NBER WORKING PAPER SERIES

VOLATILITY AND PASS-THROUGH

David Berger
Joseph S. Vavra

Working Paper 19651
http://www.nber.org/papers/w19651

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2013

This research was funded in part by the Initiative on Global Markets at the University of Chicago Booth School of Business. We would like to thank seminar participants at Princeton, Chicago Booth, Cleveland Fed, Duke Macro Jamboree, Chicago Fed, SED, NBER SI IFM, Columbia, Purdue, Northwestern, Dallas Fed Uncertainty Conference, EIEF, CESifo, and Michigan. We would also like to thank our discussant Linda Tesar as well as Rudi Bachmann, Nick Bloom, Ariel Burstein, Jeff Campbell, Gabe Chodorow-Reich, Larry Christiano, Allan Collard-Wexler, Marty Eichenbaum, Tarek Hassan, Matthias Kehrig, Oleg Itskhoki, Amy Meek, Emi Nakamura, Brent Neiman, Sergio Rebelo, Johannes Stroebel and Rozi Ulics. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Volatility and Pass-through
David Berger and Joseph S. Vavra
NBER Working Paper No. 19651
November 2013
JEL No. E10,E30,E31,F31

ABSTRACT

What drives countercyclical volatility? A large literature has documented that many economic variables are more disperse in recessions, but this could either occur because shocks get bigger or because firms respond more to shocks which are the same size. Existing evidence that the dispersion of endogenous variables rises in recessions cannot tell us which of volatility or responsiveness is getting bigger, and these two explanations have very different policy implications. However, we document new facts in the open economy environment and show that they can be used to disentangle these explanations. In particular, we use confidential BLS micro data to show that there is a robust positive relationship between exchange rate pass-through and the dispersion of item-level price changes. We then argue that changes in responsiveness can explain this fact while volatility shocks cannot.

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

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

1 Introduction

What drives countercyclical volatility? A large literature has documented that many economic variables are more dispersed during recessions than in booms.[1] The dominant explanation for this phenomena is that firms face more volatile shocks during recessions. However, increased dispersion could occur either because shocks get bigger or because firms respond more to shocks which are the same size. The existing evidence that shows that the dispersion of *endogenous*[2] variables rises in recessions cannot tell us which of volatility or responsiveness is getting bigger. Indeed, a recent theoretical literature has emerged challenging the role of volatility shocks and instead emphasizing the potential role of the responsiveness channel.[3] It is important to distinguish between these two explanations because they have very different implications for the source of business cycles and how policy should mitigate them. However, this is a challenging identification problem, so there exists little empirical evidence on this point. In this paper, we show that one can directly separate changes in responsiveness from changes in volatility using the open economy environment.

We make this argument in two steps: 1) We first use confidential BLS import price data to show that item-level price dispersion is countercyclical, confirming the results in Bloom et al. (2012) and Vavra (2014). Next, we document a strong positive relationship between price change dispersion and exchange rate pass-through. This positive relationship holds both at the item-level (cross-section) and at the month-level (time-series). That is: a) Individual items with high dispersion of price changes across time have greater exchange rate pass-through. b) During times when the cross-sectional dispersion of price changes across items is high, there is greater exchange rate pass-through. This provides suggestive "model free" evidence of a positive relationship between responses to the exchange rate and dispersion. The open-economy is key to this first result as it provides a large and observable cost shock (the nominal exchange rate) which can be used to measure responsiveness.

2) Next, we expand on this evidence and use these empirical moments to infer the quantitative role of different shocks across firms and time. We build a structural price-setting model with various sources of heterogeneity and formally estimate their importance for explaining the comovement between dispersion and pass-through which we document in the first part of the paper. While various forces affect pass-through and price change dispersion in isolation, we show that the positive comovement of pass-through and price dispersion in the data allows us to discriminate between explanations. We find strong evidence that time-variation in responsiveness rather than volatility explains countercyclical dispersion. In particular, volatility shocks imply a counterfactual negative or zero relationship between dispersion and pass-through while responsiveness shocks imply a positive relationship, consistent with the data. This quantitative result is quite robust and holds regardless of whether price adjustment is frictionless, time-dependent or state-dependent.

[1] For example, Bloom (2009), Bloom et al. (2012), and Vavra (2014).

[2] For example, Bloom (2009) (sales growth), Bloom et al. (2012) (revenue based TFP and employment growth), and Vavra (2014) (prices).

[3] See Bachmann and Moscarini (2012) and Decker et al. (2015) for some theoretical models in which countercyclical dispersion arises from greater endogenous responses.

Given the central role the positive correlation between dispersion and pass-through plays in our identification strategy, we address its robustness in two complementary ways. First, we do as much as possible directly in the data to control for observable co-variates that could induce spurious correlation between pass-through and dispersion. Second, we explicitly take into account censoring and sample selection biases when estimating our quantitative model to mitigate concerns that our results are driven by these channels, which are more challenging to address directly in the data.

Broadly speaking, in the empirical section of the paper our results are robust to controls for "micro" observables, "macro" observables, and to a variety of measurement error related concerns. Amongst "micro" controls, we show that the positive relationship between pass-through and dispersion is not driven by differences in an item's frequency of adjustment or product substitution, degree of product differentiation, country or currency of import or the volatility of that currency. Our "macro controls" include the average frequency of price changes and product substitution, real GDP growth, exchange rate volatility, seasonality and long-run trends. We also run a wide range of placebo checks, outlier robust regressions, and alternative pass-through specifications to argue that our fact is not driven by small samples or other measurement error related issues.

In our final controls for observable differences, we show that the positive correlation between pass-through and dispersion remains strong, both economically and statistically, even when we exclude data from the Great Recession. This is an important robustness check because the 2008 recession is a large outlier in price change dispersion, which makes time-variation in dispersion appear small in the rest of the sample by comparison. There were also obviously many other ways in which the Great Recession differed from usual business cycles. Thus, if our results were driven by this extreme event then it would be difficult to extrapolate our conclusions to other periods. Fortunately, we find that even after excluding 2008, time-variation in pass-through and dispersion is large in both absolute and statistical terms: constant pass-through specifications are overstated by 50 percent during the mid-1990s and understated by 50 percent during the 2001 Recession. Price change dispersion also has fluctuations of greater than 50 percent in this earlier part of the sample. Thus, the economic significance of our conclusions does not rely on identification from this extreme period.

In addition to these empirical robustness checks, we deal with concerns related to sample selection and censoring explicitly when we estimate our theoretical model using indirect inference. In particular, when we move from pure empirical results to their theoretical interpretation, we always a) generate data from our model using the same sample selection procedures that are in the data b) simulate the model with sample sizes and stickiness as in the BLS data and c) always treat simulated and actual data identically and we use no information from simulated data that is not available in actual data in the estimation. Thus, if sampling, selection or censoring drove our empirical results, they should drive similar effects in our simulated data. Finally, the fact that we find similar relationships in the cross-section and the time-series (as predicted by the structural model) also provides some reassurance that our results are not spurious.

The first half of the paper is purely empirical and focuses on establishing the positive relationship between pass-through and dispersion. In the second half of the paper, we provide a structural interpretation of this fact and assess the role of many theoretical channels that can affect pass-through and price change dispersion. The exchange rate pass-through literature often emphasizes variable markups or strategic complementarities as a channel that can affect firms' desired exchange rate pass-through. Changes in the strength of this channel will also affect price change dispersion by changing firms' desired responses to idiosyncratic shocks. Thus, this "responsiveness" channel could generate movements in pass-through and dispersion even with no changes in underlying volatility.

Conversely, price change dispersion can also move with changes in the volatility of idiosyncratic shocks, even if there is no change in endogenous responses to those shocks. As previously mentioned, a large literature including Bloom (2009), and Vavra (2014) uses such "uncertainty" shocks to explain countercyclical dispersion of outcomes such as firm growth and price changes. Their ability to explain time-variation in measured exchange rate pass-through has heretofore been unexplored.

While the above channels have received the most attention in the literature, they are far from exhaustive. Differences in the size of adjustment costs, the sensitivity of firms' costs to exchange rates, changes in the volatility of exchange rates, or changes in the "commonality" of aggregate shocks all have the potential to affect price change dispersion and exchange rate pass-through.

We use indirect inference to formally estimate the importance of these channels in a quantitative menu cost model which builds on Gopinath and Itskhoki (2010) and Burstein and Gopinath (2013). While the basic modeling framework is intentionally standard, our emphasis is not on the model itself but is instead on what it can tell us about the underlying nature of firm-level shocks. The open economy environment is key to answering this question as it provides the source of identification in the model. For example, with only data on outcomes and no data on observable shocks, changes in "uncertainty" and changes in "responsiveness" are observationally equivalent: an increase in either implies an increase in the dispersion of observable firm decisions. However, these two theories have very different implications for how firms will respond to observable exchange rate shocks.

Using our identified quantitative model, we show that heterogeneous import shares, menu costs, changes in the volatility of exchange rates, and shocks to the "commonality" of aggregate shocks are all unable to explain our results.[4] The model also strongly rejects uncertainty shocks. In contrast, our estimated model assigns an extremely important role to changes in responsiveness arising from variable markups.[5] Together, our new empirical results and their structural interpretation suggest that the literature studying countercyclical dispersion has embraced time-varying volatility too quickly. At least for our data, time-varying responsiveness is much more important.

Our results provide further support for the mechanisms emphasized by Gopinath and Itskhoki (2010) and Atkeson and Burstein (2008). Our paper is closely related to Gopinath and Itskhoki (2010), which argues that variable markups (which generate heterogeneous responsiveness) are necessary to explain cross-sectional differences in pass-through. An important distinction between our empirical work and theirs is that they study only permanent differences across items, while we show there is also striking time-variation in pass-through.[6] In addition, we formally estimate our struc-

[4] We also model various sources of measurement error and show that these cannot explain our empirical patterns.

[5] For concreteness, our variable responsiveness arises from strategic-complementarities that arise under Kimball demand, but this is largely for illustrative purposes. Other forms of strategic-complementarity have similar implications for responsiveness, price change dispersion, and exchange rate pass-through.

[6] Our cross-item facts are also distinct in two dimensions: 1) We show that there is a relationship between dispersion

tural model via indirect inference, which allows us to more precisely measure the joint-distribution of various sources of heterogeneity and to explicitly deal with sample selection and censoring concerns.

Many recent papers have argued that the distribution of price changes has important implications for aggregate price flexibility. For example, Midrigan (2011) and Alvarez et al. (2014) show that theory assigns a large role to the price change distribution in shaping the average response of inflation to nominal shocks. Vavra (2014) argues that increases in the dispersion of price changes during recessions should lead to increases in aggregate price flexibility. However, this existing literature has two limitations: First, the link between the distribution of price changes and price flexibility is theoretical rather than empirical. Second, existing models must make strong assumptions about the sources of variation in dispersion across time. In this paper, we address these limitations by first providing model free empirical evidence that time-variation in price change dispersion strongly predicts time-variation in import price flexibility. We then provide additional insight into the underlying forces that shape this relationship and argue that time-variation in the competitive structure of markets or other shocks that induce time-variation in firm responsiveness appear more consistent with our data than the "uncertainty" shocks typically assumed. Understanding the source of time-varying dispersion is important for policy design as policies designed to reduce uncertainty almost certainly differ from policies designed to alter market structure and firms' responsiveness.

The remainder of the paper proceeds as follows: Section 2 contains our main empirical findings. Section 3 discusses the implications for time-varying pass-through. Section 4 lays out a basic flexible price model that demonstrates how primitives can potentially generate a positive relationship between pass-through to price change variance. Section 5 estimates a quantitative structural model to argue that variation in responsiveness best explains the data, and Section 6 concludes.

2 Empirical Results

2.1 Data Description

In this section we describe the data employed in this study. We use confidential micro data on import prices collected by the Bureau of Labor Statistics for the period 1994-2011. This data is collected on a monthly basis and contains information on import prices for very detailed items over time. This data set has previously been used by Gopinath and Rigobon (2008), Gopinath et al. (2010), Gopinath and Itskhoki (2010), Neiman (2010), and Berger et al. (2012). Below, we provide a brief description of how the data is collected. See the IPP (Import Price Program) Data Collection

and pass-through that is independent of frequency. 2) We study what Gopinath and Itskhoki (2010) call medium-run pass-through (MRPT) rather than long-run pass-through (LRPT). MRPT measures the fraction of exchange rate movements passed-through into an item's price after one price adjustment whereas LRPT captures pass-through over an item's entire life. While much of the literature has moved towards the use of LRPT, MRPT is the relevant pass-through concept for measuring time-varying price flexibility at business cycle frequencies. MRPT measures how shocks today are passed into price changes today whereas LRPT measures how shocks will transmit to prices potentially years into the future. By construction, LRPT cannot measure time-varying aggregate dynamics since LRPT is fixed across time for each item. Nonetheless, MRPT presents additional empirical challenges because sampling error or mismeasured timing of price changes are much more important for MRPT than they are for LRPT. We address these measurement error issues explicitly in both our empirical and modeling sections.

Manual for a much more detailed description (U.S. Department of Labor, 2005).

The target universe of the price index consists of all items purchased from abroad by U.S. residents (imports). An "item" in the data set is defined as a unique combination of a firm, a product and the country from which a product is shipped. An example of the type of item in our data is “Lot # 12345, Brand X Black Mary Jane, Quick On/Quick Off Mary Jane, for girls, ankle height upper, TPR synthetic outsole, fabric insole, Tricot Lining, PU uppers, Velcro Strap.”[7]

Price data are collected monthly for approximately 10,000 imported items. The BLS collects "free on board" (fob) prices at the foreign port of exportation before insurance, freight or duty are added, and almost 90% of U.S. imports have a reported price in dollars.

The BLS collects prices monthly using voluntary confidential surveys, which are usually conducted by mail. Respondents are asked for prices of actual transactions that occur as close as possible to the first day of the month. Typically a company specifies if a price has been contracted and the period for which it is contracted, including the months in which actual trade will take place. For the periods when the price is contracted, the BLS will use the contracted price without contacting the firm directly and enters a flag for whether the good is traded or not in those months.[8]

As with all surveys, there are some concerns about data quality. However, there are many reasons to believe that reporting is accurate. First, the BLS is very concerned with data quality, so in the first step of data collection, the BLS negotiates with the company over the number of price quotes reported so that the company is not overburdened. The BLS also contacts a respondent if the reported price has not changed or the item has not traded for 12 months, which helps reduce misreporting. Second, Gopinath and Rigobon (2008) uses the Anthrax scare of 2001, which forced the IPP to conduct interviews by phone, as a natural experiment. They found almost no difference in reported price setting around these months, which helps reduce concerns about misreporting.

Nonetheless, in the appendix we explore the robustness of our quantitative results to four types of measurement error: sampling error in the price collection process, errors in reporting the correct size of the price change, unreported price changes, and variation in shipping lags of goods. We find that all of our conclusions are robust to various assumptions about the magnitude of these errors.

We focus on a subset of the data that satisfies the following three criteria: 1) We restrict attention to market transactions and exclude intrafirm transactions, as we are interested in price-setting driven by market forces.[9] 2) We require that a good have at least one price adjustment during its life.[10] This is because the goal of the analysis is to relate the standard deviation of price changes to an item’s pass-through and this requires observing at least one price change. This is the same sample

[7] This example is taken from Gopinath and Rigabon (2008).

[8] According to Gopinath and Rigobon (2008), the BLS contacted 87% of the items at least once every 3 months, with 45% of the items contacted on a monthly basis. 100% of the items are contacted at least once a year.

[9] Neiman (2010) shows that pass-through depends on whether transactions take place within or between firms.

[10] Items with one price change necessarily have a population standard deviation of zero. If we generalize to the broader population of items from which the BLS is sampling and use the sample standard deviation, dividing by (#price changes-1) restricts the sample to items with at least 2 price changes. In later results we show that all results are robust to this restriction as well as restricting to items with many price changes. This choice of sample vs. population standard deviation makes no substantive difference for our item-level conclusions and is irrelevant for our month-level dispersion results where we have thousands of price changes per month. We prefer a consistent sample across both month and item specifications and focus on the broadest possible sample in our benchmark results.

restriction used by Gopinath and Itskhoki (2010) in their study of frequency and exchange rate pass-through. 3) We restrict attention to imports whose prices which are invoiced in dollars rather than in foreign currency. We use data from all countries and all products, however we exclude commodities. We restrict attention to dollar-priced items, so as to focus on the relationship between dispersion and pass-through after removing variation due to currency choice. Gopinath et al. (2010) has shown large differences in pass-through across goods invoiced in different currencies, but the vast majority of products in the database are invoiced in dollars rather than foreign currency. Our benchmark results include all countries and all products excluding commodities so as to include the broadest possible sample for which there is still some pricing power. Throughout the paper and appendices, we show that all of our results are robust to a plethora of alternative sample selection criteria.

2.2 Baseline Dispersion Results

2.2.1 Measuring Dispersion and Pass-through

We measure price change dispersion using two distinct but related empirical objects. First, we construct a measure of "item-level" dispersion. For each item j we define item-level dispersion as $DI_j = disp(\Delta p_{i,t}|i = j)$. That is, we calculate the dispersion of all non-zero price changes for item j across time. Since individual items typically have a small number of price changes, we measure item-level dispersion using the standard deviation of that item's price changes.

The second measure of dispersion we construct is "month-level" dispersion. For each month k we define month level-dispersion as $DM_k = disp(\Delta p_{i,t}|t = k)$. To calculate month-level dispersion, we fix a particular month and then calculate the dispersion of price changes across all items in that month.[11] Since there are thousands of price changes each month, we can calculate various different measures of dispersion including the standard deviation and interquartile range of price changes.

Summarizing our two measures of dispersion, "item-level" dispersion is calculated using a single item but all time-periods while "month-level" dispersion is calculated using all items but a single time-period. Since item-level dispersion varies across items rather than time, we refer to "cross-sectional" differences in item-level dispersion. Similarly, since month-level dispersion varies across time-periods rather than items, we refer to "time-series" variation in month-level dispersion.

There is a very large amount of variation in both item-level and month-level dispersion in our dataset. For example, item-level dispersion increases from 1.8% to 23.1% when moving from the first quintile of DI_j to the fifth quintile. Similarly, there is a lot of variation in month-level dispersion across time. Figure 1 shows that the interquartile range of price changes varies from around 8.5% in the late 1990s to 12% during the 2001 Recession, an increase of 40%. In the Great Recession, month-level dispersion nearly doubles.

We measure exchange rate pass-through using micro price data in a standard way. In particular, we focus on what Gopinath and Itskhoki (2010) calls medium-run pass-through (MRPT), which measures the fraction of exchange rate movements passed through into an item's price after one price adjustment. This is the relevant pass-through concept for measuring price flexibility at the

[11] Similar results obtain if we calculate month-level dispersion only within sectors.

Figure 1: Cross-Sectional IQR of Price Changes Across Time

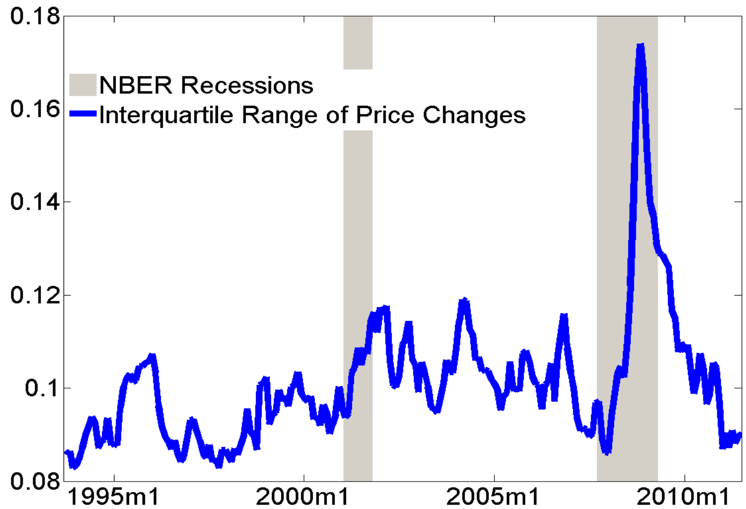

business cycle frequencies relevant for studying countercyclical dispersion. (See footnote 6 for additional discussion). Specifically, we estimate the following regression on adjusting prices:

$$\Delta p_{i,t} = \beta \Delta e_t + Z'_{i,t} \quad + \epsilon_{i,t} \tag{1}$$

Here, $\Delta p_{i,t}$ is item i's log price change, Δe_t is the cumulative change in the bilateral exchange rate since item i's last price change, and $Z'_{i,t}$ is a vector of item and country level controls.[12] We estimate this regression with country and sector fixed effects.[13] The coefficient β measures the fraction of cumulated exchange rate movements "passed-through" to an item's price when adjusting.

We estimate (1) conditional on price adjustment so that our pass-through estimates are not contaminated by nominal rigidities. If we did not condition on adjustment, pass-through would be very close to zero and dominated by variation in the frequency of adjustment. More importantly, conditioning on adjustment allows us to identify the presence of strategic complementarities since they affect how firms want to respond to shocks in the absence of nominal rigidities.

The results from estimating (1) for all price changes in our sample are shown in Table 1. Consistent with the previous literature, we find that average MRPT for dollar denominated items is low. When a price changes, it passes through only 0.14% of a 1% change in the nominal exchange rate.[14]

[12] As usual, there are some concerns about interpreting exchange rate movements as exogenous, which is one reason for including controls for macro conditions. In addition, we are mainly interested in the relative ranking of pass-through across firms and time-periods rather than the absolute level, so endogeneity is less of a concern. Finally, our monthly data means we are identifying off of high frequency variation in exchange rate movements, which are hard to relate to anything observable.

[13] The sector fixed effects are at the primary strata lower (PSL) level, defined by the BLS as either the 2 or 4-digit harmonized tariff code. The other baseline controls are U.S. GDP and CPI and foreign country CPI.

[14] Existing papers typically find pass-through coefficients closer to 0.24. Our slightly lower number is due to the use of bilateral exchange rates, all countries rather than OECD countries, and the use of a moderately longer sample. Using trade-weighted currencies and OECD countries increases MRPT to close to 0.3.

2.2.2 Item-Level Dispersion Results

In this section, we document empirically that there is a strong relationship between medium-run pass-through and item-level price change dispersion.

Figure 2: Medium-run pass-through across item-level XSD quintiles

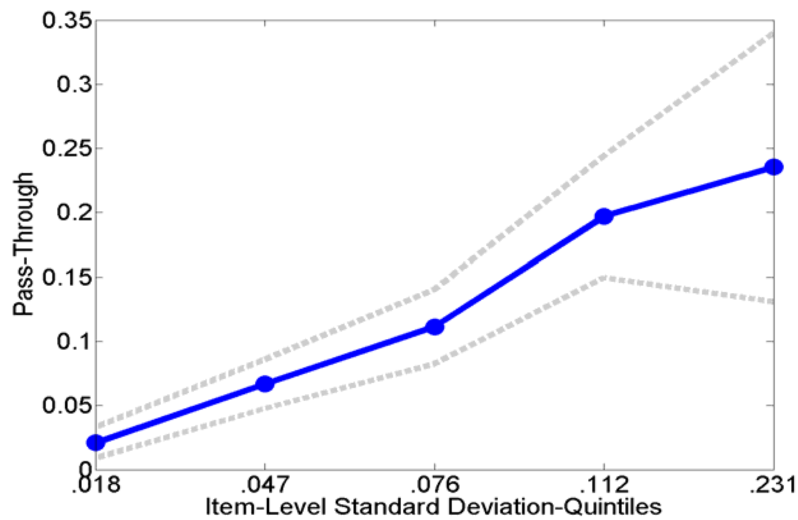

Let $XSD_i = std(\Delta p_{i,t})$ be the standard deviation of item i's price changes (conditional on adjusting). To explore the relationship between MRPT and item-level dispersion, we split our sample into XSD_i quintiles and estimate equation (1) separately for each quintile. Figure 2 shows the baseline results with 95% confidence bands. Average pass-through increases from 2% in the lowest XSD_i quintile to close to 25% for the highest quintile, an increase that is both economically and statistically significant. While we only show this baseline specification for a very broad set of countries and products and it includes no additional controls, in the following sections and appendices we show that this result is extremely robust and is not driven by other item-level features like the frequency of adjustment or degree of product differentiation.

2.2.3 Month-Level Dispersion Results

We previously documented in figure 1 that there is large variation across time in price change dispersion. These changes in dispersion could arise because the size of shocks hitting firms changes over time or because firms responsiveness to shocks varies over time. It is impossible to tell from figure 1 which story is correct. We now take the next step in our argument and show that time periods characterized by greater price change dispersion also exhibit greater exchange rate pass-through. In subsequent sections we use this moment to discriminate between these two theories.

To test for a time-series relationship between price change dispersion and MRPT, we begin by calculating the cross-sectional interquartile range of price changes for each month in our sample. Then, just as we did for the item-level dispersion results, we sort our sample into quintiles by this month-level dispersion and calculate separate pass-through regressions in each quintile.

Figure 3: Medium-run pass-through across month-level IQR Quintiles

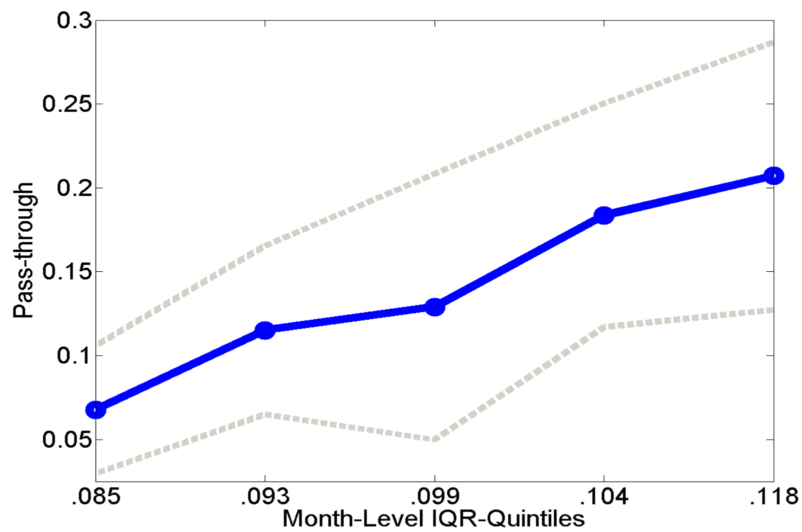

Figure 3 shows that pass-through more than triples from the lowest quintile of month-level dispersion to the highest quintile. This increase in pass-through is highly significant and the absolute increase across quintiles is quite similar to that observed for item-level dispersion.[15] We assess this relationship in more detail in the Appendix 1 and show that this same result obtains for various alternative measures of month-level dispersion, including the cross-sectional standard deviation of price changes as well as census level measures of dispersion computed in Bloom et al. (2012). If we split the sample into deciles, we find even bigger variation across time, with pass-through in the highest dispersion months approaching 50%. In Section 3 we return to a detailed discussion of time-series variation in pass-through.

2.3 Robustness of Item-Level Relationship

Figure 2 shows a strong positive relationship between item-level dispersion and MRPT, but it does not control for any other characteristics that vary across items. We now show that our item-level result is robust to a host of alternative explanations and controls. This is important, because when we move from empirics to theory, we will argue that our result reflects heterogeneity in fundamental economic primitives. In particular, we will argue that our result is driven by heterogeneity in the importance of strategic complementarities which generate variation in responsiveness, rather than by heterogeneity in volatility. In order to reach this strong conclusion, it is important to show that our result is not driven by other observable differences across items.

[15] Standard errors are larger than for the item-level relationships because our panel has a very large number of items but a much smaller number of time-periods.

2.3.1 Ruling Out Mechanical Explanations

Looking at the MRPT specification in (1), one might be concerned that the positive relationship between pass-through and item-level dispersion reflects a mechanical relationship. In particular, note that the MRPT regression coefficient is equal to

$$\widehat{\beta} = \frac{cov(\Delta p_{i,t}, \Delta e_t)}{var(\Delta e_t)} = \beta + \frac{cov(\epsilon_{i,t}, \Delta e_t)}{var(\Delta e_t)}.$$

where β is the desired response of prices to exchange rate movements. Taking the variance of both sides of (1) gives an expression for the variance of price changes across time[16]:

$$var\,(\Delta p_{i,t}) = \beta^2 var\,(\Delta e_t) + var\,(\epsilon_{i,t}) + 2\beta cov(\epsilon_{i,t}, \Delta e_t). \tag{2}$$

Thus, if there is heterogeneity in β or $cov(\epsilon_{i,t}, \Delta e_t)$, then this can generate a positive relationship between $\widehat{\beta}$ and $var\,(\Delta p_{i,t})$.[17] However, in a flexible price environment, it is straightforward to show that these effects are quantitatively irrelevant.[18] The intuition for why these mechanical effects cannot explain our empirical result is straightforward: generating the variance of price changes observed in the data requires the variance of idiosyncratic shocks to be two orders of magnitude larger than the variance of exchange rate shocks. That is, $var\,(\epsilon_{i,t}) >> var\,(\Delta e_t)$. This implies that changing only β or $cov(\epsilon_{i,t}, \Delta e_t)$ has negligible effects on $var\,(\Delta p_{i,t})$.

While it is straightforward to rule out the quantitative importance of these effects in a flexible price environment, we also argue that these effects are unimportant even in a more empirically realistic setting. In particular, when we move to a structural interpretation of our results, we show that the same flex price intuition applies in a richer quantitative model with price-stickiness. Importantly, in our indirect inference estimation, we simulate the model with sample sizes and stickiness as in the BLS data and show that the above mechanical relationship cannot explain our results. This indirect inference procedure also rules out concerns that our results might be driven by mechanical statistical confounders, since these should also arise in our simulated data.

Finally, it is also important to note that the mechanical relationship is irrelevant for our month-level results in the flexible price environment if exchange rate movements are common to items within a given month, which is the case in our model.[19] That is in month-level regressions with

[16] Here we suppress the $Z'_{i,t}$ related terms. These terms have little quantitative importance empirically and so do not change anything about the qualitative conclusions we derive, but removing them simplifies the exposition.

[17] β might differ due to heterogeneous sensitivity of costs to exchange rates. Even with no selection effects, $cov(\epsilon_{i,t}, \Delta e_t)$ might vary since in small samples the sampling covariance can differ from the population covariance.

[18] Formally, we have data on $var\,(\Delta p_{i,t})$, $var\,(\Delta e_t)$, and β, so we can use equation 2 to measure $var\,(\epsilon_{i,t})$ under the null hypothesis that our fact is explained by heterogeneity in β. Substituting empirical results from the data and using equation 2 yields $\beta = 0.144$, $var\,(\Delta e_t)$ =6.25e-4 and $var\,(\epsilon_{i,t})$ =1.83e-2. Using these values for $var\,(\Delta e_t)$ and $var\,(\epsilon_{i,t})$, varying β from 0.021 to 0.235 (as in the data) generates less than 0.1% of observed variation in dispersion.

[19] In our model we make the simplifying assumption that exchange rates are common across items. This is not an important restriction in terms of matching our empirical evidence, because in the empirical appendix we show that all our results hold within countries with a common currency, hold using a broad trade-weighted exchange rate common to all items, and are not driven by differences in exchange rate volatility across currencies. Thus, heterogeneity in exchange rates adds modeling complexity but does not explain our empirical results.

fully flexible prices and a shared exchange rate, $var\left(\Delta e_t\right)$ and $cov(\epsilon_{i,t}, \Delta e_t)$ are both identically zero so that the problematic terms are dropped from (2). We again verify the irrelevance of this channel in an environment with price-stickiness and realistic sample sizes in our indirect inference procedure.

2.3.2 Item-Level Dispersion Interactions

Is the positive relationship between pass-through and dispersion driven by other observables? To explore this, we run regressions on continuous measures of price change dispersion instead of the previous binned regressions. These more structured specifications allow us to include a variety of additional controls. Let the change in an item's price be given by:

$$\Delta p_{i,t} = \beta^{avg} \Delta e_t + \beta^{Vol}\left(Vol_i \times \Delta e_t\right) + \delta Vol_i + Z'_{i,t} \quad + \epsilon_{i,t} \tag{3}$$

The coefficient β^{avg} captures average pass-through in the sample and β^{Vol} gives the effect of item-level price change volatility on MRPT.[20] Results are shown in Table 2.

The first row shows results for our baseline sample. Average exchange rate pass-through is 14%. β^{Vol} is significantly greater than zero, which means that items with higher price dispersion have higher MRPT. The price dispersion effect is economically meaningful: a one standard deviation increase in price dispersion implies a 37% (0.05/0.14) increase in MRPT relative to average.

In the second row, we explore whether our results are driven by difference in the frequency of adjustment. It is well-known that the frequency of adjustment varies dramatically across items and sectors. This means that some items in our MRPT regression will be adjusting after long periods of stickiness and will have built up large cumulative exchange rate pressure while other items will have adjusted recently and will have little need to respond to respond to exchange rate movements. If the frequency of adjustment is correlated with item-level price change volatility, then our volatility relationship might be proxying for differences in the frequency of adjustment. Row 2 shows that this is not the case. Specifically, we estimate the following extension of (3):

$$\Delta p_{i,t} = \beta^{avg} \Delta e_t + \beta^{Vol}\left(Vol_i \times \Delta e_t\right) + \delta^{Vol} Vol_i + \beta^{freq}\left(freq_i \times \Delta e_t\right) + \delta^{freq} freq_i + Z'_{i,t} \quad + \epsilon_{i,t}$$

Controlling for an item's frequency of adjustment has no effect on our estimate of β^{Vol}. In Appendix 1 we provide additional empirical evidence on this point, and we also address it directly in our structural model, as our estimation directly targets the joint relationship between frequency, price change dispersion and pass-through.

In addition to showing that controlling for frequency does not change our results, Table 2 also shows that our results are not driven by a particular set of products or countries. Restricting to OECD countries or manufactured items only strengthens the importance of item-level dispersion. In both subsamples, the price dispersion effect is economically and statistically significant.

[20] In all specifications, the measure of item level price dispersion is the standard deviation of price changes (XSD) and robust standard errors are clustered by country and primary stratum lower (4 digit import type) pair. Measures of volatility in this and the following specifications are standardized to ease interpretation. This implies that β^{avg} is equal to the level of pass-through for an item with average dispersion, which is a natural measure of average pass-through.

For brevity, we leave a variety of other item-level robustness checks for the appendix. In those results, we show that the positive relationship between MRPT and item-level dispersion holds for individual countries, for various product subsets, using broad rather than bilateral exchange rates, is robust to controls for outliers and is not driven by small sample issues.

2.4 Robustness of Month-Level Relationship

In this section, we argue that the time-series relationship between month-level dispersion and pass-through is not driven by various other observables or confounding shocks. Since this time-series relationship is of direct consequence for policy making, it is particularly important to show that it is a deep relationship in the data rather than a spurious relationship driven by failing to control for some other observable. However, in contrast to the item-level data, which contain a limited set of covariates, we are able to show robustness results for a rich set of time-varying macro controls.

In order to control for other time-varying features of the data, we estimate an interaction specification between MRPT and month-level dispersion. More specifically, we run the regression

$$\Delta p_{i,t} = \beta^{avg}\Delta e_t + \beta^{IQR} IQR_t \times \Delta e_t + \lambda IQR_t + Z'_{i,t} \quad + \epsilon_{i,t} \tag{4}$$

where IQR_t is the interquartile range of all (non-zero) price changes in month t and $Z'_{i,t}$ is the same vector of controls as in the cross-sectional regressions.[21] Table 3 shows that increasing IQR by one-standard deviation increases pass-through by 43% (0.06/0.14). This positive relationship is highly significant, with a t-statistic of 7.01. We find similar effects when using the cross-sectional standard deviation instead of the interquartile range, as well as when restricting to OECD countries and manufactured items. Importantly, we again find that controlling for the frequency of adjustment has no effect on our estimates of β^{IQR}. That is, the effects of month-level dispersion on pass-through are unrelated to movements in the frequency of adjustment across time.

Using specification (4) allows us to control for other things that might vary across time, and Table 4 shows a number of these results. An existing literature has debated whether there have been secular declines in aggregate pass-through over time. For example, Marazzi et al. (2005) argues that aggregate measures of pass-through have declined, but Hellerstein et al. (2006) show this is largely driven by commodities. Since we exclude commodities from our analysis, we find no trends in pass-through across time. Even more importantly, these studies use only aggregate data and so are not directly relevant for our results. This is because aggregate pass-through measures $frequency\times$ MRPT rather than just MRPT.[22] Since the frequency of adjustment is low, aggregate pass-through has little relationship to MRPT, our empirical object of interest.[23] Nevertheless, if there were trends or seasonality in both MRPT and price change dispersion, our time-series results

[21] As in the cross-sectional regression we standardize all dispersion numbers to ease the interpretation of our results.

[22] There are other more subtle differences, since we also use bilateral exchange rates and have additional country and sector controls in our MRPT regressions.

[23] See Gopinath and Rigobon (2008) for additional evidence that aggregate pass-through is uninformative for MRPT. As argued in the introduction, MRPT is the relevant object for measuring the importance of strategic-complementarities and identifying time-varying responsiveness.

could be spurious. We directly address this concern in Table 4 by re-estimating regression 4 with a linear time-trend and month dummies. These controls do not significantly affect our results.

Table 4 also shows that our result is robust to controlling for the frequency of product substitution[24] and time-series volatility of exchange rates. Since price change dispersion is countercyclical, we also check that the positive relationship between dispersion and pass-through is not just proxying for a recession-pass-through relationship. Again, we find it is not: the final row of Table 4 shows that controlling for the state of the business cycle does not significantly affect our main result.[25]

Finally, we want to make sure that our results are not just being driven by data from the Great Recession. We already saw in figure 1 that price change dispersion in the Great Recession is a significant outlier. Since our identification strategy relies heavily on the positive correlation between pass-through and dispersion, it is important that this correlation is a general feature of the data and not just an artifact of this one large recession. Are our time-series results then driven by this particularly dramatic period? Table 5 shows that they are not. Restricting our analysis to data prior to 2008 mildly reduces the relationship between pass-through and dispersion, but it remains highly quantitatively and statistically significant.[26] In particular, a one-standard deviation increase in IQR increases pass-through by 30% (0.038/0.125). This is smaller than the 42% increase observed when using the full sample, but it is still very substantial. Thus, neither the economic nor the statistical significance of our relationship is an artifact of the Great Recession, so we believe our results are informative for understanding time-varying dispersion more generally.

2.5 Controls for Composition

While our empirical result holds under a variety of controls, these may not pick up compositional changes across time in our sample. In this section, we address several potential composition concerns. For brevity, we leave the actual results to our empirical appendix and summarize them here.

First, is our item-level fact actually distinct from our month-level dispersion fact? We do not have a balanced panel due to sample rotation, so it is possible that the high dispersion time-periods in our data are just times when the sample contains items with unusually high dispersion. In Tables A3 and A4, we document that our two facts are indeed distinct by showing that in a joint-regression, cross-item and cross-month dispersion both independently lead to increases in pass-through. We also show that even when restricted to a balanced panel there is a positive relationship between month-level dispersion and MRPT. This means that pass-through for the same products rises with

[24] Nakamura and Steinsson (2012) argue that missing price changes at the time of product substitution leads to downward bias in pass-through. Since our measure of MRPT conditions on price changes, the presence of product substitution is not directly relevant for our results. Nevertheless, product substitution rises mildly with the dispersion of price changes. This means that the pass-through increase we document may understate the true increase in aggregate pass-through: accounting for product substitution would, if anything, amplify our results.

[25] We find similar results using alternative samples and dispersion measures. In particular, using the standard deviation of price changes instead of the interquartile range and using different country and product mixes does not change the conclusion that these additional controls make little difference. In the interest of brevity we do not report these results, but all results for different subsamples are available upon request.

[26] We only have two recessions in our sample, so it is not obvious that excluding data from after 2008 is desirable. Nevertheless, the next section shows that the moderately lower variation in implied pass-through in the earlier period is driven by lower variation in dispersion rather than a reduced relationship between dispersion and pass-through.

month-level dispersion so our results cannot be explained by a time-varying product mix.[27]

In general, the positive relationship we observe between pass-through and month-level dispersion could be driven by movements in dispersion across sectors or within a sector. Table A5 shows that our month level relationship is mostly driven by increases in within sector dispersion rather than cross-sector dispersion. That is, increases in pass-through are more closely associated with times when price change dispersion within a sector rises rather than times when sectors become more different from each other. Thus, we focus on within sector variation in our quantitative modeling.

2.6 Exchange Rate Appreciations Vs. Depreciations

The 2008 recession was also characterized by an appreciation of the U.S. dollar against most major currencies. However, our pass-through results are not sensitive to the sign of exchange rate movements. We find that both our month-level and our item-level dispersion MRPT relationships remain highly significant even when restricting our regressions solely to price changes where Δe_t is always positive or always negative.[28] Thus, our results cannot be explained by changes across time in whether the dollar is appreciating or depreciating.

2.7 Measurement Error Concerns

2.7.1 Alternative Pass-through Specifications

All results thus far have relied on MRPT specifications as in (1). This specification directly measures the extent to which exchange rate movements are passed into current prices, so it provides a snapshot of price-flexibility at a moment in time. This in turn is extremely informative for differentiating changes in firms' desired responses from alternative shocks. Nevertheless, there are two potential concerns with this specification. First, if the timing of price changes is mismeasured, then MRPT suffers from attenuation bias.[29] Second, items may differ in the number of price changes required to fully capture pass-through. If that case, estimating pass-through conditional on a single price change may provide a distorted picture of cross-item price flexibility.

With these measurement concerns in mind, we estimate several alternative pass-through specifications. First, we use a "fixed horizon" pass-through specification where we calculate $\Delta p_{i,t}^K = p_{i,t+K} - p_{i,t}$ and $\Delta e_t^K = e_{t+K} - e_t$ for fixed pass-through horizons K. We then rerun our pass-through regressions using this new measure of price and exchange rate changes. Crucially, this alternative specification does not condition on price adjustment, so individual items may have between 0 and K price changes occurring between t and $t+K$. This reflects the full extent of an item's pass-through over a fixed horizon, whether it occurs through one or many price changes. This specification shares many of the attractive features of life-long pass-through used in Gopinath and Itskhoki (2010) but

[27] Dispersion and pass-through vary more across time in the balanced panel than in our baseline unbalanced panel.

[28] For brevity we do not report these results, but they are available upon request.

[29] Mismatched timing can occur when there is a lag between the time an item is purchased and the time it ships. In the appendix we show that controlling for shipping method does not alter our conclusions. We also simulate various additional sources of measurement error and show they cannot explain our result.

allows us to calculate time-series variation in pass-through.[30] In addition, we also run our baseline MRPT regression allowing for lagged exchange rate movements to matter for current price changes.[31]

Table 6 provides results for these alternative specifications. In all cases, increases in item-level or month-level dispersion leads to economically large and statistically significant increases in pass-through.[32] Thus, our results are insensitive to particular measures of exchange rate pass-through.[33]

2.7.2 Small Samples

Given the limited time span of our panel data and low average frequency of price adjustment, it is important to verify that the relationship between pass-through and dispersion is not driven by small sample issues. In the previous section, we argued that small samples could lead to biased pass-through estimates by affecting $cov(\epsilon_{i,t}, \Delta e_t)$, since the sample covariance can differ from zero even if the population covariance does not. Our structural modeling addresses this concern head on by using an indirect inference procedure with sampling procedures that mimic the BLS. If our empirical results were driven by censoring bias or selection then they should also show up in simulated results.

Nevertheless, we also try to address these small sample concerns directly in our empirical results. In Appendix 1, we show that our item-level results are robust to restricting the sample to items which have at least 3 or at least 5 price changes. We also run placebo regressions to see if our results are driven by spurious small sample issues. In particular, we substitute a count of an item's price changes or its price observations in place of XSD. These placebo regressions show that our results are not driven by a correlation between measured dispersion and item sample sizes.

3 Time-Variation in Pass-Through

In the previous section we documented a robust link between exchange rate pass-through and microeconomic price change dispersion. Before demonstrating how our empirical fact can help discriminate between time-varying volatility and time-varying responsiveness as sources of time-varying dispersion, we first argue that our fact is also interesting per se. In particular, we show that an implication of the positive correlation between dispersion and pass-through is that there is large variation in exchange rate pass-through at business cycle frequencies. That is, pass-through is not a single number; it varies significantly over time and is high when dispersion is high. In this sense our empirical results provide model-free evidence that looking at the microeconomic distribution of price changes is crucial for predicting inflation dynamics at a moment in time.

[30] This fixed horizon pass-through specification also provides yet another piece of evidence that our empirical relationship is not driven by heterogeneity in the frequency of adjustment, since we no longer condition on adjustment.

[31] That is, we estimate: $\Delta p_{i,t} = \beta_1^{avg}\Delta e_t + \beta_1^{Vol}(XSD_i \times \Delta e_t) + \beta_1^{IQR} IQR_t \times \Delta e_t + \beta_2^{ave}\Delta e_{t-1} + \beta_2^{Vol}(XSD_i \times \Delta e_{t-1}) + \beta_2^{IQR} IQR_t \times \Delta e_{t-1} + \delta XSD_i + \lambda IQR_t + Z'_{i,t} + \epsilon_{i,t}$

[32] Unsurprisingly, there is a very significant increase in the level of pass β^{avg} with the fixed horizon over which we measure pass-through. While theory has no strong prediction for how β^{vol} should vary with the pass-through horizon, there appears to be a modestly significant increase in β^{XSD} but no change in β^{IQR} with this horizon. The main takeaway is that all measures of β^{vol} are significantly positive across all specifications.

[33] We also find that life-long pass-through is increasing in item-level dispersion. Since life-long pass-through is only measured once for each item we cannot measure time-variation, so we do not report these results.

The results from the previous section allow us to construct implied time-series for exchange rate pass-through by multiplying observed variables by their estimated effects on exchange rate pass-through. For example, using regression specification (4) we estimate pass-through in each period t by computing $\widehat{MRPT}_t = \widehat{\beta}^{avg} + \widehat{\beta}^{IQR} IQR_t$.

Figure 4: Level of Exchange Rate Pass-through Across Time (Parametric Specifications)

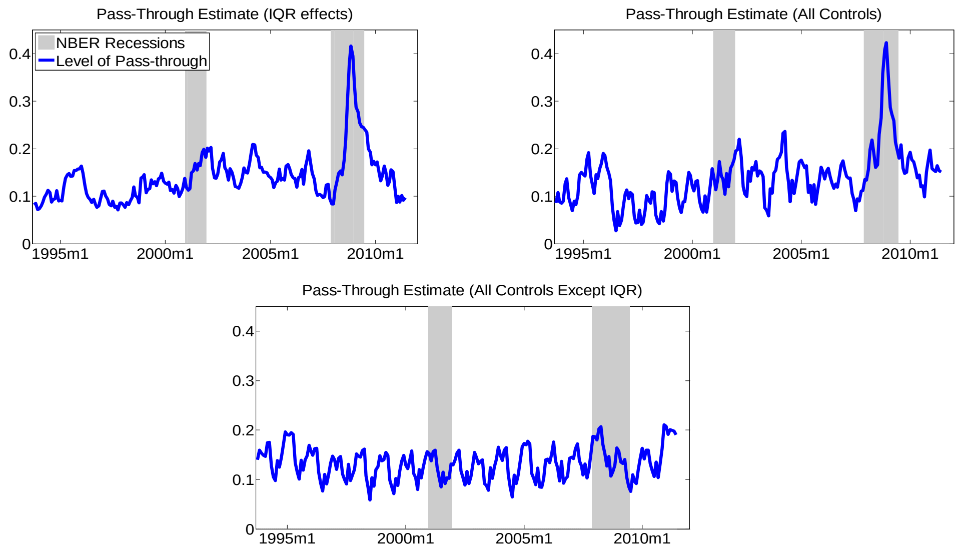

The identifying assumption in this specification is that the only thing that varies across time that affects exchange rate pass-through is IQR. The left hand panel of Figure 4 shows the resulting estimates for exchange rate pass-through under this specification.[34] During the height of the Great Recession, this estimate of exchange rate pass-through rises to 44% relative to a low of approximately 7% during the late-1990s. The assumption that time-variation in exchange rate pass-through is solely driven by variation in IQR is very strong but can be easily relaxed. In the right hand panel of Figure 4 we allow pass-through to vary with IQR, the frequency of adjustment, the frequency of product substitution, seasonal month dummies, and real GDP growth.

Allowing for these additional interactions does not change the conclusion that pass-through rose markedly during the 2008 recession. The main difference relative to the specification with only IQR is a large seasonal component. This can be seen most clearly in the bottom panel of Figure 4, which shows pass-through estimates for a specification with all controls except for IQR. Essentially all the variation in pass-through at business cycle frequencies is captured by time-series variation in price change dispersion. Interestingly, there is some seasonality in pass-through, from a high of approximately 0.16 in December to a low of approximately 0.09 in June. Understanding these seasonal patterns is an interesting topic for future work, but the bottom line is that for understanding business cycle variation, looking at price change dispersion appears essential.[35]

[34] Note that by construction, this pass-through series is perfectly correlated with the dispersion time-series shown in Figure 1, so the new object of interest is the implied scale of pass-through fluctuations.

[35] Seasonality is unlikely to be explained purely by a spike in the frequency of adjustment at the end of the year. This

Figure 5: Pass-through Estimate (Non-Parametric 12-month Rolling Window)

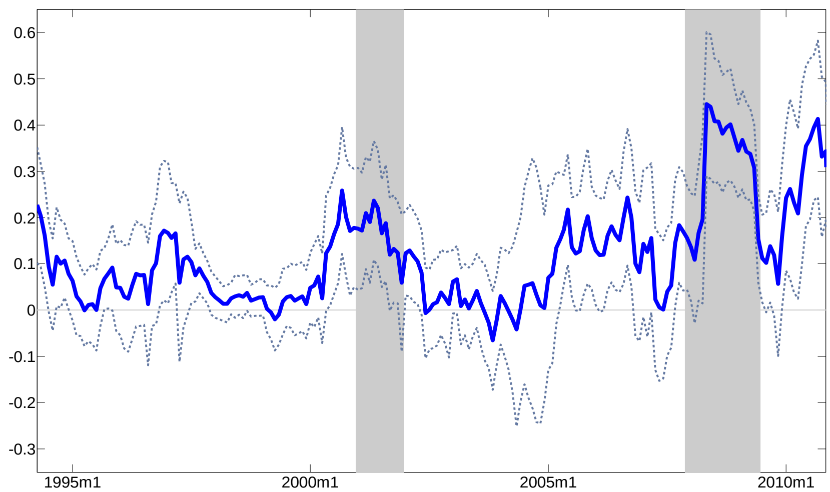

While the above results show that pass-through varies across time in a specification with a variety of controls, there is always concern that omitted variables might undo this time-series variation. That is, there may be additional variables we are not controlling for that affect pass-through and would undo the time-series variation we have found. We can assess this concern by allowing pass-through to vary across time non-parametrically. Ideally, we could re-estimate the baseline pass-through regression (1) with a full set of month dummies. However, small sample sizes make such regressions infeasible. Instead, we estimate the baseline regression using a rolling 12-month window. That is, our estimate of pass-through for period t is given by re-running regression 1 using only price changes occurring in a window 6 months before and after period t:

$$\Delta p_{i,\tau} = \beta_t \Delta e_\tau + Z'_{i,\tau} \quad + \epsilon_{i,\tau} \mid t - 6 \leq \tau \leq t + 6.$$

This allows us to construct a monthly measure of β_t that varies fully non-parametrically across time. Figure 5 shows the resulting estimates together with 90% confidence intervals. Overall the results are quite similar to the parametric specification, and again there is strong variation in pass-through at business-cycle frequencies. In particular, notice that pass-through is significantly larger than it was in the late 1990s both in the 2001 and 2008 Recessions. Running annual pass-through or 6-month pass-through regressions instead of using overlapping rolling windows produces very similar results.[36] This specification shows that being completely agnostic about what drives pass-through movements across time delivers quite similar results to our benchmark specifications.

Overall our results show that exchange rate pass-through varies dramatically across time, with

is because our measure of pass-through conditions on adjustment, so we are finding variation in how much adjusting prices respond to exchange rate movements over the season that are unlikely to be explained purely by frequency.

[36]Quarterly results (available from authors on request) are also similar although small sample sizes mean that the standard errors become extremely large and estimates are quite noisy.

microeconomic price change dispersion. This means that estimating average pass-through regressions without looking at micro data induces a significant time-varying bias, with pass-through substantially understated during periods of microeconomic churning. While a large literature tries to understand average pass-through and its implications for the nominal transmission mechanism, the above evidence shows that pass-through is not a single number and that concentrating on average pass-through may be misleading for how prices will respond to nominal shocks at a moment in time.

Beyond providing direct empirical evidence that the distribution of price changes matters for predicting pass-through, we now show that our empirical results provide additional identification that is useful for understanding the nature of heterogeneity and aggregate shocks in the economy.

4 Basic theoretical framework

4.1 Flexible price model

In this section we lay out a simple framework following Burstein and Gopinath (2013) that shows how economic primitives shape the relationship between exchange rate pass-through and price change dispersion. In this section, we do so in the simplest possible setting in order to build intuition, by assuming flexible prices, no aggregate shocks and no equilibrium effects. This allows us to develop simple formulas relating endogenous responsiveness and fundamental volatility to pass-through and the variance of price changes.[37] In the quantitative section which follows, these assumptions are relaxed but the intuition is similar.

Consider the problem of a foreign firm selling items to U.S. importers. The firm has perfectly flexible prices, set in dollars. The optimal flexible price (in logs) of item i at the border is the sum of the gross markup (μ_i) and dollar marginal cost ($mc_i(e, \eta_i)$) which depends on both the exchange rate (e) and an item-specific component orthogonal to the exchange rate (η_i):

$$p_i = \mu_i + mc_i(e, \eta_i). \tag{5}$$

Taking the total derivative of equation (5) gives:

$$\Delta p_i = \square\square_i(\Delta p_i \square \Delta p) + \alpha_i \Delta e + \epsilon_i \tag{6}$$

where $\square_i \equiv \square \frac{\partial \mu_i}{\partial(\Delta p_i \square \Delta p)}$ is the elasticity of a firm's optimal markup with respect to its relative price. We refer to this parameter as markup "responsiveness". It captures the classic pricing to market channel of Dornbusch (1987) and Krugman (1987), where firms may adjust markups in response to cost shocks, leading to incomplete pass-through. This channel implies a negative relationship between markups and relative prices, $p_i \square p$, which Burstein and Gopinath (2013) show is a robust implication of various strategic complementarities that generate incomplete pass-through. $\alpha_i \equiv \frac{\partial mc_i}{\partial e}$ is the partial elasticity of the dollar marginal cost to the exchange rate, e. We refer to this as the "import intensity" channel. Finally, $\epsilon_i = \Delta \eta_i$ captures the innovation of idiosyncratic marginal

[37] In the appendix, we consider a more general model which includes GE effects and scale-dependent marginal cost.

cost.[38] We call changes in the variance of this idiosyncratic component fundamental "volatility". Rearranging this equation gives an explicit expression for the direct effect (that is when $\Delta p = 0$) of a change in the exchange rate on prices at the border:[39]

$$\frac{\Delta p_i}{\Delta e} = \frac{\alpha_i}{1 + \Box_i} \tag{7}$$

The first factor affecting pass-through is the fraction of marginal cost denominated in dollars. If marginal cost is entirely denominated in dollars ($\alpha_i = 0$), then fluctuations in the exchange rate are irrelevant for the foreign firm's optimal dollar price and pass-through is zero. In general, exchange rate pass-through is increasing in import intensity.

The second factor affecting pass-through is the response of the foreign firm's optimal markup to changes in its relative price. If $\Box_i = 0$ (the CES case) the firm's optimal markup does not change as its price deviates from its competitors and pass-through is at its maximum. If $\Box_i > 0$, then as the price of the firm increases relative to its competitors, the elasticity of its demand rises, lowering its optimal markup. Similarly, when the firm's price is relatively low, its optimal markup rises. Thus, if $\Box_i > 0$, the foreign firm will move its price less than one-for-one in response to cost shocks.

Since lowering $\Box_i$ means that firms will be more responsive to all cost shocks, we refer to lowering $\Box_i$ as increasing total "responsiveness". That is, firms with low $\Box_i$ will respond strongly to both idiosyncratic shocks as well as exchange rate shocks. In contrast, firms with high α_i will respond more to exchange rate shocks but not to idiosyncratic cost shocks. We use the term responsiveness to differentiate general cost pass-through from exchange rate specific pass-through.

In addition to its implications for pass-through, we can also use equation (6) to show how α and $\Box$ affect the variance of Δp_i. Solving for Δp_i and computing its variance gives:

$$var(\Delta p_i) = \left(\frac{\alpha_i}{1 + \Box_i}\right)^2 var(\Delta e) + \left(\frac{1}{1 + \Box_i}\right)^2 var(\epsilon_i), \tag{8}$$

where we have used the fact that exchange rate and idiosyncratic shocks are uncorrelated.

Intuitively, the variance of the firm's optimal price is larger if it faces a more volatile exchange rate or idiosyncratic shocks. In addition, using equation (8), it follows that factors that increase exchange rate pass-through ($\alpha_i \uparrow$, $\Box_i \downarrow$) also increase the variance of price changes. Moreover, using equation (8) it can be shown that for empirically relevant values of α_i and $\Box_i$, changing $\Box_i$ has much larger effects on price change variance than changing α_i.[40] The intuition, as previously discussed in Section 2.3.1, is that empirical estimates of $var(\epsilon_i)$ greatly exceed $var(\Delta e)$. In addition, α_i is

[38] Since we do not observe this shock, it is without loss of generality to normalize the price response to η to be one.

[39] We also set the innovation of the idiosyncratic shock to its average value (zero).

[40] More formally, combine the two formulas in elasticity form to get:

$$\frac{\left|\left(\frac{\partial var(\Delta p_i)}{\partial \Box} \frac{\Box}{var(\Delta p_i)}\right)\right|}{\left(\frac{\partial var(\Delta p_i)}{\partial \alpha} \frac{\alpha}{var(\Delta p_i)}\right)} = \frac{\Box}{1 + \Box}\left(1 + \frac{1}{\alpha^2}\frac{var(\Delta \eta_i)}{var(\Delta e_i)}\right)$$

Substituting calibrated values from the modeling section yields a ratio of approximately 200.

typically small. (See Figure 2) This means that the first term in (8) contributes little to the overall variance of price changes, so changing its size also has little effect. In the quantitative modeling section, we show that this simple intuition survives in a realistic model. That is, the mechanical link between heterogeneity in α_i and heterogeneity in $var(\Delta p_i)$ is not empirically important.

4.2 Modeling Price Stickiness

Price stickiness is a pervasive feature of micro price data. For example, Gopinath and Rigobon (2008) find that the median price duration for imports to the U.S. is 10.6 months. More importantly, the price adjustment mechanism can have direct effects on measured pass-through. For example, in menu cost models, where price adjustment is endogenous, conditioning on price adjustment will induce a selection bias in MRPT estimates.[41] In contrast, this bias is absent in Calvo pricing models where price adjustment is exogenous.[42]

To understand how the primitives of a menu cost model can affect measured pass-through, it is useful to examine our baseline MRPT regression shown in equation (1). By definition, the estimated MRPT regression coefficient is equal to:

$$\widehat{\beta} = \frac{cov(\Delta p_{i,t}, \Delta e_t)}{var(\Delta e_t)} = \beta + \underbrace{cov(\epsilon_{i,t}, \Delta e_t)/var(\Delta e_t)}_{\text{selection bias}}$$

where β is the "true" responsiveness of desired prices to exchange rate movements.[43] Menu cost models induce $cov(\epsilon_{i,t}, \Delta e_t) > 0$ for firms that choose to adjust, even if the unconditional covariance is zero. This is because in a menu cost model, firms are more likely to choose to adjust when the idiosyncratic shock and the exchange rate movement reinforce each other. Thus, $cov(\epsilon_{i,t}, \Delta e_t) > 0$, for adjusters. This implies that estimated pass-through conditional on price adjustment, $\widehat{\beta}$, is biased upward relative to true desired pass-through, β.[44]

Higher menu costs lead firms to adjust less often and by larger amounts (which increases the dispersion of price changes) as firms economize on the number of times they adjust prices. Increases in the menu cost lead to a wider range of inaction, which leads the importance of selection effects and $cov(\epsilon_{i,t}, \Delta e_t)$ to increase. This then leads to an increase in measured MRPT.

Conversely, increasing the variance of idiosyncratic cost shocks lowers MRPT because the magnitude of the selection bias is decreasing in the size of these shocks. The intuition is simple: as the size of the idiosyncratic shocks increases, firms are more likely to adjust their prices for purely idiosyncratic reasons, which lowers $cov(\epsilon_{i,t}, \Delta e_t)$, conditional on adjustment. At the same time,

[41] We are not first to notice this bias. See the brief discussion in footnotes 7 and 26 of Gopinath et al. (2010).

[42] Both models suffer from attenuation bias in pass-through driven by censoring at sample rotation, but this will not affect the dispersion of price changes. Furthermore, it is captured explicitly in our model simulation since we rotate our simulated panel as in the BLS.

[43] This underlying β is determined by α and □, as shown in the previous section. It is also declining with price stickiness if exchange rate movements are not permanent, but exchange rates are close to a random walk in the data so that the flexible price expression provides a close approximation even for firms with relatively sticky prices.

[44] It is worth noting that this is only a "bias" if one is interested in measuring desired pass-through in the population. But if one is interested in measuring how much actual prices will respond to exchange rate movements, the relevant object is $\widehat{\beta}$ not β.

larger shocks mean larger price dispersion. Thus variation in the size of idiosyncratic shocks induces a counterfactual negative correlation between MRPT and dispersion and so already suggests that volatility shocks have difficulty replicating empirical facts in the open economy environment.

Reviewing the conclusions from this and the previous section, it follows that heterogeneity in α, □ or in the size of menu costs should generate a positive relationship between measured MRPT and the dispersion of price changes. However, we now show that only the "responsiveness" channel arising from variation in □ is quantitatively successful.

5 Quantitative Model

We now formally assess the theoretical link between price change dispersion and exchange rate pass-through in an estimable quantitative model. The model allows for all the theoretical channels discussed in the previous section and also includes indirect equilibrium effects that the simple model in Section 4.1 ignored. The main model we explore builds heavily on the menu cost model of Gopinath and Itskhoki (2010). This model has been successful at matching a variety of cross-sectional and steady-state empirical facts. We build on it by formally estimating various forms of heterogeneity in the cross-section as well as by adding aggregate shocks to explain our month-level dispersion evidence. We intentionally build on this workhorse model of incomplete pass-through to show that it implies extremely tight links between our empirical facts and underlying economic primitives. The model features heterogeneity in import sensitivity, idiosyncratic volatility, and responsiveness with less than full responsiveness driven by the strategic complementarities which arise under Kimball demand.[45] While we model strategic-complementarities using Kimball demand this is not an important assumption, as other channels such as pricing to market in Atkeson and Burstein (2008) yield similar reduced form implications. We discuss this more fully in Section 5.4.3.

5.1 Model Description and Calibration

5.1.1 Industry Demand Aggregator

The industry is characterized by a continuum of varieties indexed by j. There is a unit measure of domestic varieties and a measure $\omega < 1$ of foreign varieties available for domestic consumption, which captures the idea that not all varieties are traded internationally.

We generate variable markups by utilizing a Kimball (1995) style aggregator:

$$\frac{1}{|\quad|}\int_{\Omega} \Psi\left(\frac{|\quad| C_j}{C}\right) dj = 1 \tag{9}$$

with $\Psi(1) = 1, \Psi'(.) > 0$ and $\Psi''(.) < 0$. C_j is the quantity demanded of variety $j \in$, where is the set of all varieties available domestically. has measure $1 + \omega$. Individual varieties are aggregated into a final consumption good C. This intermediate aggregator contains the CES specification as a

[45] A Calvo model delivers similar conclusions about the importance of responsiveness but fits the data less well. Results are available upon request.

special case. The demand function for C_j implied by equation (9) is:

$$C_j = \varphi\left(D\frac{P_j}{P}\right)\frac{C}{|\Omega|}, \text{ where } \varphi(.) \equiv \Psi'^{-1}(.) \tag{10}$$

Here P_j is the price of variety j, P is the sectoral price index and $D \equiv \left[\int_\Omega \Psi'\left(\frac{|\Omega|C_j}{C}\right)\frac{C_j}{C}dj\right]$. P is defined implicitly by the following equation

$$PC = \int_\Omega P_j C_j dj$$

5.1.2 Firm's problem

Consider the problem of a firm producing variety j. Foreign and domestic firms face symmetric problems and we label foreign variables with asterisks. The firm faces a constant marginal cost:[46]

$$MC_{jt} = \frac{W_t^{1-\alpha}(W_t^*)^\alpha}{A_{jt}}$$

where W_t is the domestic wage and the parameter α is the share of foreign inputs in the firm's cost function. A_{jt} denotes idiosyncratic productivity, which follows an AR(1) in logs:

$$\log(A_{jt}) = \rho_A \log(A_{j,t-1}) + \mu_{jt} \text{ with } \mu_{jt} \sim iid\ N(0, \sigma_A)$$

Combining yields firm profits from selling variety j in the domestic market:

$$\Pi_{jt} = \left[P_{jt} - \frac{W_t^{1-\alpha}(W_t^*)^\alpha}{A_{jt}}\right]C_{jt}$$

Firms are price-setters but face a menu cost κ when adjusting prices. Let the state vector of firm j be $S_{jt} = (P_{j,t-1}, A_{jt}; P_t, W_t, W_t^*)$ where $P_{j,t-1}$ and A_{jt} are idiosyncratic state variables and P_t, W_t, and W_t^* are aggregate state variables. The value of a firm selling variety j is characterized by the following Bellman equation:

$$\begin{aligned}
V^N(S_{jt}) &= \Pi_{jt}(S_{jt}) + E\{Q(S_{jt+1})V(S_{jt+1})\} \\
V^A(S_{jt}) &= \max_{P_{jt}}\{\Pi_{jt}(S_{jt}) + E\{Q(S_{jt+1})V(S_{jt+1})\}\} \\
V(S_{jt}) &= \max\{V^N(S_{jt}), V^A(S_{jt}) - \kappa\}
\end{aligned}$$

where $V^N(.)$ is the value function if the firm does not adjust its price, $V^A(.)$ is the value function if it adjusts, and $V(.)$ is the value of making the optimal price adjustment decision. $Q(S_{jt+1})$ is the stochastic discount factor. Each period the firm chooses whether to adjust its price by comparing the value of not adjusting to the value of adjusting net of the menu cost.

[46] This cost function can be derived from a CRS production function in domestic and foreign inputs.

5.1.3 Sectoral equilibrium

We define $e_t \equiv \ln(W_t^*/W_t)$ as the log real exchange rate. Sectoral equilibrium is characterized by a path of the sectoral price level, $\{P_t\}$, consistent with optimal pricing policies of firms given the exogenous idiosyncratic productivity process and wage rates in the two countries. This sectoral equilibrium allows for indirect effects that we shut down in Section 4.1 but explore in our model appendix. Following Krusell and Smith (1998) and its open economy implementation in Gopinath and Itskhoki (2010), we assume that $E_t \ln P_{t+1} = {}_0 + {}_1 \ln P_t + {}_2 e_t$. We then solve the firm's Bellman equation for a given conjecture for , simulate the model and iterate to convergence. As in Gopinath and Itskhoki (2010), this forecasting rule is highly accurate in equilibrium.

We assume that all prices are set in the domestic currency, since our empirical analysis is restricted to dollar prices. Following Gopinath and Itskhoki (2010), we assume that $W_t = 1$ and that all fluctuations in the real exchange rate arise from fluctuations in W_t^*. In economic terms, these assumptions derive from assuming that the value of the domestic currency and real wage are stable relative to the exchange rate. These are good assumptions for the U.S.

5.1.4 Calibration

While there are a number of strategic complementarities that can generate variable markups (and thus incomplete pass-through), the specific form we explore in our quantitative results is the Klenow and Willis (2006) specification of the Kimball aggregator (equation 9):

$$\Psi = \left[1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)\right]^{\frac{\sigma}{\varepsilon}}, \quad \text{where} \quad x_j \equiv D\frac{P_j}{P}$$

This demand specification is governed by two parameters: $\sigma > 1$ and $\varepsilon > 0$. The elasticity and the super-elasticity of demand are given by:

$$\widetilde{\sigma}(x_j) = \frac{\sigma}{1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)} \quad \text{and} \quad \widetilde{\varepsilon}(x_j) = \frac{\varepsilon}{1 - \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)}$$

Under these assumptions the markup is given by

$$\widetilde{\mu} = \frac{\sigma}{\sigma - 1 + \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)}$$

so that when $\varepsilon \longrightarrow 0$, we get a CES demand structure with an elasticity of substitution equal to σ and a markup equal to $\frac{\sigma}{\sigma - 1}$. The price elasticity of desired markups is given by:

$$\square \equiv \square \frac{\partial \ln \widetilde{\mu}}{\partial \ln P_j} = \frac{\varepsilon}{\sigma - 1 + \varepsilon \ln\left(\frac{\sigma x_j}{\sigma - 1}\right)}.$$

Thus, responsiveness is decreasing in ε and increasing in σ (if $\varepsilon > 0$). Since we do not directly observe σ or ε we cannot separately identify heterogeneity in these two parameters. For simplicity and following Gopinath and Itskhoki (2010), we assume that variation in □ is driven solely by ε but note that variation in σ would yield similar results. We return to this point in Section 5.4.3.

The calibrated values for all parameters are reported in Table 7. The period in our model is one month so we calibrate the discount rate to generate an annual 4% real interest rate ($\beta = 0.96^{1/12}$). We set the elasticity of demand, σ, equal to 5. This implies a steady-state markup of 25%, which is the middle of the range estimated by Broda and Weinstein (2006) using U.S. import data from 1990-2001. We assume that the log of the real exchange rate, e, follows a random walk in logs. Empirically this series is highly persistent. We set the mean increment of the innovation of the real exchange rate equal to 2.5% following Gopinath and Itskhoki (2010). To calibrate the share of imports, $\frac{\omega}{1+\omega}$, we use the share of imports as a percentage of GDP from the Bureau of Economic Analysis.[47] The four year average (2008-2011) of this import share for the U.S. is 16.5%, which implies that $\omega = 0.2$.We set the persistence of the idiosyncratic shock process, ρ_A, to be equal to 0.85, which is in between the values used by Gopinath and Itskhoki (2010) and Nakamura and Steinsson (2008), and we set κ to target a frequency of 16%.

Finally, the parameters α, ε,and σ_A are jointly calibrated to match three moments of the data: average pass-through, the R^2 from our MRPT regression and the mean standard deviation of item level price changes. To get intuition for why these moments separately identify our parameters, it is useful to remember the intuition from our simple model and our baseline MRPT regression:

$$\Delta p_{i,t} = \beta \Delta e_t + \epsilon_{i,t} \tag{11}$$

Decreasing ε means that firms respond more to both exchange rate movements and idiosyncratic shocks when adjusting prices. This increases the average level of pass-through and the standard deviation of price changes but has a negligible effect on the R^2 from estimating equation (11). This is because lowering ε increases both explained variance coming from Δe_t and unexplained variance coming from $\epsilon_{i,t}$ by roughly equal amounts so that the ratio of the residual sum of squares to the total sum of squares remains unchanged. Increasing σ_A leads to a large increase in the variance of price change and a decrease in estimated pass-through since the selection bias conditional on price adjustment is decreasing in σ_A. Increasing σ_A also leads to a large decrease in R^2, since amplifying $\epsilon_{i,t}$ increases the residual sum of squares. Finally, increasing α leads to large increases in measured pass-through but has little effect on the variance of price changes since the variance of price changes is almost entirely driven by idiosyncratic shocks. At the same time, increasing α leads to an increase in R^2 since it increases the signal to noise ratio in the pass-through regression.

Thus, movements in these three parameters produce distinctly different effects on the average level of pass-through, the R^2 from our MRPT regression, and the mean standard deviation of item level price changes so that these three moments allow us to identify our parameters of interest. We find that the best fit parameters for α, ε,and σ_A are 0.18, 2.5 and 0.07, respectively.

[47] Calibrating this import share is important to allow for realistic sectoral equilibrium effects, as discussed in the Modeling Appendix.

5.2 Simple Comparative Statics

To understand the role of various channels in explaining the empirical relationship between MRPT and the dispersion of price changes, we begin with a simple comparative statics exercise. Each panel of Figure 6 shows results when we fix three of $\varepsilon, \kappa, \alpha$ and σ_A at their steady state values and vary the fourth parameter. For each set of parameters, we simulate a panel of firms with the same number of observations as in the BLS data and compute MRPT and the standard deviation of price changes exactly as in Section 2. For comparison, the empirical relationship between the standard deviation of price changes and MRPT that we documented in the IPP microdata is shown in blue.

Figure 6: Menu Cost Comparative Statics

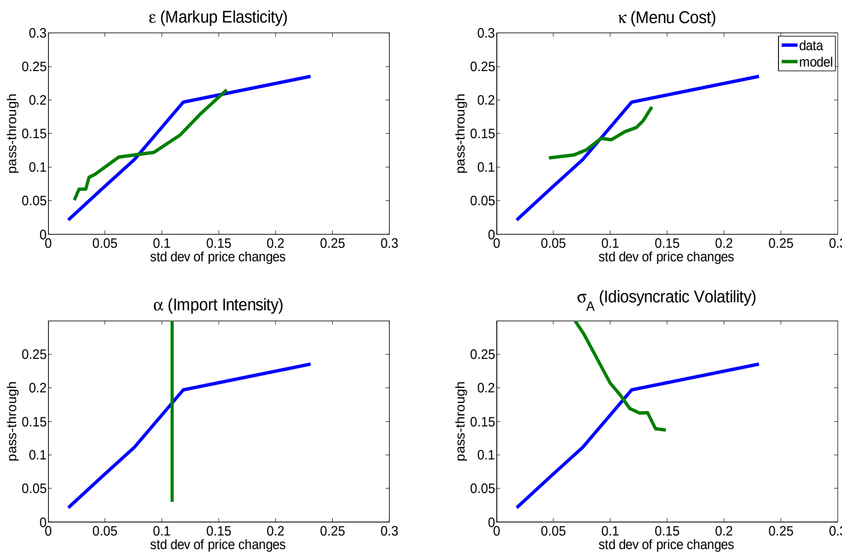

The top-left panel of Figure 6 shows the results from varying ε from 0 to 100. It is apparent that variation in ε generates a strong positive correlation between price change dispersion and MRPT. Moreover, the quantitative fit is quite good: the model matches the slope, level and much of the quantitative variation of this relationship. The bottom-left panel Figure 6 shows what happens when we vary α from 0 to 1. This leads to large changes in MRPT but negligible movements in the variance of price changes. This is consistent with the discussion in Section 2.3.1.

The top-right panel shows the model-simulated results when we vary κ from 0 to 0.2. Consistent with the discussion in the previous section, variation in κ generates a modest positive relationship between MRPT and the standard deviation of price changes. This positive correlation occurs because higher menu costs lead firms to tolerate wider price imbalances before adjusting, which amplifies selection effects. This increases price change dispersion as well as measured pass-through.

Finally, the bottom-right side panel shows results when we vary the standard deviation of idiosyncratic shocks from 0 to 0.2. Variation in σ_A generates a strong negative relationship between MRPT and the standard deviation of price changes for reasons similar to menu cost variation but in reverse: larger σ_A increases price change dispersion but implies that firms are more likely to adjust their prices for purely idiosyncratic reasons, which reduces selection effects and MRPT.

Thus, the comparative statics imply that variation in ε or κ can potentially replicate the observed relationship between MRPT and the standard deviation of price changes. However, explaining the relationship through variation in κ yields grossly counterfactual implications for the frequency of adjustment. In the data there is a mild positive correlation between adjustment frequency and price change dispersion. In contrast, variation in κ induces an almost perfect negative correlation between dispersion and frequency: as menu costs rise, the inaction region widens, frequency falls and price change dispersion rises. We return to this point in our indirect inference results.

The relation between our comparative statics and those in Gopinath and Itskhoki (2010) Proposition 3 bears some mention. They find that in a simple static model, pass-through increases with α, declines with ε and is unaffected by κ or σ_A. Our conclusion for α and ε is identical, but our results for κ and σ_A differ because we study MRPT while they study LRPT. LRPT is not subject to the selection effects that induce $cov(\epsilon_{i,t}, \Delta e_t) > 0$ but these effects are important for MRPT.[48]

While we view this comparative statics exercise as highly informative, it has several weaknesses: 1) In the data, we are sorting firms into bins by the standard deviation of price changes. Since our comparative statics exercise instead computes results for a series of models that vary by a single parameter, we are implicitly sorting firms by this (unobserved) parameter rather than by the standard deviation of price changes. Thus, there is not a perfect match between our comparative statics simulations and our empirical exercise. 2) In the data, firms are likely to differ along many dimensions simultaneously so that heterogeneity is unlikely to be well-captured by a single parameter. 3) The comparative statics exercise is intrinsically qualitative and informal. For example, both κ and ε generate positive relationships between MRPT and dispersion and there is little formal guidance for which is a better fit even along this single moment.

We now turn to a formal estimation strategy that squarely addresses each of these weaknesses.

5.3 Indirect Inference

In this section, we allow for permanent firm heterogeneity, which we assume is unobserved by the econometrician. We then formally estimate the importance of different forms of heterogeneity in explaining our empirical results using indirect inference.[49] Motivated by tractability as well as the results from the comparative statics exercise, we allow for three dimensions of heterogeneity across firms. In particular, we allow firms to differ by κ, ε and σ_A.[50] We assume that each parameter

[48] In addition, in the working paper version they explore implications for the size of price adjustment. As we showed in the theoretical results, in a flexible price environment, increases in volatility or responsiveness unambiguously increase price change dispersion. With menu costs this need not occur. Nevertheless, we find that the qualitative result from the frictionless model strongly obtains for all estimated parameters in our quantitative model.

[49] See Collard-Wexler (2013) and Keane and Anthony (2003) for examples of indirect inference in similar problems.

[50] While it would also be possible to allow for heterogeneity in α, our comparative statics exercise suggests that this parameter plays no role in explaining the relationship between MRPT and dispersion or in the relationship between dispersion and frequency. However, shutting down heterogeneity along this dimension substantially reduces the computational burden involved in estimation.

takes on one of two values uniformly distributed around the previous mean.[51] For example, we assume that for a particular firm, κ is either equal to $\kappa_h = .043 + \kappa_\Delta$ or $\kappa_l = .043$ □ κ_Δ where κ_Δ is a parameter to be estimated which governs the degree of menu cost differences across firms. We allow for a similar two point symmetric distribution for each source of heterogeneity so that we have three parameters which must be estimated: $\theta = (\kappa_\Delta, \sigma_\Delta, \varepsilon_\Delta)$.

Fixing $\kappa_\Delta, \sigma_\Delta, \varepsilon_\Delta$ there are then eight different types of firms in our model (taking on high or low values for each parameter), and we assume an equal number of firms of each type.[52] After solving for the sectoral equilibrium with these eight firm types we simulate a firm panel, which we sample exactly as in the BLS microdata to account for any small sample issues which might arise in our empirical specification. From this firm panel we calculate an auxiliary model that consists of fifteen reduced form moments $g(\theta)$ which capture essential features of the data. We then try to match these simulated moments to their empirical counterparts.

This indirect inference estimation procedure explicitly addresses the concerns identified with the comparative statics exercise: simulated and actual data are treated identically and we use no information from simulated data that is not available in actual data. In addition, we explicitly allow for the presence of multiple sources of heterogeneity and formally assess their relative importance.

To construct our empirical moments, we first sort firms into five bins by their standard deviation. We then calculate the relative standard deviation of price changes, the relative MRPT, and the relative frequency for each standard deviation bin.[53] The first five moments test the model's ability to capture the heterogeneity in price change dispersion observed in the data. The second five moments capture the relationship between this dispersion and pass-through. The final five moments capture the relationship between dispersion and frequency, which we previously argued helps identify heterogeneity in menu costs from heterogeneity in responsiveness.[54]

Given these 15 moments, we pick our 3 parameters to solve $\widehat{\theta} = \arg\min_\theta g(\theta)' W(\theta) g(\theta)$ where $W(\theta)$ is a positive definite weight-matrix.[55] Table 8 displays resulting parameter estimates as well as several measures of model fit. The first take-away from Table 8 is that the estimated level of ε heterogeneity is large and significant. In contrast, heterogeneity in σ_A is significant but not strongly so, and there is no evidence for heterogeneity in κ. We can also assess the overall model fit. Using standard over-identification tests our full model cannot be rejected at 99% confidence levels. We

[51] When relevant, we bound the value of $\kappa_l, \varepsilon_l, \sigma_l$ at 0.

[52] While it would be desirable to allow for more than a 2-point distribution of heterogeneity for each parameter, allowing for a 3-point distribution would require solving the model for 27 different types of firms while allowing for a 4-point distribution would require 64 firm types, so it is clear that the problem rapidly rises in difficulty. Since we want to estimate the model, we must resolve it for a large number of $\kappa_\Delta, \sigma_\Delta, \varepsilon_\Delta$ which rapidly becomes infeasible. Allowing for different numbers of each firm also greatly increases the parameter space.

[53] We concentrate on the relative values rather than the absolute values because our benchmark calibration is not perfectly able to match the level of XSD, MRPT and freq. We think of both our empirical exercise and our exercise with heterogeneity largely as being about matching the relative differences across firms. Nevertheless, redoing the results using absolute rather than relative moments did not qualitatively change the conclusions.

[54] As is standard in indirect inference and in contrast to typical simulated GMM implementations, our auxiliary model need not have any structural interpretation. For example, we have already noted that our OLS MRPT regression will pick up both direct effects of parameters on β as well as indirect effects on covariance terms.

[55] We pick $W(\theta)$ to be the standard efficient weight matrix so that we can apply asymptotic formulas for standard errors but using an identity weight matrix did not change our qualitative conclusions.

can also investigate restricted models that turn off various sources of heterogeneity. The results for the restricted models show that the model with no heterogeneity in ε can easily be rejected while models with no heterogeneity in κ or in σ_A cannot be rejected in favor of the full model.

Figure 7: Full and Restricted Model Fits

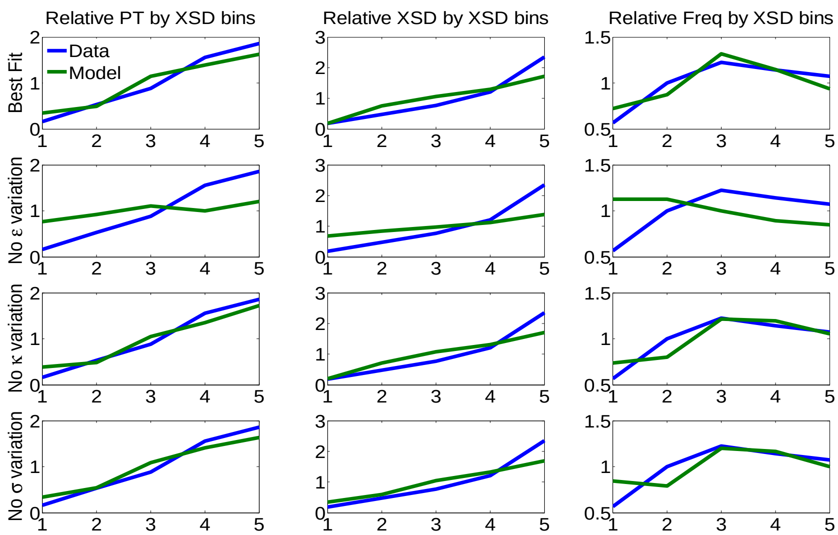

The numerical results can be seen more easily in Figure 7, which shows the model fit to all fifteen moments as well as the fit of restricted models which shut down various sources of heterogeneity.

The main take-away from this visual inspection is that the fit in the second row is dramatically worse than the fit in the first row. Turning off heterogeneity in ε means the next-best model fit does not generate enough heterogeneity in price change dispersion, fails to generate enough of a positive relationship between price change dispersion and pass-through, and it implies a negative rather than positive correlation between dispersion and pass-through. In contrast, turning off heterogeneity in menu costs or in volatility has only negligible effects on the model fit.

5.4 Aggregate Shocks

5.4.1 Baseline Results

Now that we have shown that variation in responsiveness is crucial for explaining the empirical relationship between item-level price change dispersion and pass-through, we turn to understanding our month-level dispersion results. In general, our conclusions about the role of various economic primitives in the cross-section transfer almost directly to the time-series. For this reason, our discussion of these simulations is intentionally brief.

In the previous section, we assumed that there was heterogeneity across firms that was constant across time. Instead, we now assume that firms are identical but are subject to various aggregate shocks.[56] We consider aggregate shocks to each of our parameters in turn. For expositional purposes we describe only ε shocks, but we treat other shocks analogously.

For simplicity, we assume that ε_t follows a two-state Markov process with transition probabilities $\begin{bmatrix} \Pi_{11} & \Pi_{12} \\ \Pi_{21} & \Pi_{22} \end{bmatrix}$ and allow the Krusell-Smith forecast for the sectoral price level to depend on ε_t.[57] We have little guidance on the size or persistence of our aggregate shocks, so rather than taking a strong stand, we explore several different parameterizations. Under the "small" shock calibration, ε_t moves between $(1+.6)\,\overline{\varepsilon}$ and $\frac{1}{1+.6}\,\overline{\varepsilon}$ where $\overline{\varepsilon}$ is the previous baseline calibration. This small shock calibration implies that time-series variation in ε is roughly one-fifth as large as the cross-sectional variation estimated in the previous section. In addition, we consider a "large" shock calibration that moves ε_t between $4\overline{\varepsilon}$ and $\frac{1}{4}\,\overline{\varepsilon}$. This produces time-series variation in ε comparable to the cross-sectional variation in the previous section. We have computed results for both a low monthly shock persistence of $\Pi_{11} = \Pi_{22} = 0.90$ and a high persistence of $\Pi_{11} = \Pi_{22} = 0.975$. Changing the persistence barely affected our results, so for brevity we report only the high persistence case.

Table 9 shows results with different aggregate shocks. In all cases, we divide months in thirds by their month-level dispersion and calculate pass-through in high and low dispersion months. Overall, aggregate shocks to ε_t are most consistent with our empirical time-series results. Increases in ε reduce pass-through and the standard deviation of price changes. Under the "large shock" calibration, ε variation roughly accounts for the variation in MRPT observed in the data. However, the movements in price change dispersion across time are somewhat too large: the model produces a cross-sectional standard deviation of price changes that ranges from 0.05 to 0.13. In the data, the comparable range is 0.12 to 0.15.[58] In ongoing work we plan to explore whether alternative shocks that relax the binary assumption can provide a better fit to the data. Nevertheless, shocks to ε produce variation in pass-through and dispersion that is reasonably consistent with the data.

In contrast, shocks to σ_A induce the wrong correlation between the standard deviation of price changes and pass-through. In addition, they produce time-series variation in both price change dispersion and frequency that are substantially too large relative to the data.

Shocks to κ generate some comovement between dispersion and pass-through, but they imply a strong negative relationship between frequency and dispersion. In the data, frequency, dispersion, and pass-through all comove. Shocks to κ also generate too much time-series variation in frequency.

Shocks to α induce lots of movement in pass-through but almost no movement in the standard deviation of price changes. In addition, the small movement in price change dispersion induced by α goes in the wrong direction. As α rises, pass-through rises but the cross-sectional standard deviation of price changes falls. That is because large α effectively increases the size of the exchange rate shocks relative to idiosyncratic shocks. Since the exchange rate shock is common to all firms,

[56] Including item-level heterogeneity together with aggregate shocks did not alter our conclusions but makes the model somewhat more complicated.

[57] That is, we assume that $E_t \ln P_{t+1} = {}_0 + {}_1 \ln P_t + {}_2 e_t + \varepsilon_t \times [{}_4 + {}_5 \ln P_t + {}_6 e_t]$. Again we find that the Krusell-Smith forecasting rule is highly accurate.

[58] As previously mentioned, our baseline calibration mildly underpredicts the average standard deviation of price changes in the data

this reduces the cross-sectional dispersion of price changes.

Thus, as in the cross-sectional results, only shocks to ε do a reasonable job of reproducing the empirical evidence.

5.4.2 Additional Shocks

In addition to the above aggregate shocks, which we also explore in the cross-section, we study two additional aggregate shocks which are more applicable to the time-series. First, we allow the volatility of exchange rates to change across time, since the 2008 recession was also associated with greater exchange rate volatility. However, we find that even large increases in exchange rate volatility have only mild quantitative effects, and qualitatively have the wrong sign relative to the empirical evidence. That is, increasing the volatility of exchange rates mildly increases pass-through, but (very mildly) decreases month-level dispersion. This is for the same reason that increases in α decrease the dispersion of price changes.

It is also possible that the large degree of pass-through observed during the Great Recession was driven by the fact that the recession was a large shock which affected many firms. If a shock is common to more firms, then it might have greater general equilibrium effects and generate more pass-through. To assess the role of the "commonness" of shocks, we introduce time-variation in the fraction of firms that are sensitive to the exchange rate, ω. As ω rises, exchange rate shocks affect more firms and general equilibrium effects increase in importance. However, the quantitative effect of changes in ω on pass-through is relatively small and there are no effects of ω on the dispersion of price changes: increasing ω from 0.2 to 0.9 only increases pass-through from 16% to 23% and has no effect on dispersion. Thus, general equilibrium effects in our model cannot account for the empirical relationship between month-level dispersion and exchange rate pass-through.

5.4.3 Interpreting ε shocks

Does time-variation in ε have a natural economic interpretation? It is difficult to provide direct evidence of time-varying super elasticity of demand, but Gopinath and Itskhoki (2010) argue for the importance of such variation in the cross-section. It is easy to imagine that forces which drive super elasticity differences across firms might also vary across time. Furthermore, it is important to note that time-variation in the elasticity of substitution rather than super elasticity also delivers similar quantitative results. This is because on average, $\square = \frac{\varepsilon}{\sigma \square 1}$, so that σ shocks generate time-variation in responsiveness, as long as firms exhibit some incomplete markup adjustment ($\varepsilon > 0$).

What is crucial for explaining the time-series relationship between pass-through and dispersion is reduced-form time-variation in $\square$; the structural source of this time-variation is less important. For example, business cycle shocks to "market competitiveness" would equally well explain our time-series patterns through the same $\square$ channel. If certain periods of time such as recessions are characterized by increased competition, with larger σ and lower markups, they will also be times of greater responsiveness and price change dispersion. These cyclical demand elasticity stories have received some recent theoretical attention in Kaplan and Menzio (forthcoming), and Stroebel and

Vavra (2015) provide microeconomic empirical evidence for exactly this form of markup variation.

While we model time-varying responsiveness using a Kimball demand framework, this mechanism is much more general. Any shock to strategic complementarities across time that affects markup adjustment (□) will deliver similar predictions. As reviewed in Burstein and Gopinath (2013), a number of other mechanisms can also generate less than perfect responsiveness including desire to maintain market share (Atkeson and Burstein (2008)), customer shopping concerns (Paciello et al. (2013)), or local distribution costs. We believe that better understanding the source of "responsiveness" shocks and time-varying strategic complementarities is an interesting avenue for future research. While these shocks help fit the data, we think there is much work to be done exploring their plausibility, size and implications for business cycles more generally.

5.5 Relationship to Existing Studies of Countercyclical Dispersion

Our paper joins a long literature documenting countercyclical dispersion of various economic variables. At the same time, theory has emerged trying to match this empirical evidence and explore its macroeconomic implications. This theoretical work has focused largely on "uncertainty" or "volatility" shocks that raise the variance of shocks hitting agents in the economy.

For example, in the context of retail prices Vavra (2014) documents that the dispersion of price changes in the CPI is countercyclical and explains this fact using volatility shocks. However, looking at equation 8, it is clear that the dispersion of price changes could increase either because $var(\epsilon_i)$ rises or because □ falls. That is, greater dispersion of price changes could be explained by greater volatility of shocks and constant responsiveness, or it could be explained by greater responsiveness and constant volatility. Greater volatility and greater responsiveness both lead to an increase in aggregate price flexibility, but they do so through different underlying mechanisms with different implications for policy. By only looking at data on the dispersion of price changes, it is fundamentally impossible to differentiate time-varying volatility from time-varying responsiveness, so existing models such as Vavra (2014) have proceeded by assumption rather than empirical evidence.[59]

In contrast to the existing literature, our open economy environment allows us to separately identify changes in volatility from changes in responsiveness. This identification result was precisely the point of the previous sections, which showed that our import price data strongly supports time-variation in responsiveness rather than volatility shocks as an explanation for countercyclical dispersion. Increases in volatility are unable to explain increases in pass-through. In contrast, greater responsiveness increases both price change dispersion and exchange rate pass-through in a manner consistent with the data. This result holds across a variety of price-setting environments, whether price adjustment is frictionless, time-dependent or state-dependent.

Together, our results suggest that the literature studying countercyclical dispersion has embraced time-varying volatility too quickly. Time-variation in the strength of strategic complementarities (and thus responsiveness) appears to be more relevant, at least for import price-setting. Under-

[59] While we frame this discussion in terms of price-setting, an identical argument applies to the dispersion of any outcome that includes an endogenous component. For example revenue TFP dispersion depends on firms' price decisions.

standing the generalizability and empirical relevance of our results for other sectors of the economy and other economic outcomes is an important avenue for future research.

6 Conclusion

In this paper, we exploit the open economy environment to provide evidence on the response of inflation to cost shocks at a moment in time and use this response to infer the shocks which drive price change dispersion. We start by documenting a strong positive relationship between exchange rate pass-through and the dispersion of item-level price changes. Through a battery of robustness checks, we argued that this relationship was not driven by other observables or confounding variables. Furthermore, price change dispersion varies dramatically across time. Ignoring this variation induces large time-varying bias when estimating pass-through. In other words, pass-through is not a single number and ignoring time-variation produces a misleading picture of how prices will respond to shocks at a particular point in time.

This result has important implications for the nominal transmission mechanism more generally. One reason for studying exchange rate pass-through is because it can shed light on the nominal transmission mechanism using a large, observable shock to nominal cost. Our results provide "model-free" evidence that the transmission mechanism varies systematically across time with the distribution of price changes in the economy. We believe this result is important for policy making, because monetary policy is not conducted at random times. Monetary easing is likely to occur during times when we have shown that price flexibility is systematically higher than average which suggests that the effects of monetary policy are time-varying.

After documenting the empirical relationship between dispersion and pass-through, we estimate a quantitative price-setting model to understand it. Our estimated model strongly rejects volatility shocks as a source of countercyclical dispersion but also suggests a promising alternative. In particular, time-variation in responsiveness driven by strategic-complementarities better fits the data. A large existing literature has argued that imperfect responsiveness is important for explaining low average pass-through, and we find the idea that responsiveness might also vary across time to be quite plausible.

While we showed that "responsiveness" matters, there are many mechanisms that map into responsiveness as a reduced form. Trying to disentangle these mechanisms is an interesting avenue for future research. Our BLS data has limited firm-level covariates which can shed light on underlying mechanisms, but alternative data sets with more firm-level characteristics exist. Together our empirical and modeling results suggest that exploring time-variation in the competitive structure of markets or in strategic-complementarities and trying to test these ideas in alternative data is a promising research topic. In contrast, volatility shocks should imply a negative relationship between pass-through and price change dispersion, which is strongly at odds with the data.

Our results have both obvious and more subtle implications for policy. Most obviously, if policy makers want to understand how prices are likely to respond to exchange rate changes or predict the real responses to nominal exchange rate changes, they cannot ignore individual price-setting behavior.[60] More subtly, if policy makers want to mitigate the adverse effects of price change volatility then it is important to understand what leads to this volatility. Our results suggest that policy makers should focus on policies that affect market structure and firms' responsiveness rather than on policies that reduce uncertainty and volatility.

References

Alvarez, F., H. Le Bihan, and F. Lippi (2014). Small and Large Price Changes and the Propagation of Monetary Shocks. *NBER Working Paper 20155*.

Atkeson, A. and A. Burstein (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review 98*(5).

Bachmann, R. and G. Moscarini (2012). Business Cycles and Endogenous Uncertainty.

Berger, D., J. Faust, J. Rogers, and K. Steverson (2012). Border Prices and Retail Prices. *Journal of International Economics 88*(1).

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*(3).

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry (2012). Really Uncertain Business Cycles. *NBER Working Paper 18245*.

Broda, C. and D. Weinstein (2006). Globalization and the Gains from Variety. *Quarterly Journal of Economics 121*(2).

Burstein, A. and G. Gopinath (2013). International Prices and Exchange Rates. *Handbook of International Economics 4*.

Cavallo, A. (2012). Scraped Data and Sticky Prices.

Collard-Wexler, A. (2013). Demand Fluctuations in the Ready-Mix Concrete Industry. *Econometrica 81*(3).

Decker, R., P. Derasmo, and H. Boedo (2015). Market Exposure and Endogenous Firm Volatility Over the Business Cycle. *AEJ: Macro*.

Dornbusch, R. (1987). Exchange Rates and Prices. *American Economic Review 77*(1).

Gopinath, G. and O. Itskhoki (2010). Frequency of Price Adjustment and Pass-Through. *Quarterly Journal of Economics 125*(2).

[60] One might wonder if our empirical observations are implementable for actual predictions. We note that in our BLS data, dispersion is much more precisely estimated at high frequencies than is pass-through. In addition, new online data sets such as those in the Billion Prices Project introduced by Cavallo (2012) can potentially calculate daily measures of price change volatility.

Gopinath, G., O. Itskhoki, and R. Rigobon (2010). Currency Choice and Exchange Rate Pass-through. *American Economic Review 101*(1).

Gopinath, G. and R. Rigobon (2008). Sticky Borders. *Quarterly Journal of Economics 123*(2).

Hellerstein, R., D. Daly, and C. Marsh (2006). Have U.S. Import Prices Become Less Responsive to Changes in the Dollar? *NY Fed: Current Issues in Economics and Finance*.

Keane, M. and S. Anthony (2003). Generalized Indirect Inference for Discrete Choice Models.

Klenow, P. and J. Willis (2006). Real Rigidities and Nominal Price Changes.

Krugman, P. (1987). Pricing to Market When the Exchange Rate Changes. *Real-Financial Linkages Among Open Economies*.

Krusell, P. and A. A. Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy. *The Journal of Political Economy 106*(5).

Marazzi, M., N. Sheets, R. Vigfusson, J. Faust, J. Gagnon, J. Marquez, R. Martin, T. Reeve, and J. Rogers (2005). Exchange Rate Pass-Through to U.S. Import Prices: Some New Evidence. *International Finance Discussion Papers*.

Midrigan, V. (2011). Menu Costs, Multi-Product Firms and Aggregate Fluctuations. *Econometrica 79*(4).

Nakamura, E. and J. Steinsson (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics 123*(4).

Nakamura, E. and J. Steinsson (2012). Lost in Transit: Product Replacement Bias and Pricing to Market. *American Economic Review 102*(7).

Neiman, B. (2010). Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices. *Journal of Monetary Economics*.

Paciello, L., A. Pozzi, and N. Trachter (2013). Price Setting with Customer Retention.

Stroebel, J. and J. Vavra (2015). House Prices, Local Demand, and Retail Prices. *NBER Working Paper 20710*.

Vavra, J. (2014). Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation. *Quarterly Journal of Economics*.

Table 1: Average medium-run pass-through

β	$se(\beta)$	t-stat	N_{obs}	R^2
0.144	0.014	10.17	95284	0.067

Table 2: Interaction Specification: Item-Level Volatility

	Average pass-through		Volatility (Item-Level)		Frequency		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
All countries, all items ex petroleum								
	0.14	0.01	0.05	0.02			95284	0.07
	0.14	0.01	0.05	0.02	0.02	0.01	95284	0.07
OECD countries, all items ex petroleum								
	0.18	0.02	0.09	0.03			53469	0.08
	0.19	0.02	0.08	0.03	0.07	0.02	53469	0.08
All countries, all manufacturing items								
	0.14	0.01	0.06	0.02			78439	0.09
	0.13	0.01	0.06	0.02	0.03	0.01	78439	0.09

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. In all specifications, volatility is measured using the item-level standard deviation of price changes.

Table 3: Interaction Specification: Month-Level Volatility

	Average pass-through		Volatility (Month-Level)		Frequency		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
All countries, all items ex petroleum								
- Vol=Month-level IQR	0.14	0.01	0.06	0.01			95284	0.07
	0.14	0.01	0.06	0.01	0.01	0.01	95284	0.07
- Vol=Month-level XSD	0.13	0.01	0.05	0.01			95284	0.07
	0.13	0.01	0.05	0.01	0.03	0.01	95284	0.07
OECD countries, all items ex petroleum								
- Vol=Month-level IQR	0.17	0.01	0.05	0.01			53469	0.08
	0.17	0.01	0.05	0.01	-0.01	0.01	53469	0.08
- Vol=Month-level XSD	0.17	0.01	0.06	0.01			53469	0.08
	0.18	0.01	0.07	0.01	-0.03	0.02	53469	0.08
All countries, all manufacturing items								
- Vol=Month-level IQR	0.13	0.01	0.05	0.01			78437	0.09
	0.13	0.01	0.05	0.01	0.00	0.01	78437	0.09
- Vol=Month-level XSD	0.13	0.01	0.05	0.01			78437	0.09
	0.13	0.01	0.05	0.01	0.00	0.01	78437	0.09

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes.

Table 4: Interaction Specification: Month-Level Dispersion Robustness

	Average pass-through		Volatility (Month-Level)		Frequency/Subs		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
Controls								
- Time trend + Month	.135	.025	.058	.012			95284	.075
- GDP growth	.146	.013	.054	.009			95284	.071
- Erate SD	.156	.015	.056	0.010			95284	.072
- Frequency	.140	.013	.063	.010	.011	.012	95284	.072
- Product subs	.143	.013	.062	.010	.0004	.011	95284	.071
- Time trend + Month + Frequency	.128	.025	.054	.012	.011	.014	95284	.076
+ GDP growth + Erate SD								

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. In all specifications month-level price change volatility is measured using the interquartile range. The sample is all countries, all items ex petroleum. Erate SD is the std dev of the exchange rate in a 12-month rolling window around current date

Table 5: Interaction Specification: Month-Level Dispersion Robustness: Pre-2008

	Average pass-through		Volatility (Month-Level)		Frequency/Subs		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
Controls								
- Time trend + Month	.127	.029	.039	.015			74153	.075
- Frequency	.118	.014	.042	.011	.016	.013	74153	.072
- Product subs	.115	.014	.036	.012	.007	.012	74153	.071
- Time trend + Month + Frequency	.121	.029	.040	.015	.023	.016	74153	.076
- Time trend + Month + Product subs	.125	.027	.038	.015	-.001	.011	74153	.075

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. In all specifications month-level price change volatility is measured using the interquartile range.

Table 6: Alternative Pass-through Specifications

	Average pass-through		Item-Level Volatility		Month-Level Volatility			
Fixed Horizon:	β^{avg}	$se(\beta^{avg})$	β^{XSD}	$se(\beta^{XSD})$	β^{IQR}	$se(\beta^{IQR})$	N_{obs}	R^2
1 Month	.027	.007	.034	.013	.023	.006	496060	.018
3 Month	.054	.011	.048	.023	.026	.007	448400	.049
6 Month	.085	.016	.069	.033	.026	.009	384827	.098
12 Month	.113	.018	.093	.022	.023	.009	282572	.169
Lagged Specification:								
Current Ex. Rate (β_1)	.146	.015	.040	.020	.063	.010		
Previous Ex Rate (β_2)	.082	.010	.040	.017	.054	.010	83043	.082

Note: Robust standard errors clustered by country*PSL pair.

Table 7: Parameter Values

Parameter	Symbol	Menu Cost Model	Source
Discount Factor	β	$0.96^{1/12}$	Annualized interest rate of 4%
Fraction of imports	$\omega/(1+\omega)$	16.5%	BEA input-output table
Cost sensitivity to ER shock			
Foreign firms	α^*	0.18	Estimation (see text)
U.S. firms	α	0	
Menu cost	κ	4.3%	Estimation (see text)
markup elasticity	ε	2.5	Estimation (see text)
Demand elasticity	σ	5	Broda and Weinstein (2006)
Std. dev. Exchange rate shock, e_t	σ_e	2.5%	Match bilateral RER
Idiosyncratic productivity process, a_t			
Std. dev. of shock	σ_A	7.0%	Estimation (see text)
Persistence of shock	ρ_A	0.85	Gopinath and Itshkoki (2010)

Table 8: Estimated Parameters and Fit

Parameter	Estimate	95% Confidence Interval
ε_Δ	10	(8.14,11.86)
σ_Δ	.03	(.0035,.0565)
κ_Δ	.014	(-.0125,.0405)

Models	Wald-Statistic/Likelihood Ratio	95% Critical Value	99% Critical Value
Unrestricted Model	25.76	21.03	26.22
$\varepsilon_\Delta = 0$	46.9	3.84	5.63
$\sigma_\Delta = 0$	3.23	3.84	5.63
$\kappa_\Delta = 0$	3.57	3.84	5.63

Asymptotic s.e.'s for parameters in parantheses. Unrestricted model Wald-Statistic: $g\left(\widehat{\theta}\right)' W\left(\widehat{\theta}\right)' g\left(\widehat{\theta}\right) \sim \chi^2(12)$

Restricted models: $2\left[g\left(\widehat{\theta}_r\right)' W\left(\widehat{\theta}_u\right)' g\left(\widehat{\theta}_r\right) \square\, g\left(\widehat{\theta}_r\right)' W\left(\widehat{\theta}_u\right)' g\left(\widehat{\theta}_r\right)\right] \sim \chi^2(1)$

Table 9: Aggregate Shocks

	Data (Low XSD)			Data (High XSD)		
	XSD	MRPT	FREQ	XSD	MRPT	FREQ
	0.12	0.08	0.14	0.15	0.17	0.18
	High ε			Low ε		
	XSD	MRPT	FREQ	XSD	MRPT	FREQ
Small	0.08	0.15	0.11	0.11	0.20	0.13
Large	0.05	0.12	0.08	0.13	0.22	0.15
	Low σ			High σ		
	XSD	MRPT	Freq	XSD	MRPT	FREQ
Small	0.07	0.19	0.08	0.13	0.14	0.21
Large	0.07	0.20	0.08	0.23	0.11	0.45
	High κ			Low κ		
	XSD	MRPT	FREQ	XSD	MRPT	FREQ
Small	0.10	0.21	0.08	0.09	0.15	0.16
Large	0.10	0.30	0.04	0.07	0.13	0.29
	High α			Low α		
	XSD	MRPT	FREQ	XSD	MRPT	FREQ
Small	0.09	0.29	0.12	0.09	0.13	0.12
Large	0.08	0.66	0.14	0.09	0.06	0.13

Small = ± 0.6 factor, Large = ± 3.0 factor. Binary agg shock has persistence .975

.

Online Appendix Materials - Not For Publication

7 Empirical Appendix - Not for Publication

In this empirical appendix, we provide a number of additional robustness checks that extend the baseline results in the body of the text.

7.1 Additional Item-Level Results

In this section we perform a variety of robustness checks for our item-level results. We begin by further discussing the role of adjustment frequency, then return to compositional issues. Finally, we discuss a battery of additional robustness checks, alternative samples, and placebo regressions for our baseline results.

In the main text we estimate interaction specifications to argue that differences in frequency across items do not explain our relationship between cross-item dispersion and MRPT. However, that specification assumes that the effects of dispersion are linear and does not allow for the effects of other controls to vary with item-level characteristics. In this robustness check we provide further evidence that the relationship between cross-item dispersion and MRPT is not driven by frequency. This is an important concern to address because Gopinath and Itskhoki (2010) showed that there is a robust relationship between LRPT and the frequency of adjustment. In order to further address this concern, we split items first into equal weighted frequency quintiles then examine the relationship between MRPT and item-level dispersion within each frequency quintile. In other words, we examine the relationship between pass-through and dispersion holding the frequency of adjustment (roughly) constant. The results are shown in Figure 8.

Figure 8: Medium-run passthrough and XSD controlling for the frequency of price adjustment

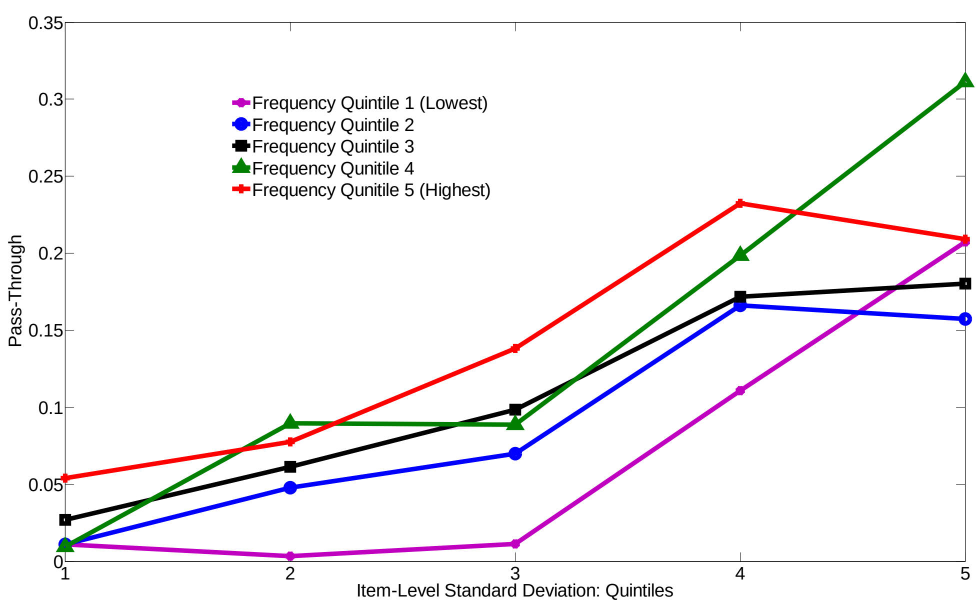

The relationship between pass-through and dispersion is increasing within each frequency quintile, and the magnitude of the increase is substantial. Average pass-through increases from 3% to 20% as we move from the lowest to highest XSD quintile. This complements the evidence in the text that the relationship between MRPT and price dispersion does not seem to be driven by differences in frequency across items.[61] Here it is worth noting that censoring might be an important concern for explaining differences in pass-through by frequency. That is, the fact that the lowest frequency items have lower pass-through than the highest frequency items might reflect censoring and thus missing price changes. However, it is difficult to explain why there is an upward sloping relationship between pass-through and dispersion within frequency bins if this is the only source of heterogeneity. While variation in strategic-complementarities can generate more complicated relationships since they affect frequency, pass-through and dispersion, this is exactly the channel which we explore in our quantitative model. Simulations in this quantitative model explicitly account for censoring, so censoring effects alone cannot explain our patterns.

We next address whether our results are driven by choice of which items we sampled. Our baseline results utilize all of the items in the IPP micro data excluding petroleum. Is the strong relationship between pass-through and dispersion affected if we split by other observable product characteristics? To address this question we first examine the sub-sample of goods that can be classified as differentiated, following Rauch's classification, as well as the sample of goods that are manufactured.[62] For differentiated goods, Figure 9 shows that moving from the lowest to highest-dispersion quintile raises MRPT from 2% to 27%. Similar results obtain when using all manufactured goods. In all cases, the difference in pass-through across XSD bins is strongly statistically significant.

In addition to splitting by product type, we can also split our sample by country of origin. Perhaps our results are driven by compositional differences in the behavior of items coming from different countries. Figure 10 shows this is not the case. For all countries and country groups with greater than 5000 price observations there is a strong upward sloping relationship between item-level dispersion and MRPT. While the relationship is insignificant for Mexico, the sample size is small relative to more aggregated country groups and Canada. Among countries with at least 5000 observations, only Japan has fewer observations (and Japan exhibits a significant upward sloping relationship).[63]

In addition to these alternative binned regressions, Table A1 shows the results from estimating equation 3 for a variety of alternative sub-samples and alternative specifications. The first robustness check only uses items which have at least 3 changes. It is difficult to precisely measure dispersion for items with few price changes, so there is some concern that our baseline specifications might be polluted by outliers and small sample issues. However, the first two rows of Table A1 show that our results are essentially unchanged when restricted to items with at least 3 price changes. We have

[61] This is not surprising since Gopinath and Itskhoki (2010) document a significant relationship between LRPT and the frequency of price adjustment but find no relationship between MRPT and the frequency of price adjustment.

[62] Items are classified as manufacturing items if their 1-digit SIC 1987 codes begin with a 2 or a 3.

[63] While it would be desirable to run this specification for each country rather than aggregating to country groups, our empirical specification requires splitting the data into fifths and then estimating second moments of price-setting on these bins, so using smaller countries becomes infeasible.

Figure 9: Medium-run passthrough by XSD

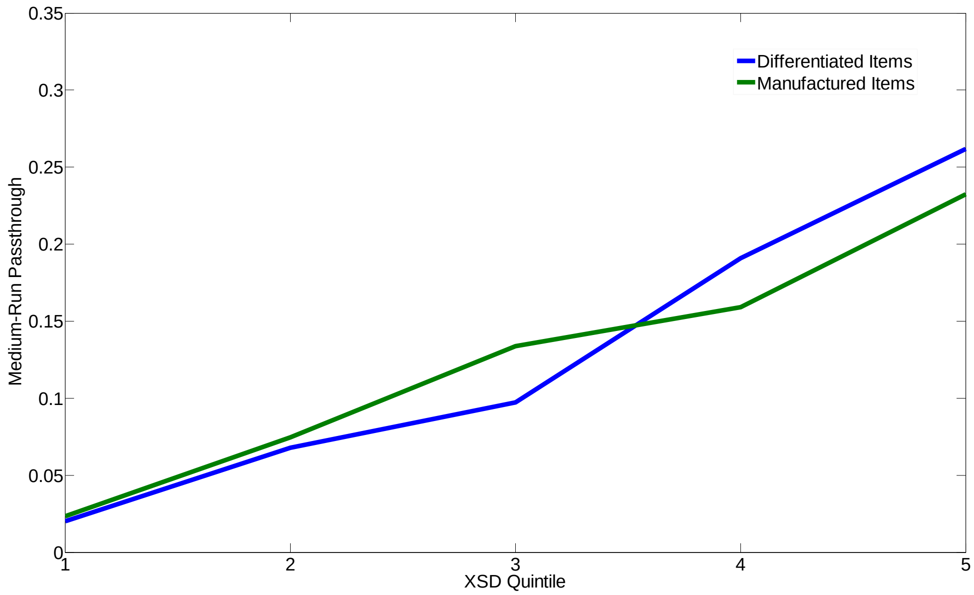

also restricted the analysis to items with at least 5 price changes and arrived at similar results.

In our second and third robustness checks, we use trade-weighted exchange rates (the broad and major currency one respectively) instead of the relevant bilateral exchange rate. As rows 3-6 show, the price dispersion is again both economically and statistically significant. A one standard deviation increase in price dispersion causes MRPT to increase relative to average pass-through by over 50%.[64]

Rows 7-10 show the results from our fourth and fifth robustness checks. In these robustness checks, we run placebo regressions to see whether our results are spuriously driven by small sample issues. In these placebo regressions, when estimating equation 3, we substitute the total number of price changes observed for an item or the number price observations respectively for XSD. These placebo regressions test whether our results are driven by a correlation between measured dispersion and item sample sizes. Table A1 shows that, as desired, the coefficient on β^{XSD} is not significant when we replace XSD with placebos. This suggests that the relationship between MRPT and price dispersion is not being driven by sampling error. Finally, rows 11 and 12 show the results from estimating equation 3 using a median regression rather than OLS. Median regressions are more robust to the presence of outliers. Once again, the price dispersion effect is strongly significant.

[64] Consistent with what was found in Nakamura and Steinsson (2012), average passthrough is significantly higher when we use broader exchange rates measures. The much larger response of prices to the trade-weighted exchange rate suggests that items respond to exchange rates beyond the bilateral one, presumably due to the role of intermediate inputs and strategic complementarities in pricing.

Figure 10: Medium-run Passthrough by XSD (Country Level Results)

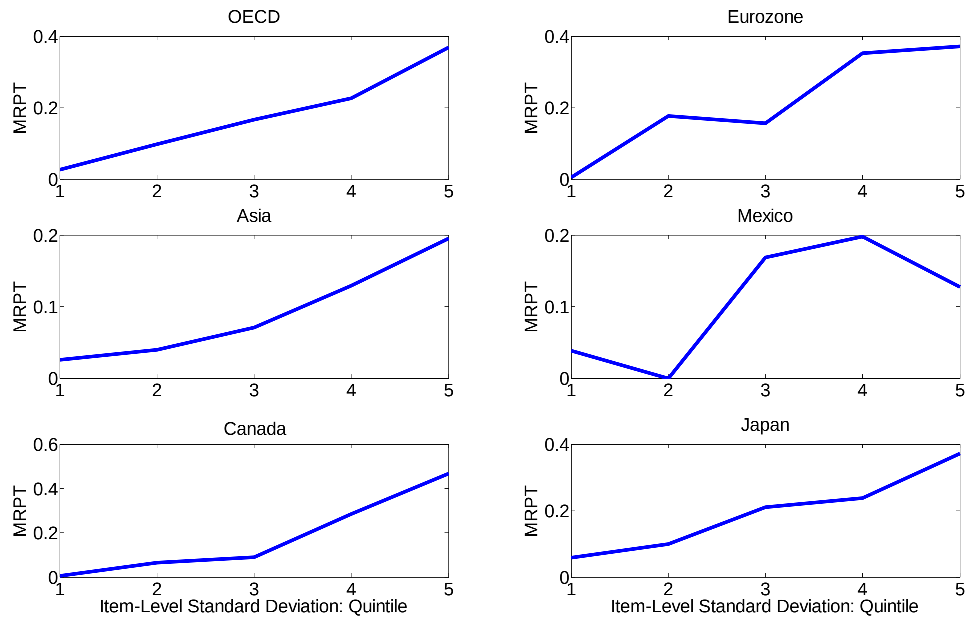

7.2 Additional Month-Level Results

In this section, we perform a variety of robustness checks for our baseline month-level results. First, in the main text, we showed that MRPT was increasing in IQR quintile. However, the standard errors for this bin approach were large due to limited sample sizes. Thus we want to more formally test for the presence of a time-series relationship between price change dispersion and MRPT. We begin by calculating the cross-sectional interquartile range of price changes for each month in our sample. We then split our sample in thirds by the interquartile range. Let I_t^{high} be an indicator for the one-third of months with the highest interquartile range in our sample. Similarly, let I_t^{low} be an indicator for the one-third of months with the lowest interquartile range in our sample. Our baseline time-series specification is then:

$$\Delta p_{i,t} = \left[\beta^{high}\Delta_c e_{i,t} + Z'_{i,t}\ ^{high}\right] I_t^{high} + \left[\beta^{low}\Delta_c e_{i,t} + Z'_{i,t}\ ^{low}\right] I_t^{low} + \epsilon_{i,t}.$$

Table A2 shows that during high dispersion months, MRPT is 21% while in low dispersion months MRPT is only 8%. This difference is both economically and statistically significant, with pass-through more than doubling between low and high dispersion months. Table A2 also shows that these differences remain significant for alternative sample selections as well as alternative measures of cross-sectional dispersion. In addition to the interquartile range, we sort months by the standard deviation of price changes. The interquartile range is more robust to outliers, so we view it as a more reliable benchmark, but using the standard deviation does not change our results. We also split our sample using Census based measures of cross-sectional TFP dispersion from Bloom et al. (2012).

When splitting by census based dispersion measures our results become even more significant, with estimated MRPT more than quadrupling between low and high dispersion months. Finally, Table A2 shows that the difference between high and low dispersion results remains significant when using alternative country restrictions as well as when restricting to a more narrow set of products.

7.3 More on Composition

While we have now shown that our empirical result holds under a variety of controls, these controls may not pick up compositional changes across time in our sample. We discussed various compositional concerns in the body of the text but suppressed the results. We elaborate on the empirical tests here and present the results.

7.3.1 2 Facts or 1 Fact?

First, is our item-level dispersion fact actually distinct from our month-level dispersion fact? Since items are periodically rotated out of the sample, we do not have a balanced panel. Thus, it is possible that the high dispersion time-periods in our data are driven by times when the sample contains items with unusually high price dispersion. To document that our two facts are indeed independent we first combine specifications (3) and (4) to allow for separate effects of cross-item and cross-month dispersion. That is, we estimate

$$\Delta p_{i,t} = \beta^{avg}\Delta_c e_{i,t} + \beta^{Vol}\left(XSD_i \times \Delta_c e_{i,t}\right) + \delta XSD_i + \beta^{IQR} IQR_t \times \Delta_c e_{i,t} + \lambda IQR_t + Z'_{i,t} \ + \epsilon_{i,t} \quad (12)$$

where XSD_i is the standard deviation of item i's price changes and IQR_t is the interquartile range of all price changes in month t. Table A3 shows that both the cross-item effects captured by XSD_i and the cross-month effects captured by IQR_t are highly significant. This conclusion is again robust to a variety of additional controls.

In addition to this double-interaction, Table A4 shows results for a binned regression as in Figures (2) and (3). We split individual items into quintiles by their item-level dispersion of price changes, and then within each item-level quintile we run a time-series regression to estimate the effect of month-level dispersion. Unlike the specification in (12) this "double-binned" regression does not impose linear effects of dispersion and allows the effect of controls to vary across bins. Nevertheless, we again find that both cross-item and cross-month dispersion effects are highly significant. Thus, simple changes in sample composition cannot jointly explain both facts.

7.3.2 Within or Between Sector Phenomena

In general, the positive relationship we observe between pass-through and month-level dispersion could be driven by either movements in dispersion across sectors or dispersion within a sector. While we believe either explanation would be interesting, they would have different implications for models. To address this, we decompose the month-level variance of price changes into a between and within-sector component: $VAR(dp_{i,t}) = VAR(dp_{i,t}^{\text{within sector}}) + VAR(dp_t^{\text{between sector}})$. We then separately interact pass-through with both between and within sector variance:

$$\Delta p_{i,t} = \beta^{avg}\Delta_c e_{i,t} + \beta^{VAR\_W} VAR\_W_t \times \Delta_c e_{i,t} + \beta^{VAR\_B} VAR\_B_t \times \Delta_c e_{i,t} + Z'_{i,t} \ \ + \epsilon_{i,t}$$

Table A5 displays results using both 2-digit and 4-digit sector definitions. For 2-digit sectors, only within-sector variance is significant while for 4-digit sectors both within and between sector variance are significant. We next perform a formal variance decomposition to describe how much of the variation in MRPT is accounted for by within vs. between sector changes in dispersion. The within-sector contribution is given by

$$W = \frac{\left(\beta^{VAR\_W}\right)^2 V\_W}{\left(\beta^{VAR\_W}\right)^2 V\_W + \left(\beta^{VAR\_B}\right)^2 V\_B},$$

where V_W is the time-series variance of within-sector price change dispersion and V_B is the time-series variance of between-sector price change dispersion. Using this decomposition, within-sector variance accounts for 99% of the time-series variation in pass-through using 2-digit sectors and 51% using 4-digit sectors. Thus, even for fairly narrow sectors, the time-series relationship between month-level dispersion and pass-through seem to be largely a within-sector phenomenon.

7.3.3 Are the Items Changing Prices During High Dispersion Months Special?

Are the items that change prices during high dispersion periods the same as the items that change prices during low dispersion periods? We now show that even when restricted to a balanced panel there is a positive relationship between month-level dispersion and MRPT. This means that pass-through for the same products rises with dispersion so our results cannot be explained by a time-varying product mix. Unfortunately, the BLS periodically rotates products in and out of the sample, so it is not feasible to construct a balanced panel that spans the entire length of our data. However, we do have enough data to construct a balanced panel that spans the Great Recession in 2008, which is the most important episode in our sample.

We restrict our analysis to a balanced panel that is in the sample continuously from 2007-2009 and then estimate separate pass-through regressions for 2007, 2008, and 2009. As in the full-sample, month-level dispersion rises dramatically in 2008. The IQR of price changes in 2007 is .08, it rises to 0.122 in 2008 and falls to 0.10 in 2009. Pass-through also exhibits large variation, rising from 0.07 in 2007 to 0.64 in 2008 and falling back to 0.22 in 2009. Time-series variation in dispersion and pass-through for the balanced panel is even stronger and more significant than for our baseline unbalanced specification.

7.3.4 Exchange Rate Appreciations Vs. Depreciations

The 2008 recession was also characterized by an appreciation of the U.S. dollar against most major currencies. Are our pass-through results sensitive to the sign of exchange rate movements? We have re-run all of our empirical specifications restricting our regressions solely to price changes where

$\Delta_c e_{i,t}$ is always positive or always negative. We find that both our month-level and our item-level dispersion MRPT relationships remain highly significant even when restricting only to exchange rate movements of a particular sign.[65] Thus, our results cannot be explained by changes across time in whether the dollar is appreciating or depreciating.

8 Modeling Appendix - Not For Publication

8.1 More General Flexible Price Results

In this section, we show that the intuition from our simple framework in Section 2, survives in a more general framework that allows for general equilibrium effects. Consider the problem of a foreign firm selling goods to importers in the U.S. The firm has perfectly flexible prices that are set in dollars. The optimal flexible price of good i at the border (in logs) can be written as the sum of the gross markup (μ_i), the dollar marginal cost (mc_i) and an idiosyncratic shock (ϵ_i):

$$p_i = \mu_i + mc_i(e_i, \eta_i)$$

Taking the total derivative of equation gives:

$$\Delta p_i = \square\,\square_i(\Delta p_i \,\square\, \Delta p) + \alpha \Delta e_i + \Delta \eta_i$$

which can be rearranged to give:

$$\Delta p_i = \frac{1}{1+\square_i}\left[\alpha \Delta e_i + \square_i \Delta p + \Delta \eta_i\right]$$

In Section 2 we explored the case when all indirect GE effects were shut off ($\Delta p = 0$). Here, we include them to show that most of the simple intuition between about the positive relationship between MRPT and dispersion survives the introduction of GE effects. The above equation can be rearranged to give the simple pass-through equation:

$$\frac{\Delta p_i}{\Delta e_i} = \frac{\alpha_i}{1+\square_i} + \frac{\square_i}{1+\square_i}\frac{\Delta p}{\Delta e_i} \tag{13}$$

We can do some comparative statics to see how parameters affect pass-through

$$\frac{\partial \frac{\Delta p_i}{\Delta e_i}}{\partial \alpha} = \frac{1}{1+\square_i} > \ 0$$

[65] For brevity we do not report these results, but they are available upon request.

$$
\begin{aligned}
\frac{\partial \frac{\Delta p_i}{\Delta e_i}}{\partial \square_i} &= \square \frac{\alpha_i}{(1+\square_i)^2} + \frac{1}{(1+\square_i)^2}\frac{\Delta p}{\Delta e_i} \\
&= \frac{\frac{\Delta p}{\Delta e_i} \square \alpha_i}{(1+\square_i)^2} < 0 \text{ if } \alpha_i > \frac{\Delta p}{\Delta e_i}
\end{aligned} \tag{14}
$$

As before, an upper bound on the level of pass-through is given by what fraction of marginal costs are denominated in units of the foreign currency, α_i. The higher this share, the higher the potential exchange rate pass-through. General equilibrium effects operating through the domestic price level do affect the comparative static with respect to the mark-up elasticity. All of things equal, if the mark-up elasticity is higher, then less of the exchange rate shock is passed into prices, which lowers $\frac{\Delta p_i}{\Delta e_i}$. This is the first term in equation (14). However, this is now an additional effect: a higher $\square_i$ means that individual prices are more sensitive to changes in the aggregate price level because strategic complementarities are higher. This is the second term in equation (14). This term is positive because $\frac{\Delta p}{\Delta e_i} > 0$ since increases in foreign marginal costs also raise the domestic price level. The total effect is ambiguous in general. However, for realistic cases (for instance all the parameter values we consider in our model), $\alpha_i > \frac{\Delta p}{\Delta e_i}$. To see this, remember that α_i is the fraction of marginal cost that is denominated in foreign currency. This gives an upper bound on the level of pass-through to individual prices from exchange rate shocks. It is hard to see how pass-through to the overall price level can be bigger than that effect since not all goods domestically are affected by the exchange rate shock and the overall-passthrough rate is affected by the level of strategic complementarities, $\square_i$, which lowers the level of pass-through.

We now show that changes in parameters that increase pass-through also increase the variance of price changes. The variance of price changes is given by:

$$
\begin{aligned}
var(\Delta p_i) &= \left(\frac{\alpha_i}{1+\square_i}\right)^2 var(\Delta e_i) + \left(\frac{\square_i}{1+\square_i}\right)^2 var(\Delta p) + \left(\frac{1}{1+\square_i}\right)^2 var(\Delta \eta_i) \\
&\quad + \frac{\alpha_i \square_i}{(1+\square_i)^2} cov(\Delta e_i, \Delta p) + \frac{\alpha_i}{(1+\square_i)^2} cov(\Delta e_i, \Delta \eta_i) + \frac{\square_i}{(1+\square_i)^2} cov(\Delta p, \Delta \eta_i)
\end{aligned}
$$

But the last terms are zero by assumption that idiosyncratic shocks are orthogonal to exchange rate shocks and will wash out in aggregate so that they do not affect the aggregate price level. This implies that

$$
var(\Delta p_i) = \left(\frac{\alpha_i}{1+\square_i}\right)^2 var(\Delta e_i) + \left(\frac{\square_i}{1+\square_i}\right)^2 var(\Delta p) + \left(\frac{1}{1+\square_i}\right)^2 var(\Delta \eta_i) + \frac{\alpha_i \square_i}{(1+\square_i)^2} cov(\Delta e_i, \Delta p) \tag{15}
$$

Using this expression, we get that

$$
\frac{\partial var(\Delta p_i)}{\partial \square_i} = \square \frac{2\alpha_i^2}{(1+\square_i)^3} var(\Delta e_i) + \frac{2\square_i}{(1+\square_i)^3} var(\Delta p) \square \frac{2}{(1+\square_i)^3} var(\eta_i) + \frac{\alpha_i(1 \square \square_i)}{(1+\square_i)^3} cov(\Delta e_i, \Delta p). \tag{16}
$$

We now show that under a mild and empirically realistic restriction, the variance of price changes is declining in $\square_i$. Empirically, we know that the variance of idiosyncratic price changes is an order of magnitude larger than the variance of aggregate price changes and exchange rate movements. With this in mind, we impose the restriction that

$$var(\Delta p_i) > var(\Delta e_i) + var(\Delta p).$$

We can substitute this restriction into (15) to get that

$$\left(\frac{\alpha_i}{1+\square_i}\right)^2 var(\Delta e_i) + \left(\frac{\square_i}{1+\square_i}\right)^2 var(\Delta p) + \left(\frac{1}{1+\square_i}\right)^2 var(\Delta \eta_i) + \frac{\alpha_i \square_i}{(1+\square_i)^2} cov(\Delta e_i, \Delta p) > var(\Delta e_i) + var(\Delta p)$$

or

$$var(\eta_i) > \left[(1+\square_i)^2 \square \square_i^2\right] var(\Delta p) + \left[(1+\square_i)^2 \square \alpha_i^2\right] var(\Delta e_i) \square \alpha_i \square_i cov(\Delta e_i, \Delta p) \tag{17}$$

Using (16) we have

$$\begin{aligned}
\frac{\partial var(\Delta p_i)}{\partial \square_i} &= \square \frac{2\alpha_i^2}{(1+\square_i)^3} var(\Delta e_i) + \frac{2\square_i}{(1+\square_i)^3} var(\Delta p) \square \frac{2}{(1+\square_i)^3} var(\eta_i) + \frac{\alpha_i(1 \square \square_i)}{(1+\square_i)^3} cov(\Delta e_i, \Delta p) \\
&\propto \square 2\alpha_i^2 var(\Delta e_i) + 2\square_i var(\Delta p) \square 2var(\eta_i) + \alpha_i(1 \square \square_i) cov(\Delta e_i, \Delta p)
\end{aligned}$$

Substituting the inequality (17) for $var(\eta_i)$ gives

$$\begin{aligned}
\frac{\partial var(\Delta p_i)}{\partial \square_i} &< \square 2\alpha_i^2 var(\Delta e_i) + 2\square_i var(\Delta p) + \alpha_i(1 \square \square_i) cov(\Delta e_i, \Delta p) \\
&\quad \square 2\left[(1+\square_i)^2 \square \square_i^2\right] var(\Delta p) \square 2\left[(1+\square_i)^2 \square \alpha_i^2\right] var(\Delta e_i) + 2\alpha_i \square_i cov(\Delta e_i, \Delta p) \\
&= \square 2\left[(1+\square_i)^2 \square \square_i^2 \square \square_i\right] var(\Delta p) \square 2\left[(1+\square_i)^2\right] var(\Delta e_i) + \alpha_i \left[\square_i + 1\right] cov(\Delta e_i, \Delta p) \\
&< \square 2\left[(1+\square_i)^2 \square \square_i^2 \square \square_i\right] var(\Delta p) \square 2\left[(1+\square_i)^2\right] var(\Delta e_i) + \alpha_i \left[\square_i + 1\right] var(\Delta e_i) \\
&< \square 2\left[(1+\square_i)^2 \square \square_i^2 \square \square_i\right] var(\Delta p) \square 2\left[(1+\square_i)^2\right] var(\Delta e_i) + (1+\square_i)^2 var(\Delta e_i) \\
&= \square 2\left[(1+\square_i)^2 \square \square_i^2 \square \square_i\right] var(\Delta p) \square \left[(1+\square_i)^2\right] var(\Delta e_i) \\
&< 0
\end{aligned}$$

The second inequality uses the result that Δp moves less than one for one with the exchange rate.

In sum, even in the case when indirect GE effects are allowed, our central theoretical prediction still holds: changes in parameters that increase exchange rate pass-through ($\alpha_i \uparrow$, $\square_i \downarrow$) also increase the variance of price changes.

8.2 The Role of Measurement Error

As mentioned in the introduction as well as in Section 3.5, measurement error is a potential concern for our empirical estimates. We attempted to address this concern in our empirical results by estimating various alternative pass-through specifications. While time-series variation in these alternative pass-through specifications is less interpretable as time-variation in price flexibility, these specifications have the advantage of reducing measurement error. Since time-series variation in our benchmark MRPT is more easily interpretable, we now assess the extent to which measurement error is indeed a serious concern for this empirical specification. To do this, we use our model to simulate three sources of potential measurement error and show that such errors cannot explain our results.

We model three sources of measurement error that are likely to be important in the BLS data: 1) Errors in aligning the timing of measured price changes with the timing of exchange rates. 2) Mis-reported prices. 3) Failure to report actual price changes.

Prices are recorded in the BLS at the time they are received rather than at the time they are ordered. Production and delivery lags mean that this price may have been set several periods in the past, under a different prevailing exchange rate.[66] To model this timing error, we assume that while the price at time t is set using information on the exchange rate at time t, the price is reported at time $t+x$ where $x \sim U[0, X]$. That is, there is a potential mismatch between the exchange rate that is actually relevant for a firm's pricing decision and the exchange rate at the time a price is reported. The left hand column of Figure 11 shows the effects of timing errors on pass-through and price change dispersion as X is varied between 0 and 6 months. As X increases, measured pass-through falls as there is additional attenuation bias in the MRPT regression. However, measured price change dispersion is not affected. This is because mismeasuring the timing of price changes has no effect on their measured size.

Thus, changes in timing errors could only explain the time-series relationship between price change dispersion and pass-through if there was some common factor that increased the dispersion of price changes at the same time that delivery lags fall. We can roughly assess this possibility by examining the composition of trade across time. Using data from the U.S. Census, we can compute the fraction of goods shipped by ocean vessels. These items are likely to have the longest delivery lags, so it would be concerning if the fraction of items shipped by vessel negatively comoved with the dispersion of price changes. However, we find that there is a positive correlation of 0.13 between the fraction of items shipped by ocean vessel and the month-level interquartile range of price changes. Thus, if anything, changes in the composition of trade across time would work against our empirical results. In the appendix we provide additional discussion of trade composition and evidence that this does not drive our results.

In addition to timing error, we allow for reporting errors by assuming that recorded price changes

[66] See Nakamura and Steinsson (2012) for additional discussion.

Figure 11: Simulating Sources of Measurement Error

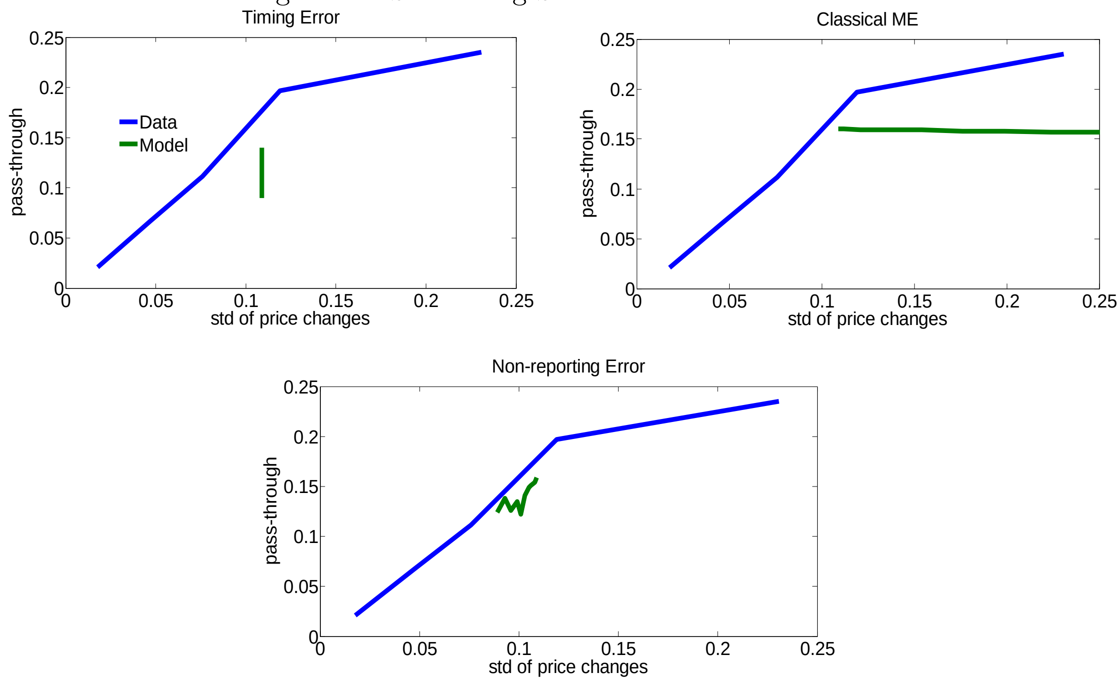

are equal to the true price plus classical measurement error. The second column of Figure 11 shows results for measurement error standard deviations ranging from 0 to 0.18. Increases in measurement error can dramatically increase the dispersion of measured price changes. However, greater measurement error leads to a decline in pass-through due to standard attenuation bias. Thus, classical measurement error is unable to explain our results.

Finally, we assume that when a price change actually occurs it is only recorded with some probability < 1. The third column of Figure 11 simulates results for non-reporting probabilities ranging from 0 to 0.9. Even huge non-reporting errors barely affect either pass-through or price change dispersion. They do not affect pass-through because once a price change is actually measured, it will reflect all of the previous pass-through that was not recorded. Furthermore, non-reporting does not affect the dispersion of price changes as long as the probability of a price change not being reported is independent of the size of the price change. The one statistic that declines dramatically with non-reporting error is the frequency of adjustment. Thus, if non-reporting error were a cause of concern for our results, this explanation would need to contend with the much smaller frequency pass-through relationships observed empirically.

Appendix Tables

Table A1:Interaction Specification - Cross-Item: Robustness

	Average pass-through		Volatility		Frequency		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
At least 3 price changes	0.15	0.02	0.05	0.02			88964	0.06
	0.15	0.01	0.05	0.02	0.01	0.01	88964	0.06
Using trade-weighted broad xrate	0.41	0.03	0.26	0.04			87383	0.08
	0.44	0.03	0.21	0.04	0.27	0.03	87383	0.08
Using trade-weighted major country xrate	0.28	0.02	0.21	0.03			96512	0.07
	0.29	0.02	0.18	0.03	0.15	0.02	96512	0.07
Placebo num changes	0.15	0.01	0.00	0.01			100871	0.09
	0.15	0.02	-0.00	0.01	0.02	0.01	100871	0.09
Placebo num obs	0.15	0.02	-0.00	0.01			100871	0.09
	0.15	0.02	-0.00	0.01	0.02	0.01	100871	0.09
Median regression	0.16	0.00	0.07	0.00			95284	
	0.16	0.00	0.07	0.00	0.01	0.00	95284	

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes

Table A2: "Binned" time series results

	High volatility		Low volatility		Difference		N_{obs}	R^2
	β^{high}	$se(\beta^{high})$	β^{low}	$se(\beta^{low})$	$\beta^{high}\square\beta^{low}$	t-stat		
All countries, all items ex petroleum								
- Interquartile range	0.21	0.03	0.08	0.02	0.12	4.35	62395	0.09
- Cross-sectional std	0.17	0.02	0.08	0.02	0.10	3.89	63095	0.09
- Bloom uncertainty	0.26	0.03	0.06	0.02	0.20	6.33	64204	0.08
OECD, all items ex petroleum								
- Interquartile range	0.22	0.03	0.12	0.03	0.10	2.88	34790	0.09
- Cross-sectional std	0.22	0.03	0.10	0.03	0.12	3.17	35288	0.09
- Bloom uncertainty	0.25	0.03	0.07	0.03	0.18	4.37	35398	0.09
All countries, all manufact. items								
- Interquartile range	0.18	0.03	0.08	0.02	0.10	3.41	51157	0.11
- Cross-sectional std	0.15	0.02	0.08	0.02	0.07	2.75	51943	0.11
- Bloom uncertainty	0.23	0.02	0.05	0.02	0.18	5.83	52661	0.10

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes

Table A3: Interaction Specifications with both Cross-Item and Cross-Month Dispersion

	Average pass-through		Item-Level Volatility		Month-Level Volatility		N_{obs}	R^2
Controls	β^{avg}	$se(\beta^{avg})$	β^{XSD}	$se(\beta^{XSD})$	β^{IQR}	$se(\beta^{IQR})$		
- No additional controls	.141	.013	.043	.017	.060	.009	95284	.072
- Item level frequency	.139	.013	.041	.017	.060	.009	95284	.072
- Aggregate frequency	.137	.013	.041	.017	.060	.009	95284	.073
- Time trend + Month	.137	.024	.042	.017	.055	.012	95284	.076
- All above controls	.125	.024	.042	.017	.055	.012	95284	.077

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents 2 to 4-digit sectoral harmonized codes. The sample is all countries, all items ex petroleum.

Table A4: Month-Level IQR Regressions by Item-Level XSD Quintiles

Item-Level XSD Quintile	$\beta^{high\ IQR}$	$\beta^{low\ IQR}$	$\beta^{high\ IQR}$□$\beta^{low\ IQR}$	t-stat	n	R^2
1 (Lowest)	.035	.029	.005	0.29	6096	.64
2	.083	.052	.031	1.46	12522	.24
3	.133	.053	.079	2.24	16630	.15
4	.277	.127	.150	3.41	16470	.13
5 (Highest)	.417	.112	.304	2.92	10942	.12

Table A5: Within and Between

Sector Definition	β^{avg}	β^{VAR_W}	t-stat W	β^{VAR_B}	t-stat B
2-digit	.141	.056	5.95	.010	0.82
4-digit	.141	.036	3.29	.034	2.59

Within is month-level variance of price changes within sectors. Between gives variance of inflation rates across sectors