2007): "Menu Costs and Phillips Curves," Journal of Political Economy, 115, 171–199.

Gopinath, G. and R. Rigobon (2008): "Sticky Borders," Quarterly Journal of Economics, 123, 531–575.

Klenow, P. J. and O. Kryvtsov (2008): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" The Quarterly Journal of Economics, 123, 863–904.

Lach, S. and D. Tsiddon (1996): "Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms," American Economic Review, 86, 1175–96.

——— (2007): "Small Price Changes and Menu Costs," Managerial and Decision Economics, 28, 649–656.

Midrigan, V. (2008): "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," Working paper.

Miranda, M. J. and P. L. Fackler (2002): Applied Computational Economics and Finance, MIT Press.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," Quarterly Journal of Economics, 123, 1415–1464.

Neiman, B. (2010): "Stickiness, Synchronization and Passthrough in Intrafirm Trade Prices," Journal of Monetary Economics, 57, 295–308.

Sheshinski, E. and Y. Weiss (1992): "Staggered and Synchronized Price Policies under Inflation: The Multiproduct Monopoly Case," Review of Economic Studies, 59, 331–59.

Stigler, G. J. and J. K. Kindahl (1970): The Behavior of Industrial Prices, Columbia University Press, New York, N.Y.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," The Review of Economics and Statistics, 86, 514–533.

5 Tables

Table 1: Summary Statistics by Bin

	1	2	3	4
Mean Employment	2996	1427	1132	1016
Median Employment	427	155	195	296
% of Prices	17.15	43.53	18.16	21.16
Mean # of Items	2.21	4.05	6.06	10.26
No. of Firms	9111	13577	3532	2160

We group firms in the Producer Price Index (PPI) by the number of goods. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods. We calculate mean and median employment by taking means and medians of the number of employees across firms in a category. % Prices denotes the fraction of prices in the PPI set by firms in each bin.

Table 2: Distribution of Firms across Sectors, by Bin

Sector	Bin 1	Bin 2	Bin 3	Bin 4
11	1.21	0.53	0.39	0.55
21	6.98	1.27	0.84	0.86
22	0.79	0.28	1.12	10.55
23	3.01	2.17	0.00	0.00
31	11.61	14.01	16.29	12.60
32	18.18	17.20	18.82	11.10
33	27.62	34.74	33.93	21.78
42	1.64	2.64	1.12	1.00
44	3.50	3.35	2.30	4.46
45	0.91	1.22	0.76	1.96
48	2.55	2.36	1.12	2.73
49	0.84	0.48	0.39	0.64
51	3.43	3.80	2.75	5.96
52	3.78	3.01	7.54	11.82
53	4.08	1.57	0.48	1.23
54	1.88	1.93	1.21	0.18
56	1.29	1.08	3.42	1.27
61	4.24	4.83	3.25	2.50
62	1.73	2.18	3.62	8.64
71	0.19	0.31	0.20	0.00
72	0.56	1.00	0.45	0.18
Total	100	100	100	100

The table shows the percentage of firms in the PPI that belong to a two-digit NAICS category in a given bin. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods.

Table 3: Distribution of Firms across Bins, by Sector

Sector	Bin 1	Bin 2	Bin 3	Bin 4	Total
11	52.86	34.76	6.67	5.71	100
21	74.13	20.19	3.48	2.20	100
22	18.85	9.95	10.47	60.73	100
23	48.25	51.75	0.00	0.00	100
31	27.73	49.90	15.15	7.22	100
32	33.79	47.65	13.62	4.95	100
33	28.23	52.92	13.50	5.34	100
42	26.18	63.00	6.98	3.84	100
44	33.51	47.70	8.56	10.23	100
45	25.94	52.19	8.44	13.44	100
48	35.67	49.09	6.10	9.15	100
49	45.03	38.60	8.19	8.19	100
51	29.57	48.87	9.23	12.34	100
52	26.91	31.96	20.92	20.22	100
53	59.08	33.97	2.69	4.27	100
54	35.68	54.56	8.92	0.83	100
56	28.37	35.58	29.33	6.73	100
61	31.83	54.14	9.52	4.51	100
62	20.41	38.37	16.67	24.55	100
71	25.37	64.18	10.45	0.00	100
72	24.52	65.87	7.69	1.92	100

The table shows the percentage of firms in the PPI that belong to a bin in a given two-digit NAICS category. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods.

Table 4: Multinomial Logit

	Negative Change				Positive Change			
	RRR	Std. Err.	z	$P > \|z\|$	RRR	Std. Err.	z	$P > \|z\|$
Fraction industry	1.0129	0.0001	92.85	0	1.0141	0.0001	109.37	0
Fraction firm	1.0710	0.0001	705.44	0	1.0750	0.0001	842.91	0
$\overline{\Delta p}$ industry	1	5.03E-09	-1.74	0.082	1	4.75E-09	11.42	0
$\overline{\Delta p}$ firm	1	2.66E-09	-4.86	0	1	2.40E-09	7.03	0
Change in product code	48.1927	3.5692	52.32	0	46.9592	3.290421	54.93	0
π_{CPI}	0.8936	0.0078	-12.88	0	1.0556	0.0081	7.07	0
Employees	1+4E-06	1.50E-07	26.57	0	1+3E-06	1.41E-07	24.5	0
Month 1	1.0060	0.0122	0.49	0.624	1.3367	0.0142	27.41	0
Month 2	0.5562	0.0073	-44.9	0	1.5790	0.0171	42.15	0
Month 3	10.8005	0.1297	198.22	0	0.8005	0.0103	-17.38	0
Month 4	1.0102	0.0122	0.84	0.4	1.0517	0.0115	4.63	0
Month 5	0.9786	0.0113	-1.88	0.06	0.9493	0.0100	-4.94	0
Month 6	0.9542	0.0112	-4.01	0	0.9484	0.0101	-4.99	0
Month 7	0.2765	0.0035	-101.95	0	1.6793	0.0172	50.69	0
Month 8	1.0547	0.0123	4.56	0	1.0129	0.0108	1.2	0.229
Month 9	1.0338	0.0129	2.67	0.008	1.0352	0.0117	3.06	0.002
Month 10	1.0699	0.0119	6.1	0	1.0773	0.0109	7.37	0
Month 11	1.0470	0.0112	4.3	0	1.0192	0.0101	1.93	0.053
R^2	47.56%							

The table reports results from the estimation of a multinomial logit model for positive and negative price changes. The base category is no price change. Coefficients show the estimated relative risk ratios, where values bigger (smaller) than 1 mean that the decision to adjust upwards or downwards is more (less) likely due to a change in the explanator and relative to the base category. Among the variables, Fraction industry and Fraction firm denote the monthly fractions of price changes within the industry or firm that have the same sign as the good under consideration which is excluded from the calculation. $\overline{\Delta p}$ industry and $\overline{\Delta p}$ firm denote the average monthly size of log price changes of all goods in the industry or firm, excluding the price of the good under consideration. We also control for two-digit industry fixed effects which we omit from the table. Change in product code is an indicator variable that takes on the value of 1 if the underlying base price has changes but the price in the data has not. Employees refers to the average number of employees per good in a firm, π_{CPI} to monthly CPI inflation, Month 1 to 11 to month dummies for January through December and Item bin 2 through 4 to dummies for the bin of each good. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods.

Table 5: Marginal Effects, Multinomial Logit

	Marginal Effects		± 1/2 Std. Dev.	
	-	+	-	+
Fraction Industry	0.06%	0.10%	1.32%	2.27%
Fraction Firm	0.32%	0.53%	8.82%	14.73%
Δp Industry	-7.16E-8%	4.30E-07%	-0.02%	0.15%
Δp Firm	-7.58E-8%	1.39E-07%	-0.06%	0.11%
π_{CPI}	-0.61%	0.48%	-0.21%	0.17%
Employees	1.93E-05%	2.52E-05%	0.26%	0.33%

Based on the regression results in Table 4, the table shows two marginal effects associated with changes in the key explanatory variables, both for upwards (+) and downwards (-) adjustment decisions. The first marginal effect is the change in the probability of adjusting upwards or downwards given a unit change in the explanatory variable, at the mean. The second marginal effect (± 1/2 Std. Dev.) denotes the change in probability associated with a change from half a standard deviation below to half a standard deviation above the mean of the explanatory variable. All reported effects are statistically significant from zero.

Table 6: Multinomial Logit, No Firm Variables

	Negative Change				Positive Change			
	RRR	Std. Err.	z	$P > \|z\|$	RRR	Std. Err.	z	$P > \|z\|$
Fraction Industry	1.0538	0.0001	503.87	0	1.0540	0.0001	556.89	0
$\overline{\Delta p}$ industry	1.0000	3.98E-09	-4.87	0	1.0000	3.65E-09	12.97	0
Change in product code	53.9963	3.6933	58.32	0	51.3618	3.2870	61.55	0
π_{CPI}	0.8876	0.0068	-15.5	0	1.0417	0.0065	6.58	0
Employees	1+6E-06	1.24E-07	44.68	0	1+5E-06	1.08E-07	46.99	0
Month 1	1.0707	0.0114	6.39	0	1.4681	0.0126	44.62	0
Month 2	0.6114	0.0071	-42.55	0	1.8924	0.0164	73.47	0
Month 3	11.1243	0.1158	231.34	0	0.9541	0.0103	-4.36	0
Month 4	0.9870	0.0106	-1.22	0.224	1.0414	0.0094	4.5	0
Month 5	0.9648	0.0099	-3.5	0	0.9317	0.0082	-8.01	0
Month 6	0.9445	0.0098	-5.5	0	0.9306	0.0083	-8.1	0
Month 7	0.3209	0.0036	-101.86	0	2.2186	0.0181	97.65	0
Month 8	1.0479	0.0108	4.53	0	1.0117	0.0090	1.32	0.187
Month 9	1.0216	0.0113	1.93	0.053	1.0355	0.0097	3.74	0
Month 10	1.0761	0.0106	7.48	0	1.0955	0.0092	10.89	0
Month 11	1.0496	0.0099	5.14	0	1.0244	0.0084	2.94	0.003
R^2	29.93%							

The table reports results from the estimation of a multinomial logit model for positive and negative price changes. The base category is no price change. Coefficients show the estimated relative risk ratios, where values bigger (smaller) than 1 mean that the decision to adjust upwards or downwards is more (less) likely due to a change in the explanator and relative to the base category. Among the variables, Fraction industry denotes the monthly fractions of price changes within the industry that have the same sign as the good under consideration which is excluded from the calculation. $\overline{\Delta p}$ industry denotes the average monthly size of log price changes of all goods in the firm, excluding the price of the good under consideration. We also control for two-digit industry fixed effects which we omit from the table. Change in product code is an indicator variable that takes on the value of 1 if the underlying base price has changes but the price in the data has not. Employees refers to the average number of employees per good in a firm, π_{CPI} to monthly CPI inflation, Month 1 to 11 to month dummies for January through December and Item bin 2 through 4 to dummies for the bin of each good. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods.

Table 7: Log Odds by Bin

	Bin 1	Bin 2	Bin 3	Bin 4
Negative Change				
Fraction Industry	1.0284	1.0184	1.0082	0.9928
Fraction Firm	1.0679	1.0712	1.0730	1.0760
π_{CPI}	0.9483	0.9120	0.9180	0.8352
Positive Change				
Fraction Industry	1.0257	1.0171	1.0094	1.0000
Fraction Firm	1.0728	1.0756	1.0769	1.0778
π_{CPI}	1.0740	1.0400	1.0470	1.0720
R^2	42.85%	47.93%	48.33%	49.10%

We repeat the multinomial logit estimation as described in Table 4 separately for each bin of goods. Bin 1 groups firms with 1 to 3 goods, bin 2 firms with 3 to 5 goods, bin 3 firms with 5 to 7 goods and bin 4 firms with more than 7 goods. The table shows estimated log odds ratios of the key variables, where values bigger (smaller) than 1 mean that the probability of a negative or positive price change is more (less) likely due to a change in the explanator and relative to the base category. All reported coefficients are statistically significant from 1.

Table 8: Marginal Effects by Bin

	Bin 1	Bin 2	Bin 3	Bin 4
Negative Change				
Fraction Industry	0.11%	0.07%	0.04%	-0.06%
Fraction Firm	0.25%	0.26%	0.35%	0.51%
π_{CPI}	-0.24%	-0.39%	-0.50%	-1.55%
Positive Change				
Fraction Industry	0.16%	0.10%	0.07%	0.01%
Fraction Firm	0.45%	0.45%	0.56%	0.78%
π_{CPI}	0.49%	0.28%	0.41%	1.02%
R^2	42.85%	47.93%	48.33%	49.10%

The table shows the bin-specific marginal effects of a unit change in the explanators around the mean on the probability of adjusting prices upwards or downwards. Marginal effects are calculated for the model estimated in Table 7. All reported effects are statistically significant from zero.

Table 9: Marginal Effects by Bin, ± 1/2 Std. Dev.

	Bin 1	Bin 2	Bin 3	Bin 4
Negative Change				
Fraction Industry	2.10%	1.41%	0.89%	-1.51%
Fraction Firm	5.38%	6.52%	9.92%	15.78%
π_{CPI}	-0.08%	-0.14%	-0.17%	-0.54%
Positive Change				
Fraction Industry	3.07%	2.12%	1.55%	0.21%
Fraction Firm	9.61%	11.46%	15.94%	24.24%
π_{CPI}	0.17%	0.10%	0.14%	0.35%
R^2	42.85%	47.93%	48.33%	49.10%

The table shows the bin-specific marginal effects of a one-standard deviation change in the explanators around the mean on the probability of adjusting prices upwards or downwards. Marginal effects are calculated for the model estimated in Table 7. All reported effects are statistically significant from zero.

Table 10: Parameters in Simulation

β	$(0.96)^{\frac{1}{12}}$
μ_P	0.21%
θ	4
ρ	0.96
σ_A	2%
σ_P	.37%

The table shows our choice of parameter values used in the simulation exercise.

Table 11: Results of Simulation

	1 Good	2 Goods	3 Goods
Frequency of price changes	15.22%	18.05%	19.72%
Absolute size of price changes	5.21%	4.29%	3.96%
Size of positive price changes	5.34%	4.46%	4.23%
Size of negative price changes	-5.02%	-4.04%	-3.57%
Fraction of positive price changes	61.68%	61.28%	59.38%
Fraction of small price changes	1.33%	20.97%	23.59%
Kurtosis	1.38	1.76	1.97
First Percentile	-6.84%	-7.60%	-8.06%
99th Percentile	7.09%	8.07%	8.63%
Synchronization measures:			
Fraction, Upwards Adjustments	-	30.25	38.11
Fraction, Downwards Adjustments	-	29.39	37.19
Correlation coefficient	-	0	0
Menu costs	0.35%	0.65%	0.75%

We perform stochastic simulation of our model in the 1-good, 2-good and 3-good cases and record price adjustment decisions in each case. Then, we calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

6 Graphs

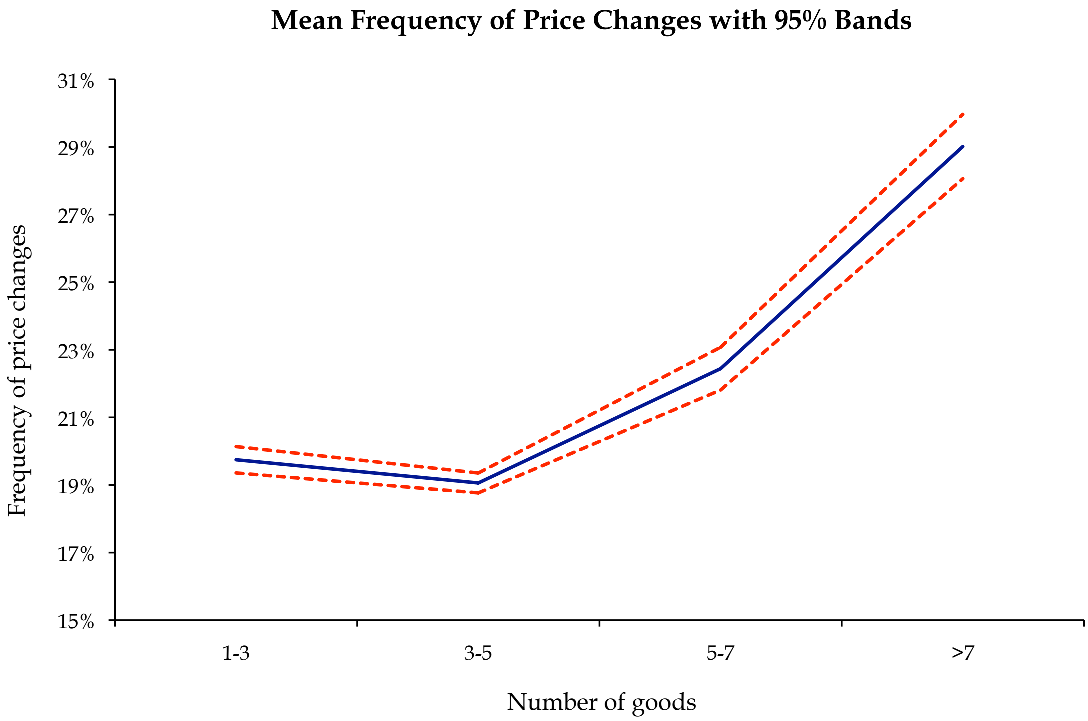

Figure 1: Mean Frequency of Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean frequency of price changes in these groups in the following way. First, we compute the frequency of price change at the good level. Then, we compute the median frequency of price changes across goods at the firm level. Finally, we report the mean across firms in a given group.

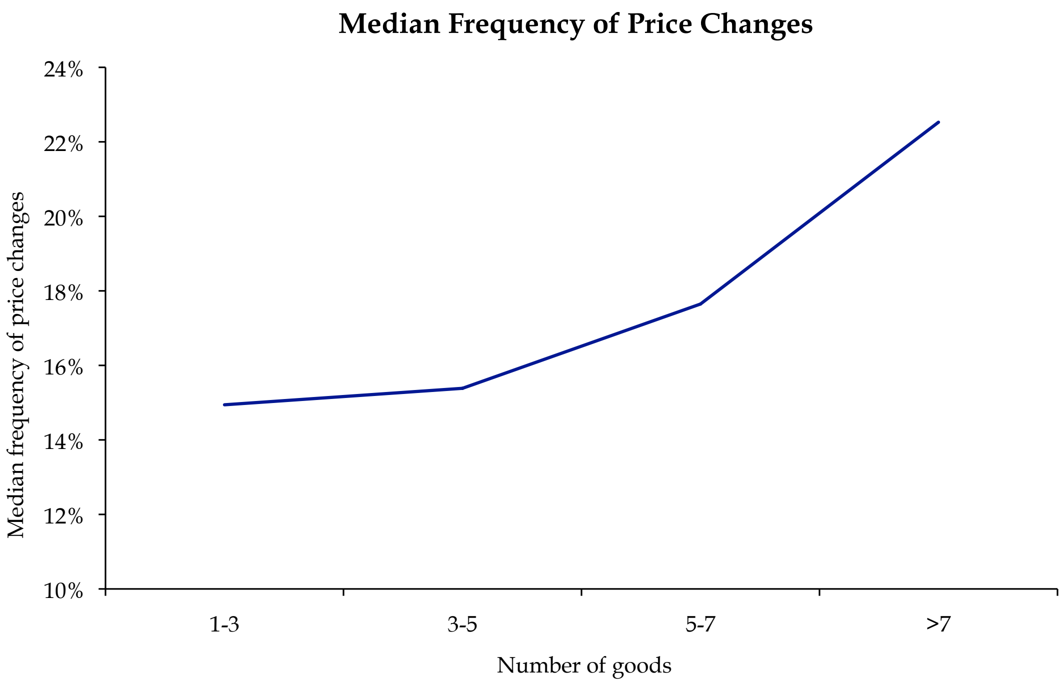

Figure 2: Median Frequency of Price Changes

Based on the PPI data we group firms by the number of goods they produce. We compute the median frequency of price changes in these groups in the following way. First, we compute the frequency of price change at the good level. Then, we compute the median frequency of price changes across goods at the firm level. Finally, we report the median across firms in a given group.

Figure 3: Mean Fraction of Positive Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean fraction of positive price changes in these groups in the following way. First, we compute the number of strictly positive good level price changes over all zero and non-zero price changes for a given firm. Then, we report the mean across firms in a given group.

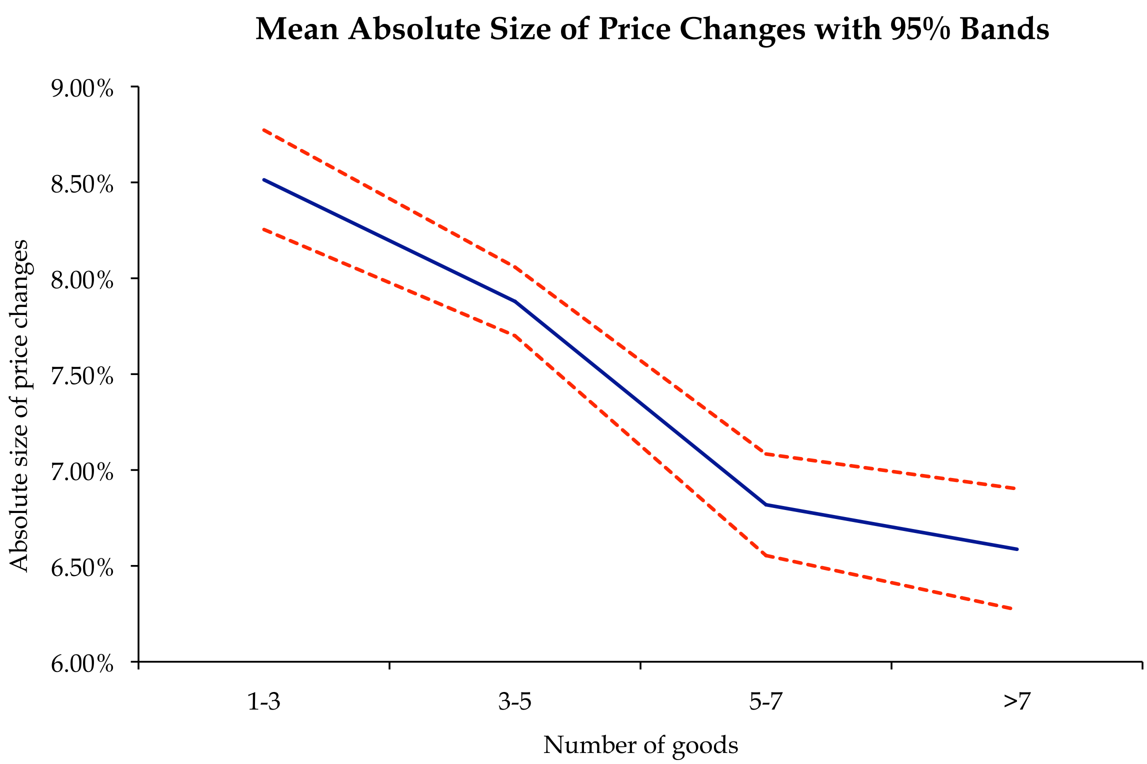

Figure 4: Mean Absolute Size of Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean absolute size of price changes in these groups in the following way. First, we compute the percentage change to last observed price at the good level. Then, we compute the median size of price changes across goods at the firm level. Then, we report the mean across firms in a given group.

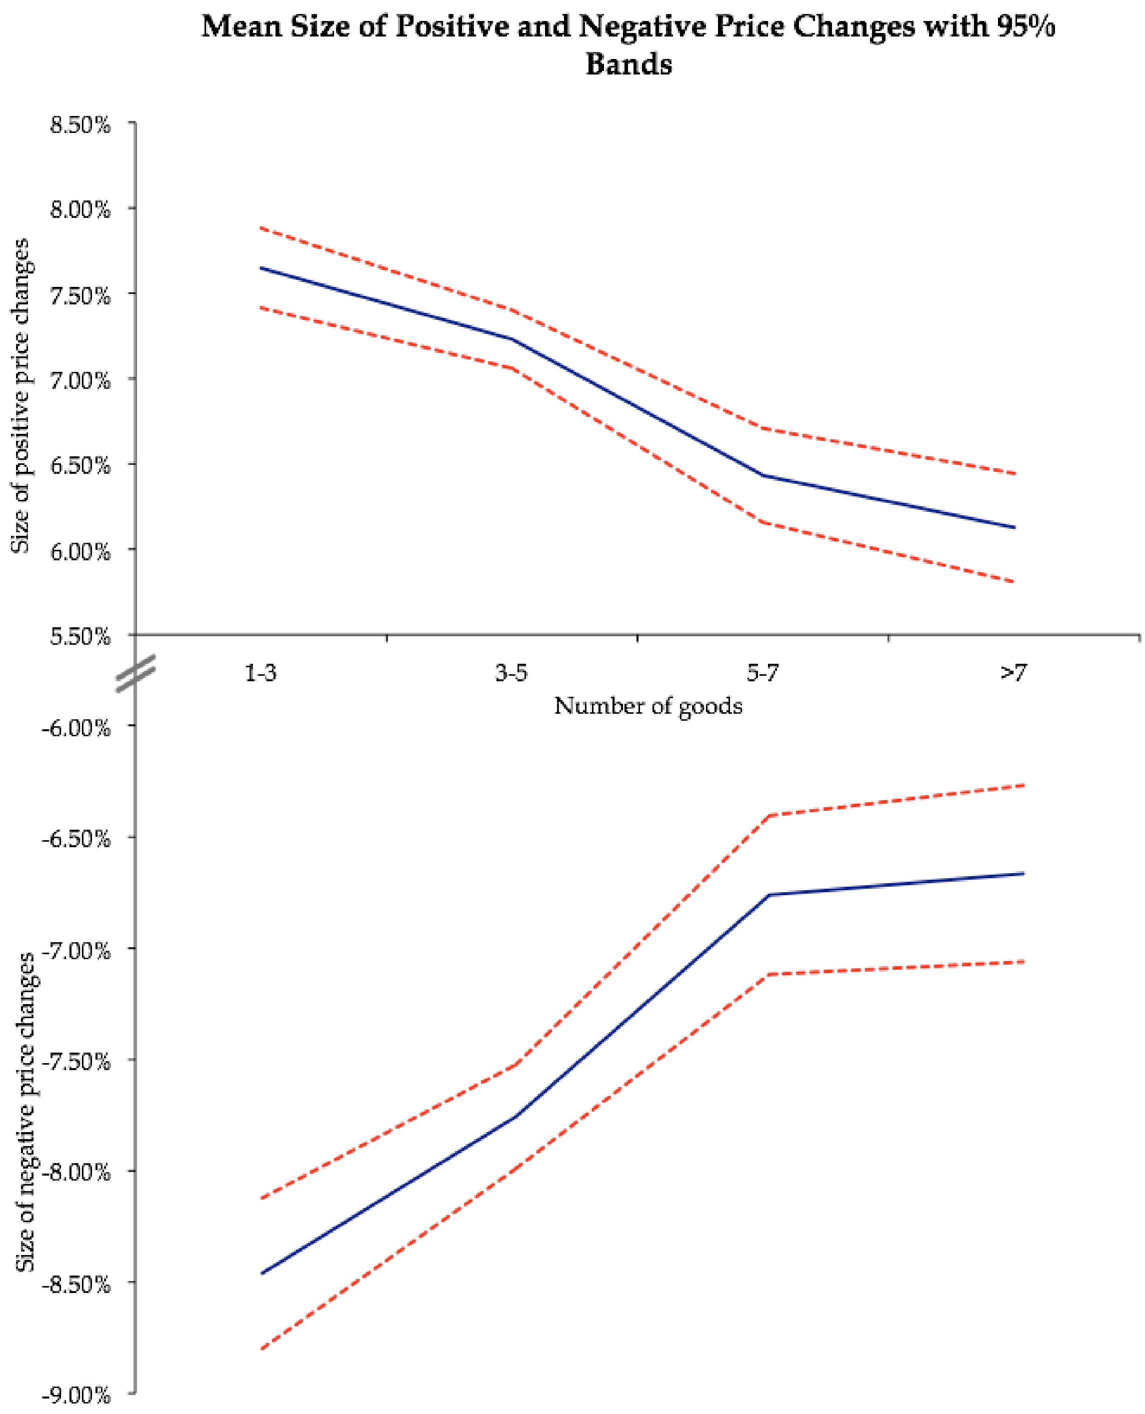

Figure 5: Mean Size of Positive and Negative Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean size of positive price changes in these groups in the following way. First, we compute the percentage change to last observed price at the good level. Then, we compute the median size of price changes across goods at the firm level. Then, we report the mean across firms in a given group.

Figure 6: Mean Fraction of Small Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean fraction of small price changes in these groups in the following way. First, we compute the fraction of price changes that are smaller than 0.5 times the mean absolute percentage size of price changes across all goods in a firm. Then, we report the mean across firms in a given group.

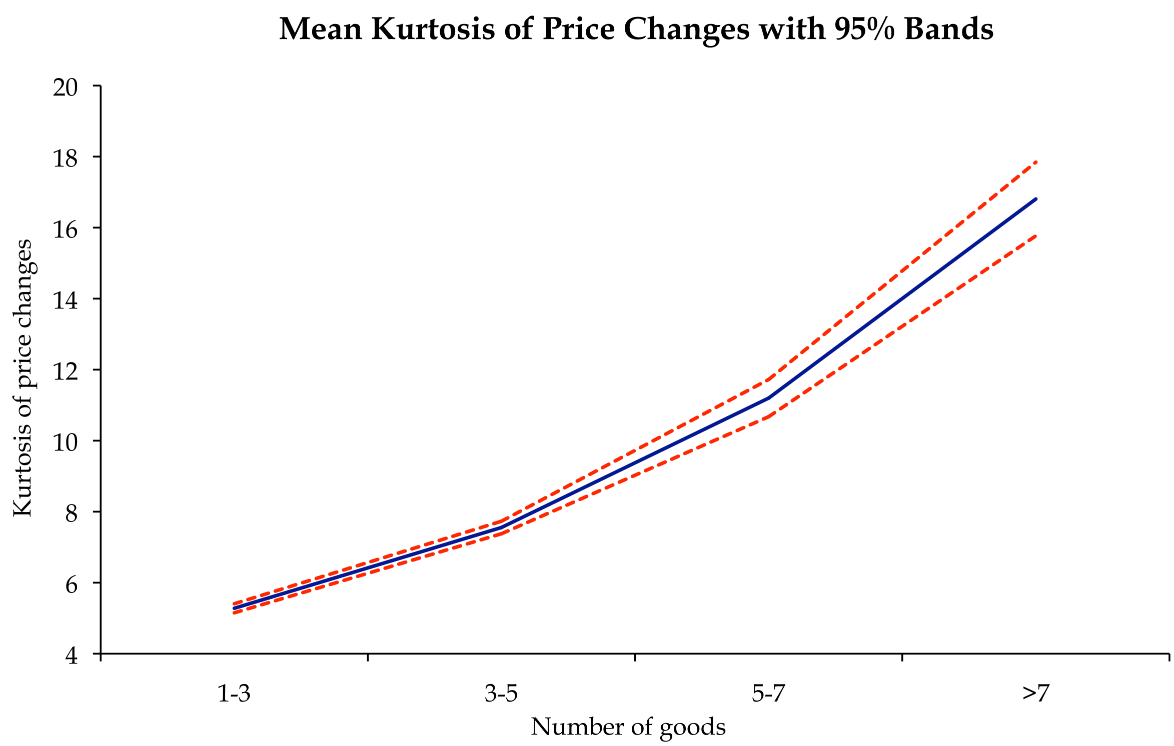

Figure 7: Mean Kurtosis of Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean kurtosis of price changes in these groups in the following way. First, we compute the kurtosis of price changes at the firm level, defined as the ratio of the fourth moment about the mean and the variance squared of percentage price changes. Then, we report the mean across firms in a given group.

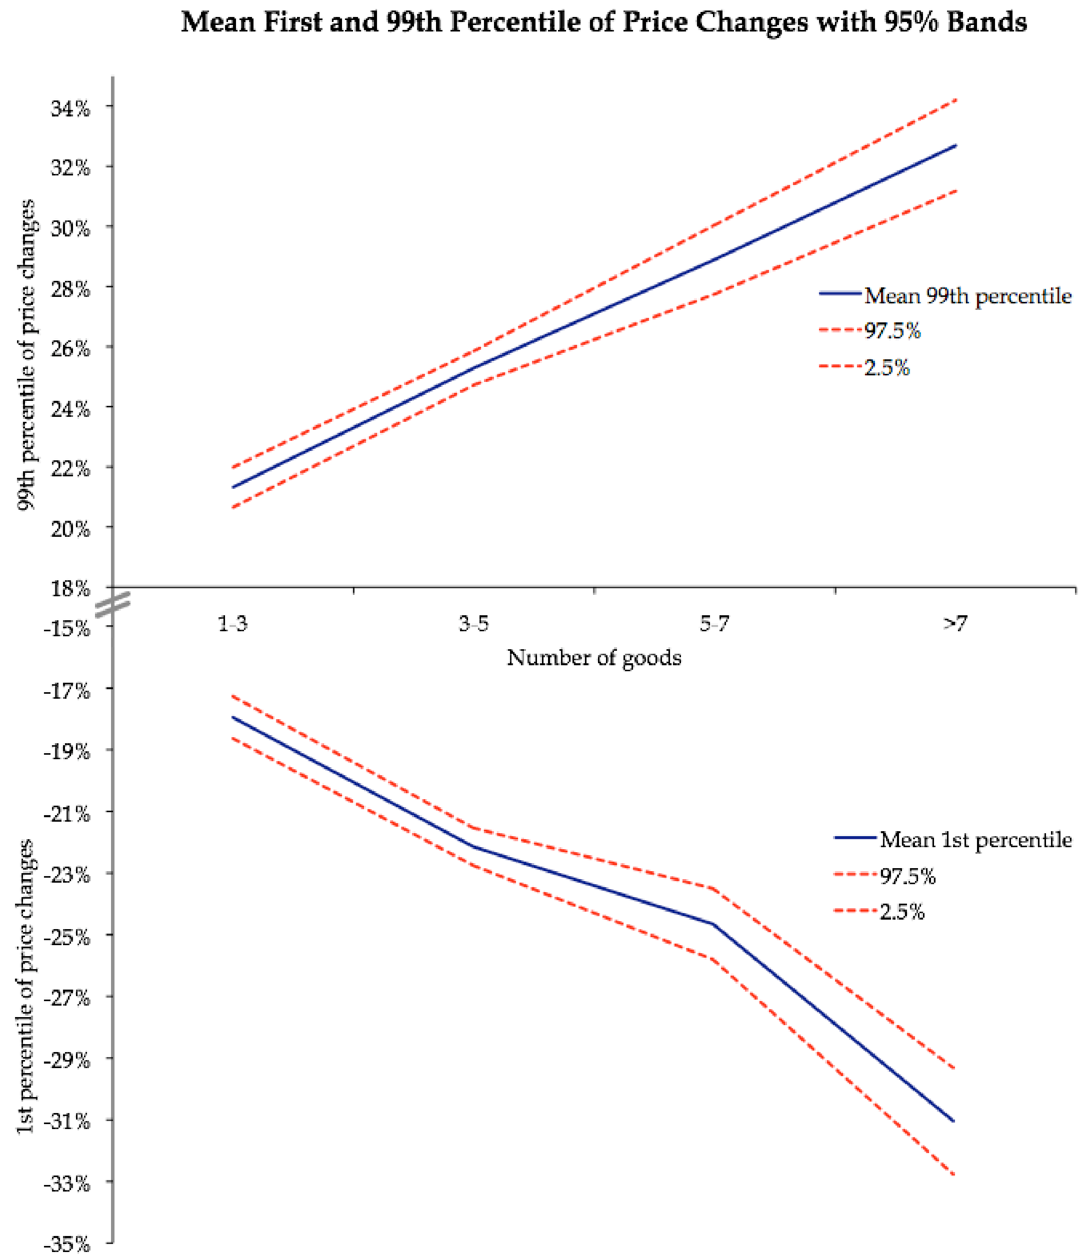

Figure 8: Mean First and 99th Percentile of Price Changes with 95% Bands

Based on the PPI data we group firms by the number of goods they produce. We compute the mean size of positive price changes in these groups in the following way. First, we compute the percentage change to last observed price at the good level. Then, we compute the median size of price changes across goods at the firm level. Then, we report the mean across firms in a given group.

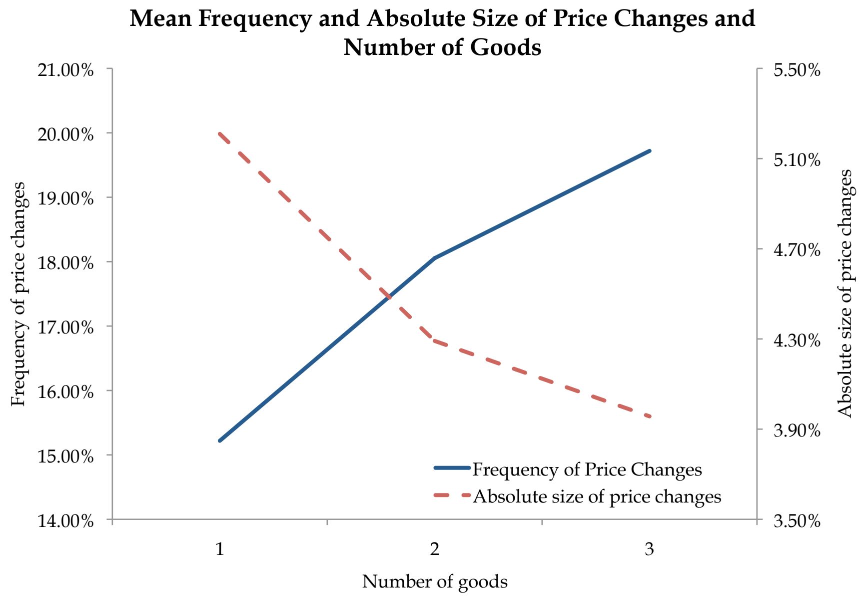

Figure 9: Fraction of Small Price Changes and Number of Goods

Figure 10: Kurtosis of Price Changes and Number of Goods

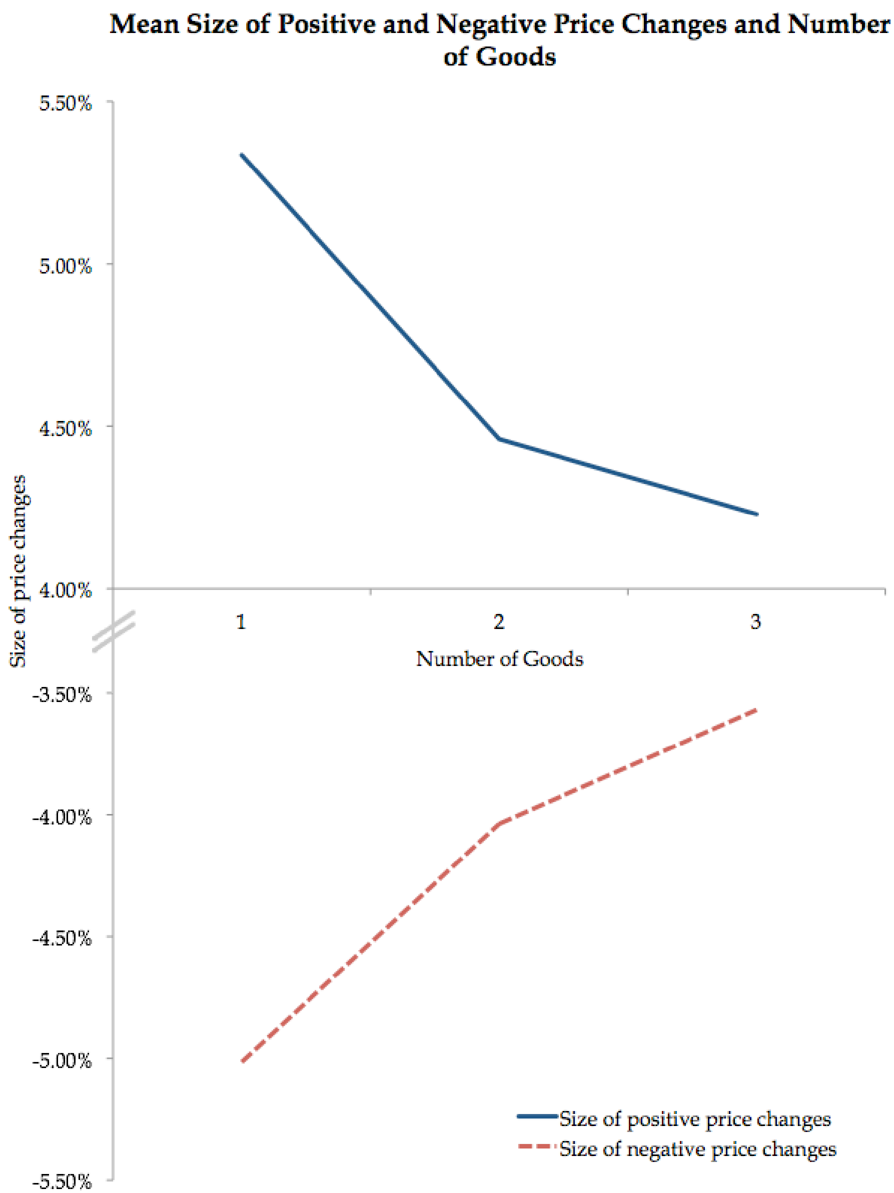

Figure 11: Mean Size of Positive and Negative Price Changes and Number of Goods

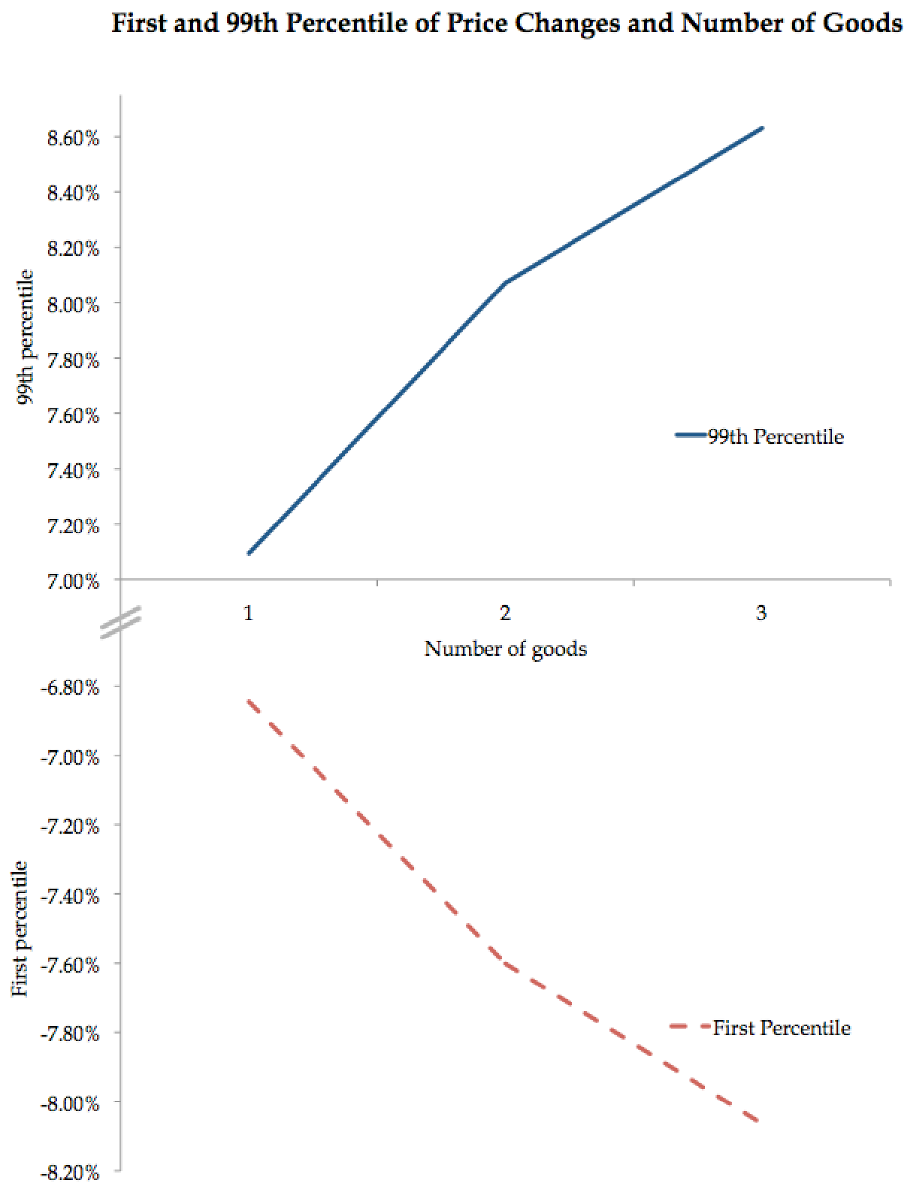

Figure 12: First and 99th Percentile of Price Changes and Number of Goods

Figure 13: Fraction of Small Price Changes and Number of Goods

Figure 14: Kurtosis of Price Changes and Number of Goods

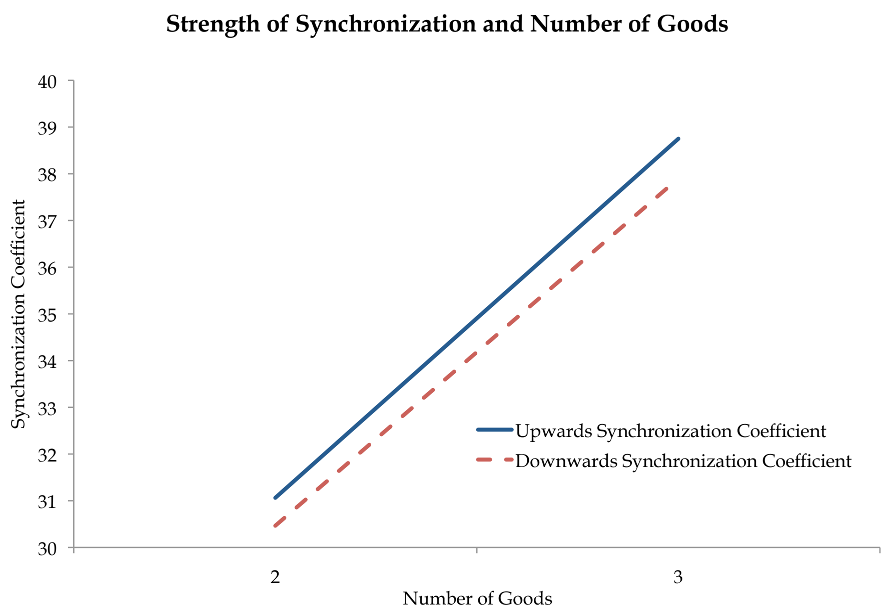

Figure 15: Strength of Synchronization and Number of Goods

APPENDIX 1

Here we describe in detail the computational algorithm used to solve the recursive problem of the firm. The state variables of the problem are last period's real prices, $\frac{p_{i,t-1}}{P_t}$, and the current productivity shocks, that is, $\mathbf{p}_{-1} = \left(\frac{p_{1,t-1}}{P_t},\ \frac{p_{2,t-1}}{P_t}, ...,\ \frac{p_{n,t-1}}{P_t} \right)$ and $\mathbf{A} = (A_{1,t}, A_{2,t}, ..., A_{n,t})$. The value functions are given by:

$$V^a(\mathbf{A}) = \max_{\mathbf{p}} \left[\pi\left(\mathbf{p};\mathbf{A}\right) - K + \beta \int\int V\left(\mathbf{p}'_{-1}, \mathbf{A}'\right) dF\left(\epsilon_A^1, \epsilon_A^2, ..., \epsilon_A^n\right) dF(\epsilon_P) \right] \tag{A-1}$$

$$V^n(\mathbf{p}_{-1}, \mathbf{A}) = \pi\left(\mathbf{p}_{-1};\mathbf{A}\right) + \beta \int\int V\left(\mathbf{p}'_{-1}, \mathbf{A}'\right) dF\left(\epsilon_A^1, \epsilon_A^2, ..., \epsilon_A^n\right) dF(\epsilon_P) \tag{A-2}$$

where $V^a(\mathbf{A})$ is the firm's value of adjusting all prices, $V^n(\mathbf{p}_{-1}, \mathbf{A})$ is the firm's value of not adjusting prices, $'$ denotes the subsequent period, and

$$V = \max\ (V^a, V^n)\,.$$

Our numerical strategy to solve for the value functions consists of two major steps. First, as described in Miranda and Fackler (2002), we approximate the value functions by projecting them onto a polynominal space. Second, we compute the coefficients of the polynomials that are a solution to the non-linear system of equations given by the value functions.

In particular, we approximate each value function, $V^a(\mathbf{A})$ and $V^n(\mathbf{p}_{-1}, \mathbf{A})$), by a set of higher order Chebychev polynomials and require (A-1) and (A-2) to hold exactly at a set of points given by the tensor product of a fixed set of collocation nodes of the state variables. This implies the following system of non-linear equations, the so-called collocation equations:

$$\Phi^a c^a = v^a(c^a) \tag{A-3}$$

$$\Phi^{na} c^{na} = v^{na}(c^{na}) \tag{A-4}$$

where c^a and c^{na} are basis function coefficients in the adjustment and non-adjustment cases and Φ^a and Φ^{na} are the collocation matrices. These matrices are given by the value of the basis functions evaluated at the set of nodes. The right-hand side contains the collocation functions evaluated at the set of the collocation nodes. Note that this is the same as the value of the right-hand side of the value functions evaluated at the collocation nodes, but where the value functions are replaced by their approximations.

We use the same number of collocation nodes as the order of the polynomial approximation. Therefore, we choose between 7-11 nodes for the productivity state variable and 15-20 nodes for the real prices. Moreover, we pick the approximation range to be $\pm$ 2.5 times the standard deviation from the mean of the underlying processes. We use Gaussian quadrature to calculate the expectations on the right-hand side, with 11-15 points for the real price transitions due to inflation while calculating the expectations due to productivity shocks exactly. For the adjustment case, we use a Nelder-Mead simplex method to find the maximum with an accuracy of the maximizer of 10^{-10}.

Next, we solve for the unknown basis function coefficients c^a and c^{na}. We express the collocation equations as two fixed-point problems:

$$c^a = \Phi^{a-1} v^a(c^a) \tag{A-5}$$

$$c^{na} = \Phi^{na-1} v^{na}(c^{na}) \tag{A-6}$$

and iteratively update the coefficients until the collocation equations are satisfied exactly.

We conduct two sensitivity analyses. First, given that zero menu costs imply flex-pricing, we verify that the approximate solution is "good" given the known analytical solution. Figure 16 shows that optimal price policies and the price policies obtained by the approximation line up a 45-degree line. The norm of the error is of order 10^{-9} and errors are equi-oscillatory, as is a usual property of approximations based on Chebychev polynomials. Second, we conduct the stochastic simulations in the maximum likelihood exercise and find that the errors between the left- and right-hand sides of (A-1) and (A-2) at points other than the collocation nodes are on average of the order of 10^{-5} or less.

Figure 16: Analytical and Numerical Optimal Flex Price

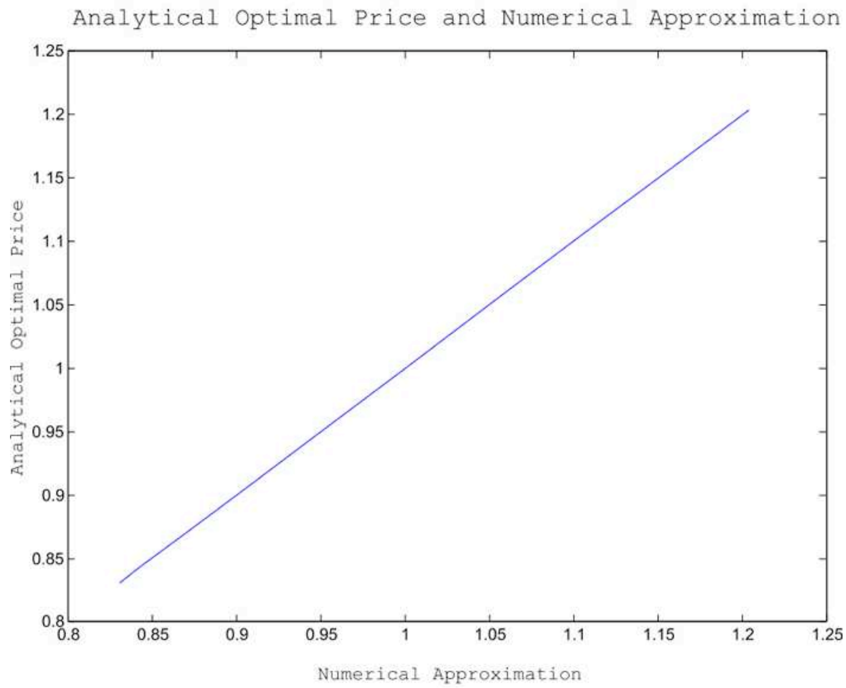

We compute the numerical solution for optimal adjustment prices given productivity shocks and zero menu costs. We compare to the analytical solution known in the flex price case. Errors are of the order of 10^{-9}.

APPENDIX 2

In this appendix, we describe in further detail the sampling procedure of the BLS which implies a monotonic relationship between the actual and sampled number of goods produced by multi-product firms. First, we document that firms with more goods have larger total sales. More goods will therefore be sampled in total from large firms. Second, we show that sales shift to goods with lower sales rank in firms with more goods. Therefore, standard survey design implies that not only more, but different goods will be sampled from larger firms. Finally, we also summarize the fraction of joint price changes in firms.

Sales Values

First, we document that firms with more goods have larger total sales in the data. Because total sales value determines the sampling probabilities of firms in the sampling selection procedure,[20] this implies that on average more goods will therefore be sampled from large firms.

We compute our measure of total sales value for each n-good-type firm in the following way. First, we compute the total dollar-value sales in a given month, year, and firm by aggregating up the item dollar-value of sales from the last time the item was re-sampled. Second, we count the number of goods for each firm in a given month and year. Third, we compute the unweighted and weighted median total sales value across all firms for a given n-good type of firm and month and year. Fourth, we calculate the mean and median sales value for an n-good type of firm

We find that firms with more goods have larger total sales: there is a strong empirical, monotonic relationship between the number of goods and the natural logarithm of the total sales. Figure 17 summarizes this relationship. Because total sales value determines the sampling probabilities of firms in the BLS sampling selection procedure, on average a higher number of goods will be collected from large firms.

Within-Firm Sales Shares

[20]Employment is another measure of firm size. The exact same results hold for employment: firms with more goods have a larger number of employees.

Second, we show that sales shift to goods with lower sales rank in firms with more goods. Therefore, standard survey design such as sampling proportional to size implies that not only more, but different goods will likely be sampled from larger firms.

We compute within-firm sales shares and sales ranks in the following way. First, we compute the total dollar-value sales for a given month, year, and firm by aggregating up the dollar-value of sales of the good from the last time the item was re-sampled. Second, we calculate the good-specific sales shares for each firm in a given month and year. Third, we rank the goods in each firm according to these sales shares. Fourth, we count the number of goods for each firm in a given month and year. Fifth, we compute the mean sales shares for an r-ranked good in an n-good firm in a given month and year, across all firms. Sixth, we compute the sales-weighted mean for an r-ranked good in an n-good firm over time. These calculations give us the sales share representative of an r-ranked good in a firm with n goods.[21]

We find that sales shift to goods with lower sales rank in firms with more goods. For example, the representative sales share of the best-selling good in a two-good firm is 63% while it is 45% for the second good. For a three-good firm, the sales shares are 45%, 35%, and 30%. Table 12 summarizes sales shares by the number of goods and rank of the goods. The table covers firms with up to 11 goods which account for more than 98% of all prices in the data. Under standard survey designs such as sampling proportional to size and a fixed survey budget not only more, but different goods are more likely to be sampled when firms produce more goods.

[21]Note that these shares do not have to sum up to 100% in an n-good firm by way of computation.

Figure 17: Log Mean Firm Sales Value by Number of Goods

Table 12: Mean Sales Shares of r-Ranked Goods for n-Good Firms, Sales-Weighted

#Goods /#Rank	1	2	3	4	5	6	7	8	9	10	11
1	100.000%	62.861%	45.110%	35.916%	37.090%	28.301%	25.517%	19.084%	25.532%	14.393%	16.642%
	(0.000)%	(0.131)%	(0.075)%	(0.086)%	(0.413)%	(0.097)%	(0.337)%	(0.381)%	(0.480)%	(0.344)%	(0.427)
2		44.952%	34.453%	27.594%	22.966%	19.963%	18.396%	15.209%	15.749%	12.297%	10.748%
		(0.081%)	(0.051%)	(0.034%)	(0.035%)	(0.049%)	(0.170%)	(0.091%)	(0.260%)	(0.089%)	(0.106%)
3			29.499%	24.051%	19.703%	17.428%	15.896%	14.223%	13.160%	11.316%	10.514%
			(0.028%)	(0.030%)	(0.023%)	(0.031%)	(0.032%)	(0.069%)	(0.134%)	(0.073%)	(0.105%)
4				22.360%	18.251%	16.162%	14.357%	13.318%	11.375%	10.615%	9.835%
				(0.030%)	(0.039%)	(0.013%)	(0.034%)	(0.050%)	(0.049%)	(0.038%)	(0.057%)
5					17.579%	15.197%	13.577%	12.707%	10.600%	10.212%	9.393%
					(0.044%)	(0.007%)	(0.028%)	(0.052%)	(0.066%)	(0.028%)	(0.032%)
6						14.748%	12.629%	11.771%	10.185%	9.824%	9.051%
						(0.011%)	(0.038%)	(0.016%)	(0.070%)	(0.032%)	(0.017%)
7							12.356%	11.417%	10.072%	9.578%	8.924%
							(1.261%)	(1.165%)	(1.028%)	(0.978%)	(0.911)
8								11.302%	9.679%	9.357%	8.731%
								(1.153%)	(0.988%)	(0.955%)	(0.891%)
9									9.943%	9.257%	8.562
									(1.015%)	(0.945%)	(0.874%)
10										9.222%	8.456%
										(0.027%)	(0.064%)
11											8.528%
											(0.057%)

Based on the PPI data, we first compute the total dollar-value sales for a given month, year, and rm by aggregating up the dollar-value of sales of each good. Second, we calculate the good-specic sales shares for each rm in a given month and year. Third, we rank goods in each rm according to the sales shares. Fourth, we count the number of goods for each rm in a given month and year. Fifth, we compute mean sales shares for an r-ranked good in an n-good rm in a given month and year, across all rms. Sixth, we compute the sales-weighted mean for an r-ranked good in an n-good rm over time.

Fractions of Joint Price Adjustment

We compute the representative monthly fraction of goods in a firm which change prices. We find that 71.36% to 75% of all prices change jointly in each month. This suggests that our assumption of complementarities in the price adjustment is a good first order approximation.

First, we compute the fraction of goods which change prices in a given firm, year, and month, conditional on observing at least one price change. Second, we compute the median of this fraction over time, for each firm. Third, we compute means and medians across all firms in the data.

Table 13: Fractions of Joint Price Adjustment

	f_{joint}	f^{+}_{joint}	f^{-}_{joint}
Mean	71.36%	64.83%	61.26%
Median	75.00%	66.67%	52.90%
Std. Error	0.17%	0.17%	0.19%
Std. Dev.	27.24%	27.49%	28.85%

The table summarizes monthly fractions of price changes in a firm, computed conditional on at least one price adjustment. We denote by f_{joint} the fraction irrespective of sign of price change, by f^{+}_{joint} upwards joint fractions and by f^{-}_{joint} downwards joint fractions.

APPENDIX 3

In this section, we present the simulation results derived from alternative modeling assumptions. In each case, we replace or supplement the assumption of economies of scope in price adjustment by one of the alternative assumptions. We compute statistics and trends as in the main model.

Table 14: Results of Simulation: Substitution

	1 Good	2 Goods	3 Goods
Frequency of price changes	15.22%	21.72%	30.75%
Absolute size of price changes	5.21%	4.31%	3.58%
Size of positive price changes	5.34%	4.45%	3.61%
Size of negative price changes	-5.02%	-4.11%	-3.53%
Fraction of positive price changes	61.68%	59.55%	59.20%
Fraction of small price changes	1.33%	2.86%	1.92%
Kurtosis	1.38	1.53	1.45
First Percentile	-6.84%	-6.19%	-5.58%
99th Percentile	7.09%	6.46%	5.64%
Synchronization measures:			
Fraction, Upwards Adjustments	-	0.48	10.39
Fraction, Downwards Adjustments	-	-0.62	10.06
Correlation coefficient	-	0	0
Menu Cost	0.35%	0.7%	1.05%
Elasticity of substitution	4	6	9

We perform stochastic simulation of our model in the 1-good, 2-good and 3-good cases and record price adjustment decisions in each case. Then, we calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. As we increase the number of goods, we increase the elasticity of substitution θ. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

Table 15: Results of Simulation: Correlation

	1 Good	2 Goods	3 Goods
Frequency of price changes	15.22%	17.09%	14.25%
Absolute size of price changes	5.21%	4.52%	4.75%
Size of positive price changes	5.34%	4.61%	4.91%
Size of negative price changes	-5.02%	-4.36%	-4.49%
Fraction of positive price changes	61.68%	62.35%	63.72%
Fraction of small price changes	1.33%	16.77%	18.69%
Kurtosis	1.38	1.77	2.05
First Percentile	-6.84%	-7.83%	-9.38%
99th Percentile	7.09%	8.28%	9.89%
Synchronization measures:			
Fraction, Upwards Adjustments	-	31.06	38.73
Fraction, Downwards Adjustments	-	30.41	37.88
Correlation coefficient	-	0.6	0.6
Menu Cost	0.35%	0.70%	1.05%

We perform stochastic simulation of our model in the 1-good, 2-good and 3-good cases and record price adjustment decisions in each case, allowing for correlation of the productivity shocks $A_{i,t}$ in the multi-good cases. Then, we calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

Table 16: Benchmark Case

	2 Goods MP Firm	2 Firms	3 Goods MP Firm	3 Firms
Frequency of price changes	18.05%	15.30%	19.72%	15.30%
Size of absolute price changes	4.29%	5.24%	3.96%	5.23%
Size of positive price changes	4.46%	5.35%	4.23%	5.35%
Size of negative price changes	-4.04%	-5.08%	-3.57%	-5.08%
Fraction of positive price changes	61.28%	62.14%	59.38%	62.14%
Fraction of small price changes	20.97%	3.01%	23.59%	3.01%
Kurtosis	1.76	1.55	1.97	1.52
1st Percentile	-7.60%	-7.47%	-8.06%	-7.47%
99th Percentile	8.07%	8.02%	8.63%	8.02%
Synchronization measures:				
Fraction, Upwards Adjustments	30.25	0.33	38.11	14.46
Fraction, Downwards Adjustments	29.39	-0.63	37.19	13.87
Correlation coefficient	0	0	0	0
Menu cost	0.65%	0.35%	0.75%	0.35%

We perform stochastic simulation of our model for the 2-good and 3-good multi-product firms as in Table 11. Results from these simulations are summarized under the columns "MP Firms." In addition, we simulate two, and respectively three 1-good firms subject to common inflationary shocks but completely independent productivity draws. We record price adjustment decisions and calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

Table 17: Results of Simulation: Menu of Menu Costs

	1 Good	2 Goods	
Frequency of price changes	15.22%	18.05%	26.98%
Size of absolute price changes	5.21%	4.29%	3.16%
Size of positive price changes	5.34%	4.46%	3.56%
Size of negative price changes	-5.02%	-4.04%	-2.67%
Fraction of positive price changes	61.68%	61.28%	55.34%
Fraction of small price changes	1.33%	20.97%	24.95%
Kurtosis	1.38	1.76	3.95
1st Percentile	-6.84%	-7.60%	-7.87%
99th Percentile	7.09%	8.07%	11.41%
Synchronization measures:			
Fraction, Upwards Adjustments	-	30.25	-0.20
Fraction, Downwards Adjustments	-	29.39	-0.48
Menu costs (K_1, K_2, K_{12})	(-,-,0.35)%	(-,-,0.65)%	(0.35,0.35,0.65)%
Correlation coefficient	-	0	0

We perform stochastic simulation of our model in the 1-good, and 2-good cases, allowing 2-good firms to adjust 0, 1, or 2 goods simultaneously. The cost of adjusting one good only is K_1 or K_2, and joint adjustment costs K_{12}. We record price adjustment decisions and calculate statistics for each case as described in the text. In the 2-good case, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

Table 18: Results of Simulation: Demand Interactions

	1 Good	2 Goods	3 Goods	2 Goods	3 Goods
		$\gamma = -0.1$		$\gamma = 0.1$	
Frequency of price changes	15.22%	18.76%	19.36%	21.66%	20.76%
Size of absolute price changes	5.21%	4.143%	4.142%	3.90%	3.71%
Size of positive price changes	5.34%	4.34%	4.41%	4.20%	3.93%
Size of negative price changes	-5.02%	-3.85%	-3.77%	-3.50%	-3.39%
Fraction of positive price changes	61.68%	61.27%	59.14%	58.15%	59.92%
Fraction of small price changes	0.98%	22.05%	24.56%	21.81%	23.61%
Kurtosis	1.38	1.88	2.01	2.03	2.02
1st Percentile	-6.84%	-7.88%	-8.61%	-7.84%	-7.77%
99th Percentile	7.09%	8.35%	9.10%	8.53%	8.11%
Synchronization measures:					
Fraction, Upwards Adjustments	-	30.28	38.13	30.04	38.02
Fraction, Downwards Adjustments	-	29.45	37.13	29.38	36.98
Correlation coefficient	-	0	0	0	0
Menu Cost	0.35%	0.65%	0.75%	0.65%	0.75%

We perform stochastic simulation of our model in the 1-good, 2-good and 3-good cases, allowing for interactions in demand through a profit function $\pi_t = \sum_i^{n=3} \left(\frac{p_{i,t}}{P_t} - \frac{\bar{w}}{A_{i,t}}\right)\left(\frac{p_{i,t}}{P_t}\right)^{-\theta}\left(\frac{p_{-i,t}}{P_t}\right)^{-\gamma}$ and record price adjustment decisions in each case. Then, we calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.

Table 19: Results of Simulation: Demand Shocks

	1 Good	2 Goods	3 Goods
Frequency of price changes	15.22%	5.14%	5.52%
Size of absolute price changes	5.21%	4.09%	3.94%
Size of positive price changes	5.34%	4.09%	3.94%
Size of negative price changes	-5.02%	0%	0%
Fraction of positive price changes	61.68%	100.00%	100.00%
Fraction of small price changes	1.33%	0.00%	0.00%
Kurtosis	1.38	3.71	3.90
1st Percentile	-6.84%	3.73%	3.49%
99th Percentile	7.09%	4.75%	4.71%
Correlation coefficient	-	0	0
Menu costs	0.35%	0.65%	0.75%

We perform stochastic simulation of our model in the 1-good, 2-good and 3-good cases, allowing for demand shocks $Z_{i,t}$ instead of productivity shocks. This implies a period profit function $\pi_t = \sum_i^{n=3} \left(\frac{p_{i,t}}{P_t} - \bar{w}\right) Z_{i,t} \left(\frac{p_{i,t}}{P_t}\right)^{-\theta}$. We record price adjustment decisions and calculate statistics for each case as described in the text. In the 2-good and the 3-good cases, we report the mean of the good-specific statistics. We obtain the synchronization measure from a multinomial logit regression analogous to the empirical multinomial logit regression. We control for inflation. Menu costs are given as a percentage of steady state revenues.