NBER WORKING PAPER SERIES

SHOCKS VS. RESPONSIVENESS:
WHAT DRIVES TIME-VARYING DISPERSION?

David Berger
Joseph Vavra

Working Paper 23143
http://www.nber.org/papers/w23143

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2017

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 the editor and anonymous referees as well as 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, Gita Gopinath, 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.

Shocks vs. Responsiveness: What Drives Time-Varying Dispersion?
David Berger and Joseph Vavra
NBER Working Paper No. 23143
February 2017
JEL No. E10,E3,E31,E52,F3,F31

ABSTRACT

The dispersion of many economic variables is countercyclical. What drives this fact? Greater dispersion could arise from greater volatility of shocks or from agents responding more to shocks of constant size. Without data separately measuring exogenous shocks and endogenous responses, a theoretical debate between these explanations has emerged. In this paper, we provide novel identification using the open-economy environment: using confidential BLS microdata, we document a robust positive relationship between exchange rate pass-through and the dispersion of item-level price changes. We show this relationship arises naturally in models with time-varying responsiveness but is at odds with models featuring volatility shocks.

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

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

1 Introduction

The cross-sectional dispersion of many economic variables is countercyclical, but there is much debate over the source of this empirical phenomenon.[1] This is because existing research measures the dispersion of *endogenous* variables, which will reflect some combination of exogenous shocks and firms' optimal responses to those shocks. As such, greater dispersion of endogenous outcomes could occur because exogenous shocks get bigger (what we refer to as greater volatility) or because firms respond more to shocks which are the same size (what we refer to as greater responsiveness).[2]

With only data on outcomes and not the separate contributions of exogenous shocks and endogenous responses, a theoretical debate between these explanations has emerged. Many models such as Bloom et al. (2012) and Vavra (2014) assume that firms draw exogenous idiosyncratic shocks with time-varying volatility in order to generate time-variation in dispersion. On the other side of the debate, papers such as Bachmann and Moscarini (2012), Ilut et al. (2014), Baley and Blanco (2016) and Munro (2016) propose a more varied set of mechanisms such as learning, ambiguity aversion, incomplete information and customer search to generate variation in dispersion through the responsiveness channel. Resolving this debate is important for understanding the nature of business cycles and shocks. At stake is the empirical viability of one potential source of economic fluctuations: exogenous changes in volatility have been proposed as a possible cause of business cycles. If it is responsiveness which is instead reacting to the cycle, then time-varying dispersion is merely a symptom of business cycles arising from some other source.

Making empirical progress differentiating time-varying responsiveness from time-varying volatility is difficult, but we show that the open economy environment can be used to provide identification. This is because it provides a large and observable cost shock, the nominal exchange rate, which can be used to differentiate between these channels. In the first half of the paper, we use confidential BLS import price data to document that item-level price change dispersion is both countercyclical and highly correlated with exchange rate pass-through. In the second half of the paper, we use a workhorse open-economy model to show that these facts strongly support time-varying responsiveness over time-varying volatility. The intuition is straightforward: increasing responsiveness increases both dispersion and pass-through. In contrast, when volatility increases, dispersion increases but pass-through actually declines as price changes become dominated by idiosyncratic forces. Indeed, we estimate our model using indirect inference and show that time-varying responsiveness can match a variety of facts in BLS micro data while time-varying volatility is strongly rejected. In more detail, our paper proceeds in three steps.

First, we start by documenting new facts. We begin by showing that, like many other economic outcomes, the dispersion of item-level price changes in BLS import price data is strongly countercyclical. For example, Figure 1 shows that the interquartile (IQR) range of price changes in our data moves substantially across time and exhibits a strong negative correlation with real GDP growth. Second, the dispersion of price changes is highly correlated with exchange rate pass-through. As a simple illustration of this fact, we divide our entire sample into 8-month long windows and compute the IQR of price

[1] Countercyclical dispersion is found in Bloom (2009) (sales growth), Bloom et al. (2012) (revenue TFP and employment growth), and Vavra (2014) (prices). Bachmann and Bayer (2014) finds procyclical dispersion of investment rates, but as we discuss in footnote 20, this is highly consistent with our results since their measure includes zeros.

[2] To avoid confusion, we distinguish between "dispersion" and "volatility" throughout the paper. We define dispersion as the spread of endogenous outcome variables, while volatility is the spread of exogenous shocks.

Figure 1: Price Change Dispersion is Countercyclical

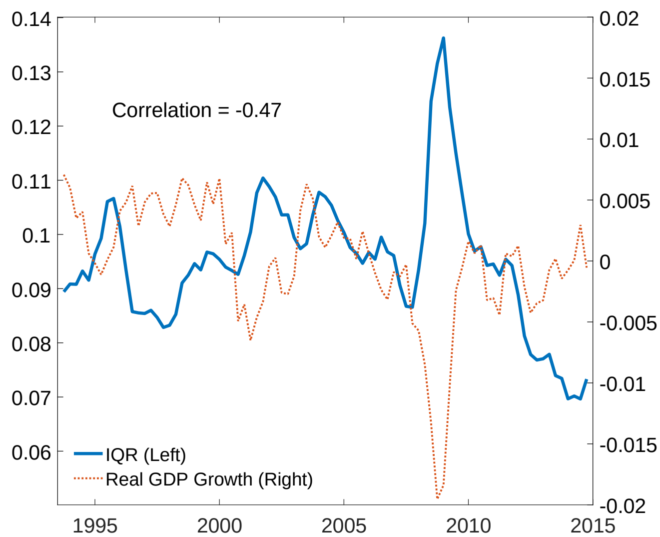

This figure shows the IQR (75-25 range) of all non-zero price changes in our benchmark sample, described below, and chained real GDP growth from 1993m10 to 2015m1. The monthly IQR is averaged quarterly for consistency with GDP measures, and both series are smoothed with a 3-period moving average.

changes and our preferred measure of exchange rate pass-through separately in each window.[3] The resulting scatter plot of IQR vs. pass-through in Figure 2 shows the strong positive relationship between these variables. Since the time-series graph in Figure 1 shows that the IQR in the Great Recession is an outlier, we include Great Recession observations separately as blue squares in Figure 2 to show that the pass-through-IQR relationship is not driven by this single recession.

Figure 2: Dispersion vs. pass-through

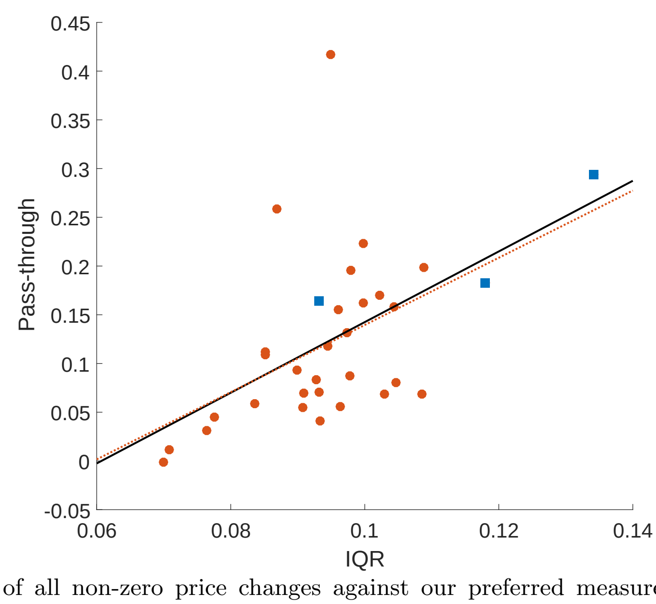

This figure shows the IQR of all non-zero price changes against our preferred measure of exchange rate pass-through, described below. Both statistics are computed separately in a series of 8-month disjoint windows which span our sample period. Windows which have a majority of months during the Great Recession, as defined by NBER, are shown in blue. The black regression line includes all observations while the red-dotted line excludes Great Recession observations. The regression with all observations has a slope coefficient of 3.625, t-stat of 3.36 and R^2 of 0.27.

[3] Appendix Figure A.1 shows similar patterns hold for different window lengths.

This positive relationship between price change dispersion and pass-through is extremely robust. We show it is not driven by changes in the frequency of price adjustment, secular trends, changes in exchange rate volatility, by particular products or countries, or by mechanical reverse causation, and it holds under a variety of alternative specifications designed to deal with certain misspecification concerns.[4] We also show that although price change dispersion is countercyclical, our patterns reflect a dispersion-pass-through relationship not a business cycle-pass-through relationship. In particular, the positive relationship between dispersion and pass-through holds after controlling for various business cycle indicators. It also holds at the sector-level after controlling for time dummies to flexibly absorb common aggregate variation, and similar results arise at the individual item level: items which exhibit disperse price changes across time also exhibit high exchange rate pass-through when they change prices. Together these results allay concerns about spurious correlation and the effects of confounding aggregate shocks.

If one views exchange rate pass-through as a simple reduced form measure of responsiveness, our empirical results then suggest an important link between countercyclical dispersion and responsiveness. However, such "suggestive" evidence should be viewed with caution. While pass-through is a widely computed moment in the open-economy literature, interpreting this moment and its relationship with dispersion requires imposing additional structure.[5]

In the second part of the paper, we move in this direction by adopting the flexible price framework of Burstein and Gopinath (2014). In this model, the mapping from structural parameters to observables is straightforward, which allows us to starkly illustrate the nature of the identification problem as well as its solution. With flexible prices, the dispersion of price changes across firms is determined by two parameters: i) the volatility of idiosyncratic shocks and ii) the response of optimal prices to shocks. This means that changes in dispersion could be explained by changes in either parameter. However, these parameters have very different implications for exchange rate pass-through. Increasing volatility increases price change dispersion but has no effect on pass-through, since optimal pass-through is scale-invariant.[6] In contrast, increasing responsiveness simultaneously increases dispersion *and* pass-through. Thus, in this simple flexible price environment, changes in responsiveness can explain the positive relationship between dispersion and pass-through observed in the data, while changes in volatility cannot.

In the final part of the paper, we turn to a quantitative environment with more realistic pricing frictions. In this environment, the mapping from structural parameters to observables is more complicated. However, we show that price frictions only amplify our previous conclusions: while increases in responsiveness continue to increase both dispersion and pass-through, increases in volatility increase dispersion but lead to a counterfactual decrease in pass-through. This is because when the volatility of idiosyncratic shocks increases, exchange rate movements become less relevant for optimal price adjustment and measured pass-through (conditional on adjusting a firms price) declines.[7]

[4]In particular, one might be concerned that this positive correlation reflects mechanical reverse causality whereby increases in pass-through make prices more sensitive to exchange rate shocks and increase price change dispersion. However, we show in our quantitative results that this mechanical effect on variance is completely negligible and plays no role in our results. The intuition is that in the data the variance of price changes is orders of magnitude larger than the variance of exchange rate changes, which means that the variance of price changes is dominated by idiosyncratic shocks and so changing the sensitivity to exchange rates has essentially no effect on overall price change dispersion. This also means that changes in the volatility of exchange rates across time have little effect on dispersion.

[5]See e.g. Burstein and Gopinath (2014) for a detailed discussion of the mapping between pass-through regressions and a variety of commonly used models of incomplete pass-through as well as associated pitfalls.

[6]That is, doubling the size of a cost shock doubles the optimal price change.

[7]More formally, as we show in section 4.2 state-dependent pricing implies an upward statistical "selection bias" in our

In addition, this quantitative model provides a laboratory which we use to explicitly test the validity of our earlier empirical methodology. In particular, there are valid concerns that our empirical patterns might be driven by censoring, small samples, sample turnover, or misspecification. We address these concerns head-on by formally estimating our quantitative model using indirect inference to match our empirical regressions. In doing so, we simulate data from our model, replicate BLS sample sizes and sampling and then run regressions on this simulated data identical to those in our empirical work. This indirect inference estimation procedure allows us to rule out many potential concerns with our empirical results, since the same biases should arise when running regressions on simulated and actual data. Put differently, as usual with indirect inference, identification does not require our empirical regressions to be correctly specified or have any structural interpretation in the true model. It merely requires that changes in structural parameters manifest themselves distinctly in our reduce form regressions, and this is indeed the case: our estimation formally rejects variation in volatility in favor of variation in responsiveness.

What does variation in responsiveness represent? As described above, a variety of mechanisms have been proposed to endogenously generate countercyclical dispersion. In Appendix B, we show that the forces in these models all map to the same responsiveness parameter that is key to our qualitative results.[8] This means that our empirical strategy rejects volatility shocks in favor of responsiveness shocks, but it cannot isolate a particular model of responsiveness. However, this also means that our *qualitative* insights do not require taking a stand on a particular model of responsiveness, and in reality it seems likely that many of these responsiveness forces coexist and complement each other.

Moving from qualitative results to quantitative results necessarily comes with some tradeoff: in order to formally estimate our model and reach conclusions about magnitudes, we must impose more structure on the data generating process. Since it is infeasible to simultaneously include all potential mechanisms that can generate countercyclical responsiveness, we focus solely on variation in responsiveness which arises from movements in the "super-elasticity" of demand in Kimball's preferences.[9] We focus on this source of responsiveness shocks for three reasons: 1) This is an extremely standard specification in the open economy literature for generating incomplete pass-through, which has been used to rationalize a number of related cross-sectional pass-through facts. It is important that matching our new facts not come at the cost of missing existing results. 2) It is parsimonious and straightforward to solve numerically. The full version of our quantitative model includes four aggregate and two idiosyncratic states and requires global solution methods in an equilibrium environment, so estimation would be infeasible in more complicated settings such as those with learning and incomplete information. 3) It fully nests existing models in the literature, such as Vavra (2014), which explain countercyclical dispersion via countercyclical volatility, and so it gives these models equal footing in matching our new empirical evidence.

It is difficult to directly assess the plausibility of our estimated super-elasticity shocks, since no empirical estimates of this statistic exist.[10] However, we show that these relatively simple shocks produce observable series which are quite reasonable in many dimensions. In particular, after picking exchange

exchange rate pass-through regression since firms are more likely to adjust when the exchange rate and idiosyncratic shock reinforce each other. As idiosyncratic volatility rises, this bias declines and our measure of pass-through falls.

[8]More broadly, any force which changes the response of firms' desired prices to cost shocks rather than the size of cost shocks should deliver similar implications.

[9]That is the elasticity of the elasticity of demand with respect to a firm's relative price. See Klenow and Willis (2006) for the first use of this terminology.

[10]Note that we would require time-series estimates of super-elasticity. Given the controversy over the cyclicality of the elasticity of demand, it is not surprising that no estimates exist measuring the behavior of the super-elasticity across time.

rates in our model to match data from 1993-2015, we show that we are able to well-match the behavior of the IQR, overall inflation, import inflation, output growth and adjustment frequency over our 1993-2015 sample with only super-elasticity and nominal demand shocks. The model also matches time-series variation in exchange rate pass-through which is not directly targeted, and it generates markup movements across time which are relatively small and well-within the range of estimates in the literature. Thus, we conclude that even though super-elasticity movements cannot be directly measured in the data, such shocks produce reasonable results along dimensions which are observable.[11]

It is important to note that our analysis focuses on import prices, so one should be cautious when extrapolating to other contexts. However, several recent papers reach similar conclusions in other environments and suggest that our conclusions indeed have external validity. First, Fleer et al. (2015) extends our analysis from imports to broader consumer prices in Switzerland and finds similar results. Moving beyond prices, Ilut et al. (2014) shows that individual firms' employment responds more to idiosyncratic TFP changes during recessions, and Decker et al. (2016) shows that secular reallocation trends in US manufacturing are driven by changes in responsiveness rather than changes in the volatility of shocks. Finally, a growing literature argues for cyclical changes in market structure and demand which should lead to exactly the sort of time-varying responsiveness necessary to explain our results.[12]

Our paper is related to some recent empirical work trying to determine if aggregate volatility is a source of, or response to, business cycles. These papers study aggregate time-series volatility rather than cross-sectional dispersion, and so they use identification strategies which focus on relationships between aggregate variables. This makes them quite distinct from our micro data based strategy.[13] For example, Baker and Bloom (2013) uses natural disasters to instrument for stock market first and second moments in order to assess their independent effects on GDP growth, and Ludvigson et al. (2016) and Berger et al. (2016) use time-series VAR strategies to explore similar questions of causality.

Within the pass-through literature, we are most closely related to Gopinath and Itskhoki (2010). The most important distinction is that we focus on time-series variation in pass-through and explore its link to the dispersion of price changes. In contrast, their paper focuses on long-run differences in pass-through across different items and links this to the frequency of price adjustment. Our focus on time-series variation leads us to alternative empirical specifications which are better suited for this purpose as well as to the estimation of a model with a variety of aggregate shocks. Nevertheless our results are broadly complementary, and we ultimately find that similar structural forces can help to jointly explain both their cross-sectional and our time-series evidence.

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 October 1993-January 2015. This data

[11] It is also useful to note that rejecting volatility shocks does not necessarily require one to accept our particular formulation of responsiveness variation. A model without either super-elasticity or volatility shocks actually does better than a model with volatility shocks at matching empirical dispersion-pass-through relationships.

[12] See e.g. Stroebel and Vavra (2016), Munro (2016) and Kaplan and Menzio (2016).

[13] Bachmann et al. (2016) explore a useful related micro exercise looking at the behavior of investment expectation errors, but since investment is highly endogenous their evidence cannot distinguish changing responsiveness from changing volatility.

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 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."[14]

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.[15]

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 ensuring high data quality. In the first step of data collection, the BLS negotiates with the company over the number of price quotes reported to limit the reporting burden. The BLS also contacts a respondent if the reported price is unchanged 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. Finally, simple forms of measurement error would, if anything, work against our finding.[16]

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.[17] 2) We require that a good have two price changes during its life so that we can measure pass-through of cumulative exchange rate movements over a completed spell into the item's new price.[18] 3) We restrict attention to imports whose prices are invoiced in dollars rather than in foreign currency. We use data from all countries and all products, however we exclude commodities since these items have little market power. 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.

[14] This example is taken from Gopinath and Rigobon (2008).

[15] 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.

[16] In a frictionless price environment, it is straightforward to show that measurement error in the exchange rate leads to a negative relationship between pass-through and dispersion, while measurement error in price changes and thus dispersion attenuates any fundamental relationship. We have confirmed in our quantitative model that similar results obtain in the presence of nominal frictions.

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

[18] Some alternative pass-through specifications we explore allow us to relax this requirement and it does not change our results. We also show that all results are robust to only including items with many price changes.

(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.

Overall these sample choices conform with the now large literature studying exchange rate pass-through from a micro perspective. In Appendix A we provide further statistics on the properties of our benchmark sample and additional information on each cut of the data. More importantly, we show that our results are robust to a variety of alternative sample selection criteria.

2.2 Measuring Dispersion and pass-through

Our primary dispersion measure is the interquartile range (IQR) of all non-zero log price changes in a given month.[19] The IQR is robust to outliers and has been widely used in the literature, but we show throughout that all results are similar when using other measures of dispersion such as the standard deviation of price changes. Since this dispersion measure varies across months as the distribution of price changes moves, we refer to it as "month-level dispersion". Measuring dispersion excluding zeros helps to isolate mechanical effects of frequency from changes in the price change distribution conditional on adjustment, and is ultimately crucial for identification.[20] More specifically, since the frequency of adjustment is low, increasing it leads to an increase in dispersion of price changes when zeros are included, even if price always change by the same amount when adjusting. Measuring dispersion excluding zeros and frequency separately allows us to explore their independent relationships with pass-through.

While we focus on changes in month-level dispersion, we also document that similar dispersion-pass-through relationships hold across items by calculating what we call "item-level" dispersion: the standard deviation of all non-zero price changes for a particular item across time.[21]

Our benchmark measure of exchange rate pass-through is standard in the literature. 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. Specifically, we estimate the following regression on adjusting prices:

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

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.[22] We estimate this regression with country and sector fixed effects.[23] The coefficient β measures the fraction of cumulated exchange rate movements "passed-through" to an item's price when adjusting.[24]

[19] Similar results obtain if we calculate the average IQR within sectors instead of across all price changes.

[20] See Bachmann and Bayer (2014) for related discussion. They show that the standard deviation of investment rates including zeros is procyclical while the standard deviation conditional on "spike" adjustment is countercyclical. This difference is driven by changes in the frequency of adjustment spikes. Overall, their conclusions are highly consistent with other patterns of counteryclical dispersion and could similarly be driven by either changing volatility or responsiveness.

[21] We use the standard deviation since items typically have a small number of price changes and IQR is undefined.

[22] 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.

[23] 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.

[24] Holding the frequency of price adjustment constant, a decline in β thus implies that the real exchange rate moves more strongly with the nominal exchange rate.

The use of MRPT, which conditions on price adjustment, is important for our identification argument: in our quantitative results we show that increasing responsiveness increases dispersion, frequency and MRPT. In contrast, increasing volatility increases dispersion and frequency but not MRPT. If one instead measures pass-through without conditioning on price adjustment, the increase in frequency when volatility rises *will* cause pass-through to rise. This means that pass-through specifications which do not condition on price adjustment cannot disentangle greater volatility from greater responsiveness.

The results from estimating average pass-through for the entire sample using (1) are shown in column 1 of Table 1. Consistent with prior literature, we find that average MRPT is low. When a price changes, it passes through only 0.154% of a 1% change in the nominal exchange rate.[25]

2.3 Baseline Results

Figure 1 shows that price change dispersion varies substantially across time. However, as discussed above, it is impossible to tell from Figure 1 whether this variation is driven by changes in the volatility of exogenous shocks or in the endogenous responsiveness to those shocks. We now document the central empirical fact of our paper, that we show in subsequent sections allows us to discriminate between these explanations: time periods with greater price change dispersion also exhibit greater exchange rate pass-through. To test for a time-series relationship between price change dispersion and MRPT we begin by splitting our sample into quintiles by the value of IQR_t and then estimate equation (1) separately using only observations in each quintile. Figure 3 shows that pass-through more than quadruples from the lowest quintile of month-level dispersion to the highest quintile.[26]

Figure 3: Dispersion vs. pass-through

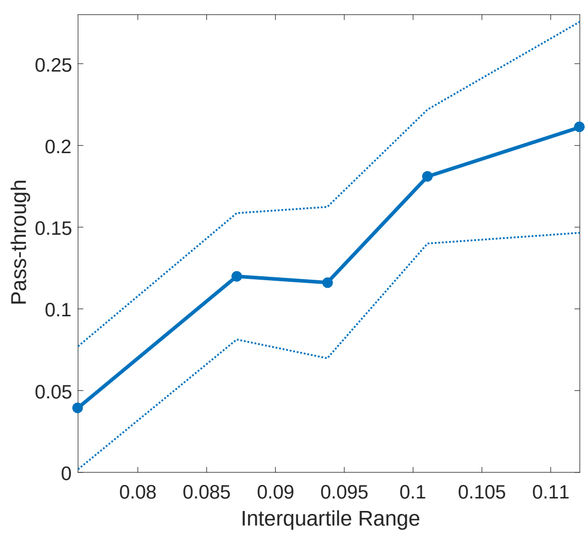

This figure shows separate estimates of regression (1) in each of 5-quintiles by the value of IQR_t. All regressions have country $\times$ PSL fixed effects and robust standard errors are clustered at the country $\times$ PSL level. We also include controls for foreign CPI growth, US gdp growth and US CPI growth. 95% confidence intervals are shown with dotted lines, and the average IQR value in each quintile is shown on the x-axis.

[25]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.

[26]Figure A.2 in Appendix A.1 shows that estimating this binned-pass-through relationship more non-parametrically using 100 overlapping bins produces extremely similar results.

Of course, the months in each IQR bin differ from each other in many ways besides their month-level dispersion. To what extent is the positive relationship between pass-through and dispersion driven by changes in other observables? To explore this, we move from binned regressions to a more structured regression that interacts exchange rate movements with dispersion which we use to show that the positive dispersion-pass-through continues to hold after controlling for a wide variety of time-varying covariates. In particular, we begin by running the following regression:

$$\Delta p_{i,t} = \beta_0 \Delta e_t + \beta_1 IQR_t \times \Delta e_t + \lambda IQR_t + Z'_{i,t}\gamma + +\epsilon_{i,t}, \tag{2}$$

where all variables are defined as in Regression (1). Table 1 column (2) shows the results of this regression with no additional covariates. Consistent with the results in Figure 3, an increase in IQR is associated with a large increase in pass-through. To ease interpretation, coefficients in all tables are standardized, so the 0.07 coefficient on $IQR \times \Delta e$ means that a one standard deviation increase in IQR is associated with an increase in MRPT of seven percentage points. This is a very substantial effect relative to average passthrough of 14.3% given by the coefficient on Δe. For example, it implies that a 10% increase in exchange rates occurring during a month at the 5th percentile of IQR will lead adjusting prices to increase by an average of only 0.3% while the same exchange rate increase occurring during a month at the 95th percentile of IQR will lead prices to rise by 2.6%.

Since our MRPT specification conditions on price adjustment and Gopinath and Itskhoki (2010) show an important relationship between frequency and long-run pass-through, it is natural to ask whether IQR effects are driven by changes in frequency. Column (3) provides evidence that this is not the case. This regression adds controls for $freq_t$ and $freq_t \times \Delta e$ and shows that IQR effects are unchanged.[27] While frequency is perhaps the most obvious potential confounding effect, many other variables also move across time. In column (4), we allow pass-through to vary with a wide array of additional controls. In particular, we introduces interactions of Δe with the frequency of product substitution, the time-series volatility of the exchange rate, seasonality, secular time trends and the business cycle, as measured by GDP growth.[28]

Product substitution can potentially affect measured pass-through as shown in Nakamura and Steinsson (2012), and changes in the volatility of exchange rates might affect both dispersion and pass-through. We allow for secular trends since a prior debate using aggregate data has sometimes found such trends, the presence of which could lead to spurious relationships with dispersion.[29] Finally, since Figure 1 shows that IQR is countercyclical, we control for GDP growth to show that our fact is indeed a dispersion-pass-through relationship not just a business cycle-pass-through relationship. In Appendix A.1, we show results are similar with a variety of other business cycle controls.

Introducing all of these controls mildly reduces the one SD effect of IQR on pass-through from 0.07

[27] Controlling for frequency also partially proxies for changes in the importance of price-spell censoring, which can in turn potentially affect measures of both pass-through and dispersion through selection effects as described in Section 4.2.

[28] The time-series volatility of the exchange rate is measured as the standard-deviation of the bilateral exchange rate associated with a particular item's country of origin in the 12-month period around the month of its price change. Seasonality is captured with 12 month dummies, interacted with exchange rate changes. Secular changes are modeled as a linear trend in pass-through, but similar results obtain when using a quadratic or cubic trend. Real GDP growth is given by chained GDP growth in the quarter corresponding to a given month's price change.

[29] 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. Using our micro data, we find no evidence of trends in MRPT regardless of the treatment of commodities. This difference between our micro results and Marazzi et al. (2005) arises in part because their study uses aggregate data which means that their pass-through statistic measures $frequency \times$ MRPT rather than just MRPT, and frequency did have a declining trend over the period they studied.

to 0.05 but our effect of interest remains economically large and highly statistically significant.[30] Thus, the positive relationship between dispersion and pass-through is robust to including a large set of time-varying covariates. The final columns of Table 1 show that our results also hold when using the standard deviation as an alternative measure of dispersion instead of the IQR. In Appendix A.1, we repeat results separately for imports from individual countries as well as for different product classifications to show that changing composition of the sample along these dimensions does not drive our results.

2.4 Robustness

2.4.1 Does the Great Recession Drive All Results?

Since the increase in IQR during the Great Recession shown in Figure 1 is a large outlier, it is important to show that our results are not driven by this single period. In Table 2, we repeat our regressions excluding the Great Recession.[31] We continue to find very large and significant effects of dispersion on pass-through even outside the Great Recession. While the coefficients are somewhat smaller, it is important to note that units in our tables are standardized so that the regression coefficients represent one-standard deviation effects. Most of the decline in the coefficient on $IQR \times \Delta e$ reflects the fact that IQR has a lower standard deviation outside of the Great Recession rather than a decline in the response of pass-through to a given change in IQR: computing the elasticity of pass-through to an increase in IQR, instead of one standard deviation effects, delivers an elasticity of 2.57 over the entire sample and 2.23 when excluding the Great Recession. These large and similar elasticities are not surprising in light of the scatter plot in Figure 2, which we return to below.

We now show that the time-series relationship between dispersion and pass-through also holds within sectors. This provides additional evidence that our results are not driven solely by the Great Recession or by any other confounding aggregate shock since each sector exhibits different IQR time-series. In column (1) of Table 3, we repeat Regression (2), but replace IQR with IQR_{sector}, which is the interquartile range of all price changes in an item's one-digit sector in a given month. The effect of IQR_{sector} is large and significant. However, it is possible that this is driven by movements in IQR_{sector} that are common across sectors. That is, if IQR_{sector} increases for all sectors then so does IQR, which means the positive coefficient in Column (1) could potentially just be picking up the previously documented IQR effects. Column (2) shows that this is not the case since IQR and IQR_{sector} both independently increase pass-through. However, this behavior could still be driven by the response to a confounding shock which increases pass-through, if IQR and IQR_{sector} also increase at the same time. To eliminate the effect of any common shock which moves both series, in Columns (3) and (4) we include only changes in IQR_{sector} relative to changes in IQR. That is, in these specifications we ask whether sectors that had a relative increase in dispersion have a relative increase in pass-through. Column (3) measures this using the absolute deviation of IQR_{sector} from IQR while Column (4) uses the percentage deviation. Indeed in both cases, relative increases in sectoral dispersion increase relative pass-through. Finally, in Column (5), we redo the regression in Column (1) but with the addition of month-date dummies and

[30] It is unsurprising that introducing a large set of covariates which have previously been shown to have importance for pass-through would absorb some of the initial effects of IQR. Nevertheless, this attenuation is small, and we actually cannot reject equality of coefficients at conventional significance levels.

[31] We exclude all price changes which occur during the Great Recession, but some price changes which occur shortly after the Great Recession end might be changing from a price previously determined during the Great Recession. Repeating results using only completed price spells which are entirely outside of the Great Recession delivers nearly identical results.

month-date dummies interacted with Δe. This is the most stringent test of sector specific effects since in this specification, month dummies absorb the effects of *any* common aggregate shocks which affect pass-through, not just shocks which change aggregate IQR. For example, if the Great Recession increases overall pass-through and IQR through any mechanism, this will be absorbed by these dummies and will not generate a positive coefficient on $IQR_{sector} \times \Delta e$, since that coefficient is only identified off of differences across sectors *within* a given calendar month.

Across all specifications, there is an economically and statistically significant positive relationship between dispersion within sectors and pass-through. Together this greatly alleviates any concerns that our results are spurious or explained by failure to control for confounding shocks.[32]

2.4.2 Misspecification

Our baseline regression 2 is intentionally simple both to illustrate effects transparently and to align our results with much of the existing micro oriented pass-through literature. In particular, we impose a simple linear relationship between pass-through and IQR and assume that only exchange rate movements accumulated over the current spell affect current price changes. The presence of large shocks to exchange rates or IQR could make the first assumption problematic while strategic complementarities or any other force which leads firms to adjust gradually to exchange rate movements could violate the second. More generally, misspecification in our exchange rate regression could lead us to falsely conclude that there is a fundamental relationship between dispersion and pass-through when none exists under a correct specification. We address this concern in two ways. 1) In this section, we show that our empirical results are robust to a wide variety of alternative specifications which are less sensitive to the above concerns. 2) When we turn to quantitative estimation, we use an indirect inference procedure which maps true structural relationships into the same reduced form regression we use in our empirical analysis to show that misspecification within that structural model cannot explain our results.

In Column (1) we show that the positive interaction between IQR and pass-through is not just picking up some non-linearity in the true pass-through specification together with correlation between IQR and Δe by including $(\Delta e)^2$ and $(IQR)^2$ in the regression. That is, we allow for non-linearities as captured by a full second order Taylor expansion for each of our effects. In Column (2) we include non-linear effects of IQR on pass-through but find they are insignificant. Column (3) instead allows for the effects of IQR on pass-through to depend non-linearly on Δe. Unsurprisingly, pass-through rises with the size of the exchange rate shock, and we find similar interactions with dispersion. In Columns (4) and (5) we include cumulative exchange rate movements over prior price spells rather than just over the current spell. Consistent with our theoretical model in 5, increasing IQR increases the response of prices to both current and lagged exchange rate movements, but the effects become weaker at longer lags as items have essentially reached their long-run pass-through after several price changes. In Columns (6) and (7) we split our sample separately into observations with positive and negative exchange rate movements since there might be asymmetry in the response of prices to exchange rate shocks of opposite signs and the sign of exchange rate movements might also be correlated with dispersion. However, we find strong positive relationships in both sub-samples. Column (8) shows results only including items with 5+ price changes

[32] The estimates in Column (2) suggest that roughly 2/3 of the pass-through-dispersion relationship is driven by factors common to all sectors while 1/3 is driven by sector specific factors. Since columns (3)-(6) remove these common effects it is not surprising that they deliver pass-through-dispersion relationships which are somewhat dampened relative to Table 1.

to address concerns that our results might be driven by censoring or biases induced by conditioning on price adjustment.[33] Since these items have many price changes, selection and censoring are much less of a concern but results are nearly identical to our baseline. In Column (9) we instead restrict to items with few price changes and again find a positive relationship.[34]. In Appendix Table A5 we also show that results are similar for pass-through specifications which do not condition on price adjustment and so are unaffected by censoring. As noted in Section 2.2, such specifications are less useful for our identification purposes but can still be helpful for diagnosing misspecification. Finally, column (10) runs a median regression instead of OLS to again address misspecification concerns as well as limit the influence of outliers. This continues to deliver a strong positive relationship between IQR and pass-through.

2.4.3 Cross-Item Evidence

In our final set of robustness results, we show that our results extend from the time-series to the cross-section. In particular, we calculate "item-level" dispersion: the standard deviation of all non-zero price changes for a particular item across time and then show that item-level dispersion is positively correlated with that item's exchange rate pass-through. This robustness check is useful in two ways: 1) When we move to structural models, we will show that variation in responsiveness drives a positive relationship between dispersion and pass-through. Gopinath and Itskhoki (2010) argue that heterogeneity in responsiveness across items is crucial for understanding long-run frequency-pass-through differences in the cross-section. Under this hypothesis, we should also see a positive dispersion-pass-through relationship across items in the data. 2) More importantly, if items differ in their dispersion and pass-through, then we want to ensure that the time-series relationship between dispersion and pass-through is not driven by changes in sample composition across time.

Table 5 shows that there is indeed a positive relationship between the standard deviation of item level price changes and pass-through. Furthermore, this relationship is not driven by differences in the frequency of adjustment across items, and both month-level dispersion (IQR) and item-level dispersion (XSD) have independent positive effects on pass-through. This means that our time-series effects are not explained by composition shifts from low dispersion and pass-through items to high dispersion and pass-through items across time. In Appendix A.1 we show similar results using a less-structured bin regression as in Figure 3 and robustness to a number of the issues raised above.

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.

[33]Results are similar as this threshold is increased to 10 or 15 price changes, but sample sizes decline rapidly.

[34]Unsurprisingly, selecting on items with few changes attenuates all results modestly since overall pass-through is also lower.

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 pass-through. For example, using regression specification (2) we estimate pass-through in each period t as $\widehat{MRPT}_t = \widehat{\beta}_0 + \widehat{\beta}_1 IQR_t$.

Figure 4: Dispersion vs. pass-through

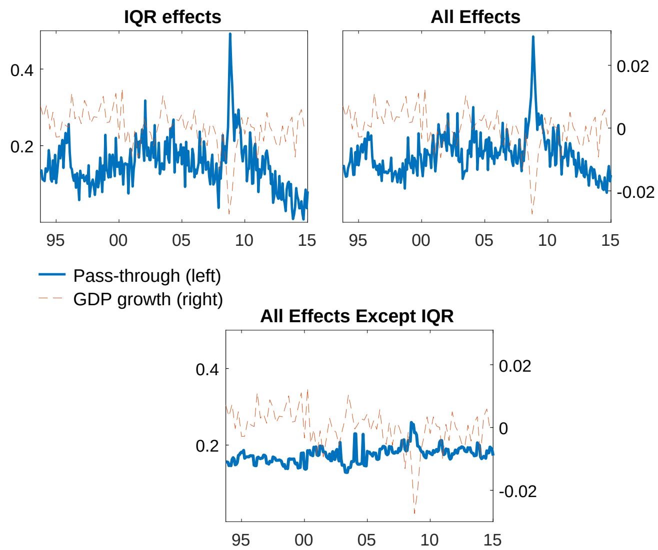

This figure shows separate estimates of pass-through across time plotted against GDP growth, using versions of (2). The upper left panel allows pass-through to vary only with IQR. The upper right panel allows pass-through to vary with IQR, frequency, the volatility of exchange rates and real GDP growth and the bottom panel allows pass-through to vary with each of these variables except IQR.

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 is strongly countercyclical. During the height of the Great Recession, this estimate of exchange rate pass-through rises to almost 50% relative to a low under 5% towards the end of the sample and in the late 90s. 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 volatility of exchange rates, and real GDP growth.[35]

Allowing for these additional interactions does not change the conclusion that pass-through is countercyclical and that time-series variation is largely driven by IQR. This can be seen most clearly in the bottom panel of Figure 4, which shows pass-through estimates for a specification with the additional interactions but excluding IQR. Essentially all the variation in pass-through at business cycle frequencies is captured by time-series variation in IQR.

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 are many possible additional variables we are not controlling for that might affect pass-through and 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)

[35] We exclude the seasonal dummies, time-trend and frequency of product substitution from our prior list of all controls since while point estimates were not zero, none of these coefficients were even marginally significant. Including them as explanatory variables thus introduces additional spurious random noise into the resulting pass-through estimates. Nevertheless, including these effects does not change the substantive conclusion.

with a full set of month dummies. However, small sample sizes make such regressions infeasible. Instead, we estimate the baseline regression using a series of disjoint, rolling windows. 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 from 4 months before to 3 months after t:[36]

$$\Delta p_{i,\tau} = \beta_t \Delta e_\tau + Z'_{i,\tau}\gamma + \epsilon_{i,\tau} \mid t-4 \leq \tau \leq t+3.$$

This allows us to construct a monthly measure of β_t that varies fully non-parametrically across time.[37] Figure 2 from the introduction shows that there remains a strong positive relationship between this fully flexible pass-through specification and IQR so that even being completely agnostic about what drives pass-through movements across time does not change our conclusions. In particular, this scatter plot of pass-through in each 8 month window against the value of IQR in this same window shows a clear positive relationship. It is also worth noting again that we identify the Great Recession observations separately in blue, as additional evidence that our relationship is not driven by this single recession. Reassuringly, best-fit regression lines are nearly identical whether we include or exclude these observations.

Overall these results show that exchange rate pass-through varies substantially with price change dispersion. This means that estimating average pass-through regressions without looking at the distribution of price changes induces a significant time-varying bias, with pass-through substantially understated during periods of high dispersion. A large literature tries to understand average pass-through and its implications for the nominal transmission mechanism, but 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.

We now turn from documenting the fact that price change dispersion is positively correlated with pass-through to showing that it can be used to distinguish changing volatility from changing responsiveness.

4 Basic theoretical framework

4.1 Flexible price model

In this section we lay out a simple framework following Burstein and Gopinath (2013) to show how economic primitives shape the relationship between exchange rate pass-through and price change dispersion. In order to build intuition, we start with the simplest possible setting by assuming flexible prices, no aggregate shocks and no equilibrium effects. This allows us to develop simple formulas relating permanent changes in responsiveness and volatility to pass-through and the dispersion of price changes.[38] 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

[36] Shorter windows allow for more time-variation but induce larger standard errors while larger windows have the reverse trade-off. However, Appendix Figure A.1 shows that results are nearly identical for 4, 6 and 12-month windows.

[37] In order to ensure that each estimate is independent and deliver correct standard errors, we run regressions only using values for t which are 8-months apart so that each regression uses disjoint observations. That is, each price change observation is attributed to a unique window. However, some prices may be changing from prices last set in a different window. Redoing results only including price spells contained entirely in each window produces very similar patterns.

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

item-specific component orthogonal to the exchange rate (η_i):

$$p_i = \mu_i + mc_i\left(e, \eta_i\right). \tag{3}$$

Taking the total derivative of equation (3) gives:

$$\Delta p_i = -\Gamma_i(\Delta p_i - \Delta p) + \alpha_i \Delta e + \epsilon_i \tag{4}$$

where $\Gamma_i \equiv -\frac{\partial \mu_i}{\partial(\Delta p_i - \Delta p)}$ is the elasticity of a firm's optimal markup with respect to its relative price. We define this Γ_i parameter as "responsiveness", for reasons described below. It captures the classic pricing to market channel of Dornbusch (1987) and Krugman (1987), where firms adjust their optimal markups in response to cost shocks, leading to incomplete pass-through. A positive value for Γ_i implies a negative relationship between markups and relative prices, $p_i - p$, which Burstein and Gopinath (2013) show is a robust implication of models 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 α_i as "import intensity". Finally, $\epsilon_i = \Delta \eta_i$ captures the innovation of idiosyncratic marginal cost.[39] We call changes in the variance of this idiosyncratic component changes in "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:[40]

$$\frac{\Delta p_i}{\Delta e} = \frac{\alpha_i}{1 + \Gamma_i} \tag{5}$$

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 $\Gamma_i = 0$ (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 $\Gamma_i > 0$, 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 $\Gamma_i > 0$, the foreign firm will move its price less than one-for-one in response to cost shocks.

Notice that in this flexible price framework, pass-through is determined exclusively by these two factors. Importantly, this means that changing volatility has zero effect on pass-through when prices are fully flexible. This is because pass-through is scale invariant: doubling the size of a cost shock doubles the size of the optimal price change leaving pass-through, which is measured in percentage terms, unchanged.

Since lowering Γ_i means that firms will be more responsive to all cost shocks, we refer to lowering Γ_i as increasing total "responsiveness". That is, firms with low Γ_i will respond strongly to both idiosyncratic and exchange rate shocks. In contrast, firms with high α_i will respond more to exchange rate shocks but not to idiosyncratic cost shocks. Thus, the term responsiveness is used exclusively to refer to Γ_i, which determines general cost pass-through of all shocks, as distinct from parameters such as α_i that affect

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

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

exchange rate specific pass-through.

The open economy literature has extensively studied mechanisms which can generate $\Gamma > 0$ and thus less than full responsiveness to explain incomplete pass-through, but explaining our empirical results require time-variation in this parameter. What can generate such variation? Interestingly, such variation in Γ is precisely what is predicted by the growing set of models in which countercyclical dispersion arises as an endogenous phenomenon. In particular, in Appendix B, we show that a variety of mechanisms such as learning, consumer search, experimentation, ambiguity aversion and market power all naturally map into variation in this responsiveness parameter. This in turn implies that there is an important and heretofore unrecognized link between models explaining pass-through and those trying to understand time-varying dispersion and that these models have very similar reduced form implications.

In addition to its implications for pass-through, we can also use equation (4) to show how α and Γ affect the variance of Δp_i. Solving for Δp_i and computing its time-series variance gives:

$$var(\Delta p_i) = \left(\frac{\alpha_i}{1+\Gamma_i}\right)^2 var(\Delta e) + \left(\frac{1}{1+\Gamma_i}\right)^2 var(\epsilon_i), \tag{6}$$

where we have used the fact that exchange rate and idiosyncratic shocks are uncorrelated.[41]

Intuitively, the variance of the firm's optimal price is larger if it faces more volatile exchange rate or idiosyncratic shocks. The variance of price changes also rises with responsiveness and with import sensitivity ($\alpha_i \uparrow$, $\Gamma_i \downarrow$). Importantly, increases in responsiveness and import sensitivity both also increase pass-through, as shown in (5). However, it can be shown that for empirically relevant values of α_i and Γ_i, changing Γ_i has much larger effects on price change variance than changing α_i.[42] The intuition is that empirical estimates of $var(\Delta p_i)$ greatly exceed $var(\Delta e)$. In addition, estimates of α_i are typically small. (See Figure 3). This means that the first term in (6) 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 effect of greater α_i on exchange rate pass-through and thus $var(\Delta p_i)$ is not quantitatively 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 with important potential effects for MRPT estimates.[43]

[41] If we instead compute the variance of price changes across items at a point in time (month-level dispersion) in this flexible price environment, the first term disappears and so dispersion is wholly determined by Γ_i and $var(\epsilon_i)$

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

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

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

[43] This is a selection bias in the classic statistical sense, in which residuals (ε) are uncorrelated with the explanatory variable (Δe) in the population but are correlated in the sample selected for our regression. We are not first to notice this bias. See the brief discussion in footnotes 7 and 26 of Gopinath et al. (2010). Note that in Calvo pricing models where price

To understand how primitives of a menu cost model affect measured pass-through, it is useful to examine our baseline MRPT specification in equation (1). By definition, the 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" uconditional responsiveness of desired prices to exchange rate movements.[44] 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 larger than true unconditional desired pass-through, β.[45]

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 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, larger shocks mean larger price dispersion. Thus changes in the volatility of idiosyncratic shocks induce a counterfactual negative relationship between MRPT and dispersion and so already suggests that volatility shocks will have difficulty replicating empirical facts in the open economy environment.

Reviewing the conclusions from this and the previous section, it follows that changes in volatility $var(\epsilon_i)$ should generate a counterfactual negative relationship between measured MRPT and price change dispersion while changes in α, Γ or in menu costs should generate a positive relationship. However, we now show that only the responsiveness channel arising from variation in Γ 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. In Appendix B we show that a variety of models which generate time-variation in dispersion endogenously ultimately do so by generating variation in Γ. Since these models all have a similar reduced form interpretation, our qualitative insights do not require taking a stand on a particular source of Γ variation. However, in order to deliver quantitative results we must specify a particular mechanism and functional form for this variation. We do so by building directly on the menu cost model of Gopinath and

adjustment is exogenous, this bias is absent but in this case volatility increases still do not increase pass-through.

[44] This underlying β is determined by α and Γ, 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.

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

Itskhoki (2010) which includes Kimball demand and introduces heterogeneity in responsiveness by including cross-sectional heteroeneity in the "super-elasticity" of demand in preferences.[46] We intentionally build on this workhorse model of incomplete pass-through and adopt this particular form of responsiveness variation for two primary reasons: 1) Gopinath and Itskhoki (2010) show that this model can hit a wide variety of cross-sectional microeconomic facts, and it is important that matching our new facts not come at the expense of missing ones hit by previous models. 2) It is highly parsimonious relative to many other models which give rise to variation in responsiveness. This is crucial in order to estimate our model with a variety of potential aggregate shocks using our previous time-series evidence and then infer the underlying nature of shocks. However, it is important to again emphasize that this does not imply that we think other mechanisms such as experimentation or incomplete information are unimportant, and in reality many of these mechanisms are likely simultaneously present in the data.

5.1 Model Description and Calibration

We begin by describing the model with no aggregate shocks, the baseline calibration and show simple comparative statics to provide a quantitative complement to the results in the previous section. We then formally introduce aggregate shocks to these parameters. In order to infer the importance of various shocks we estimate the model via indirect inference: for a given set of shocks, we solve for the sectoral equilibrium of the model and then simulate data mimicking BLS procedures, run our empirical regressions on this simulated data and compare these results to those in Section 2.3. We then repeat this process repeatedly with alternative sets of aggregate shocks until we find the best fit to the empirical data.

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:

$$\frac{1}{|\Omega|}\int_{\Omega}\Psi\left(\frac{|\Omega|\,C_j}{C}\right)dj = 1 \tag{7}$$

with $\Psi(1) = 1, \Psi'(.) > 0$ and $\Psi''(.) < 0$. C_j is the quantity demanded of variety $j \in \Omega$, where Ω is the set of all varieties available domestically. Ω has measure $1 + \omega$. Individual varieties are aggregated into a final consumption good C. This intermediate aggregator contains the CES specification as a special case. The demand function for C_j implied by equation (7) is:

$$C_j = \varphi\left(D\frac{P_j}{P}\right)\frac{C}{|\Omega|}, \text{ where } \varphi(.) \equiv \Psi'^{-1}(.) \tag{8}$$

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

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

[46] A Calvo model delivers similar conclusions about the importance of responsiveness but fits micro facts less well.

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:[47]

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

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:

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

Combining unit revenues, unit costs and total demand for variety j 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 states and P_t, W_t, and W_t^* are aggregate states. 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.

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 $E_t \ln P_{t+1} = \gamma_0 + \gamma_1 \ln P_t + \gamma_2 e_t$. We then solve the firm's Bellman equation for a given conjecture for γ, 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

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

assuming that the value of the domestic currency and real wage are stable relative to the exchange rate. It is indeed the case in the U.S. that exchange rates have little explanatory power for these variables in the U.S. since net exports are a small part of the overall U.S. economy.

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 7):

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

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:

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

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 changes in these two parameters. For simplicity and following Gopinath and Itskhoki (2010), we assume that variation in Γ is driven solely by ε but note that variation in σ would yield similar results. We return to this point in Appendix B when discussing additional sources of variation in Γ.

Calibrated values for all parameters are reported in Table 6. 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 to yield a steady-state markup of 25%. This is the middle of the range estimated for U.S. imports by Broda and Weinstein (2006). We assume that the log of the real exchange rate, e, follows a random walk. 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.[48] The average of this import share for the U.S. over our sample period is 14.5%, which implies that $\omega = 0.17$.We set the persistence of the idiosyncratic shock process, ρ_A, to be equal to

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

0.85, which is in between the values used by Gopinath and Itskhoki (2010) and Nakamura and Steinsson (2008), and we set $\kappa = 0.05$ to match the frequency of price adjustment of 17% in our sample.[49]

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. In Appendix Figure B.4 we plot the relationship between each parameter and these moments, but to get a sense for why these moments separately identify our parameters, it is useful to remember the intuition from our simple model and our baseline MRPT regression:

$$\Delta p_{i,t} = \beta \Delta e_t + \epsilon_{i,t} \tag{9}$$

Decreasing ε means that firms respond more to both exchange rate movements and idiosyncratic shocks when adjusting prices. This increases pass-through and the standard deviation of price changes but has little effect on the R^2 from estimating equation (9). 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 modest 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 0.165, 2.35 and 0.08, respectively.

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 5 shows results when we fix three of ε, κ, α and σ_A at their baseline calibrated values and vary the fourth parameter. For each value of this parameter we solve the model, simulate a panel of firms with the same number of observations as in the BLS data and compute MRPT, frequency and the standard deviation of price changes exactly as in Section 2.

This comparative statics exercise allows us to trace out how changes in structural parameters affect the joint-behavior of these 3 statistics and provides a quantitative counterpart to the intuition in Section 4.1.[50] This cannot be mapped directly to the empirical results in Section 2.3, since it is showing the

[49] Note that our sample in both the model and data only includes items with at least two price changes, so this frequency is moderately higher than the frequency of price adjust of all items in the IPP.

[50] 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.

Figure 5: Comparative Statics

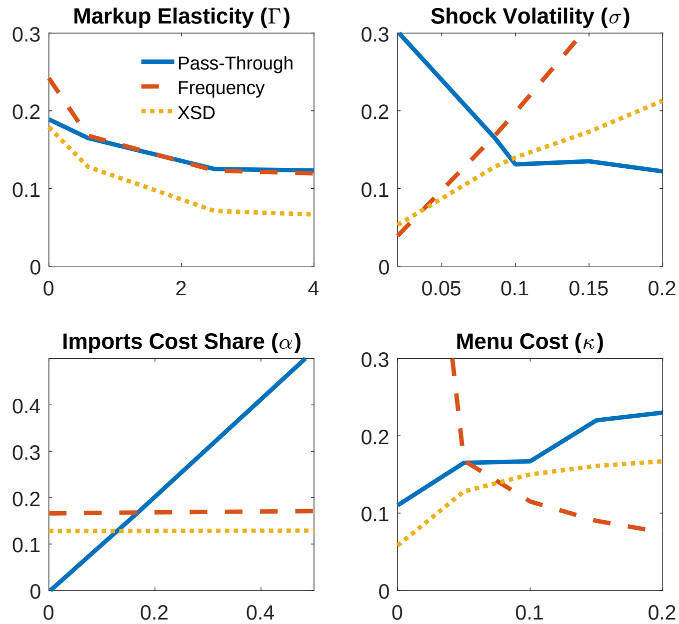

Holding all other parameters at their baseline values, this figures shows the effect of varying individual parameters on frequency, pass-through and the standard deviation of price changes. The solid blue line shows MRPT, the dashed red line shows frequency and the dotted yellow line shows the standard deviation of price changes. The x-axis in each plot shows the value of the parameter being varied.

implications of permanently changing parameters within a model and so does not correspond exactly to our empirical exercise. However, it gives a sense of the quantitative response of observable moments to underlying structural parameters and so is useful for guiding the indirect inference exercise which follows. There we introduce aggregate shocks, simulate time-series and run regressions just as in Section 2.3.

The top-left panel shows the effects of changing responsiveness by varying the markup elasticity Γ from 0 to 4 (corresponding to moving ε from 0 to 16). It is apparent that lowering responsiveness (increasing Γ) causes pass-through, frequency and the standard deviation of price changes to all fall. The upper-right panel shows that increasing the volatility of shocks σ_A also increases the standard deviation of price changes and frequency but instead lowers pass-through. This is because larger σ_A increases price makes firms more likely to adjust their prices for purely idiosyncratic reasons, which reduces selection effects and MRPT. Thus, variation in responsiveness results in a positive pass-through-dispersion relationship while variation in volatility generates a counterfactual negative relationship.

The bottom-left panel shows what happens as we vary α from 0 to 0.5. This leads to large changes in MRPT but negligible movements in the variance of price changes and frequency. This quantitatively confirms the intuition in Footnote 4 that reverse causality, in which increasing α mechanically increases dispersion, is unimportant for our results. This is because idiosyncratic shocks are much more important than exchange rate shocks for explaining price change dispersion so that increasing the sensitivity to the exchange rate barely raises price change dispersion.

The bottom-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 positive relationship between MRPT and dispersion. 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, but it also leads to a large decline in the frequency of price adjustment. Since this strong negative relationship between dispersion and frequency is counterfactual, this is what ultimately leads us to reject variation in menu costs as an explanation for our empirical results.

While we view this comparative statics exercise as very informative, it has some weaknesses: 1) In the data, we are sorting months and firms into bins by the dispersion 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 by this (unobserved) parameter rather than by price change dispersion. Thus, there is not a clean match between our comparative statics simulations and our empirical exercise. 2) In the data, firms and time periods 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 relatively 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 aggregate shocks, which we assume are unobserved by the econometrician. We then formally estimate the importance of different shocks in explaining our empirical results using indirect inference. More specifically, we assume that

$$\begin{aligned}
\ln \varepsilon_t &= \ln \varepsilon^{ss}(1-\rho) + \rho \ln \varepsilon_{t-1} + \epsilon_t \text{ with } \epsilon_t \sim N(0, \sigma_\varepsilon) \\
\ln \sigma_t &= \ln \sigma^{ss}(1-\rho) + \rho \ln \sigma_{t-1} + s_t \text{ with } s_t \sim N(0, \sigma_\sigma) \\
\ln \kappa_t &= \ln \kappa^{ss}(1-\rho) + \rho \ln \kappa_{t-1} + \gamma_t \text{ with } \gamma_t \sim N(0, \sigma_\kappa).
\end{aligned}$$

where $\varepsilon^{ss}, \sigma^{ss}, \kappa^{ss}$ are the steady-state values shown in Table 6. Since each additional aggregate shock increases the computational burden in estimation substantially, and since Figure 5 shows that changes in α do not affect dispersion, we do not model shocks to α.[51] Once we introduce aggregate shocks, we must also modify the equilibrium transition rules, which assume then take the form:

$$\begin{aligned}
E_t \ln P_{t+1} = \gamma_0 &+ \gamma_1 \ln P_t + \gamma_2 e_t \\
&+ \gamma_3 \ln \varepsilon_t + \gamma_4 \ln P_t \ln \varepsilon_t + \gamma_5 e_t \ln \varepsilon_t \\
&+ \gamma_6 \ln \sigma_t + \gamma_7 \ln P_t \ln \sigma_t + \gamma_7 e_t \ln \sigma_t \\
&+ \gamma_8 \ln \kappa_t + \gamma_9 \ln P_t \ln \kappa_t + \gamma_{10} e_t \ln \kappa_t.
\end{aligned}$$

That is, we allow the price level to have an intercept, persistence and sensitivity to exchange rates that depends on the current realization of our three aggregate shocks. For a given set of parameters, we then solve for the model equilibrium and then construct a firm panel, which we sample exactly as in BLS

[51]Previous versions of this paper, which calibrated instead of estimating shocks, also included shocks to the volatility of exchange rates and to the "common-ness" of exchange rate shocks, to reflect the fact that the Great Recession was a large, common aggregate shock. We found they were unable to explain our empirical results. Including these shocks would make estimation computationally infeasible. Furthermore, our empirical results control for the volatility of exchange rates and are not driven by the Great Recession. Similarly, while we could in principle estimate a different ρ for each aggregate shock, this increase in the parameter space would also render estimation currently infeasible, since estimating the 4 parameters in our restricted model requires roughly one month of calendar time on a computing cluster and several years of cpu time.

microdata to account for any small sample issues or other misspecification concerns which might arise in our reduced form 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, and we pick our four parameters $(\rho, \sigma_\varepsilon, \sigma_\sigma, \sigma_\kappa)$ to best 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 simultaneous shocks and formally assess their relative importance.

To construct our empirical moments, we first sort months into five bins by their month-level price change standard deviation. We then calculate the relative standard deviation of price changes, the relative MRPT, and the relative frequency for each standard deviation bin.[52] The first five moments test the model's ability to capture the time-series variation 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 showed can help identify shocks to menu costs from shocks to responsiveness.[53]

Given these 15 moments, we pick our 4 parameters to solve $\widehat{\theta} = \arg\min_\theta g(\theta)' W(\theta) g(\theta)$ with positive definite weight-matrix $W(\theta)$.[54] Table 7 shows resulting parameter estimates and measures of model fit.

The main take-away from Table 7 is that we estimate an important role for σ_ε but no role for σ_σ or σ_κ. In fact, even though we allow for simultaneous aggregate shocks to responsiveness, volatility and menu costs, our estimation ultimately prefers a single shock model with only responsiveness shocks. Inspecting standard errors around these point estimates, the model rejects essentially any role for volatility shocks while it allows for some possibility of modest shocks to menu costs. Conversely, versions of the model without responsiveness are strongly rejected: inspecting the goodness of fit, we can easily reject all models with $\sigma_\varepsilon = 0$ in favor of the unrestricted model that allows for such variation.

These numerical results can be seen more easily in Figure 6, which shows the unrestricted model fit to each moment as well as that of the restricted model with no responsiveness shocks. Clearly, the model with no responsiveness shocks is unable to match the positive correlation between dispersion and pass-through.[55]

While the bulk of the paper has focused on time-series variation in dispersion, Section 2.4.3 documented similar patterns across items. In Appendix B.5, we thus repeat our indirect inference including permanent cross-item heterogeneity rather than aggregate shocks and show that this exercise delivers similar conclusions. In particular, our estimation using cross-item empirical data finds an important role for permanent responsiveness differences across items. Reassuringly, this is also consistent with the

[52] 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 our exercise with shocks as largely about trying to match relative differences across time. Nevertheless, redoing the results using absolute rather than relative moments did not qualitatively change the conclusions but modestly reduces the overall fit of relative movements.

[53] 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.

[54] 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.

[55] More precisely, some parameter configurations with large menu cost shocks can match this relationship, but they do so at the cost of a terrible fit to frequency. Our estimation optimally weights the deviation from each moment and the restricted model prefers hitting frequency and missing the dispersion-pass-through relationship rather than hitting this relationship and missing frequency even more dramatically.

Figure 6: Indirect Inference Estimation Results

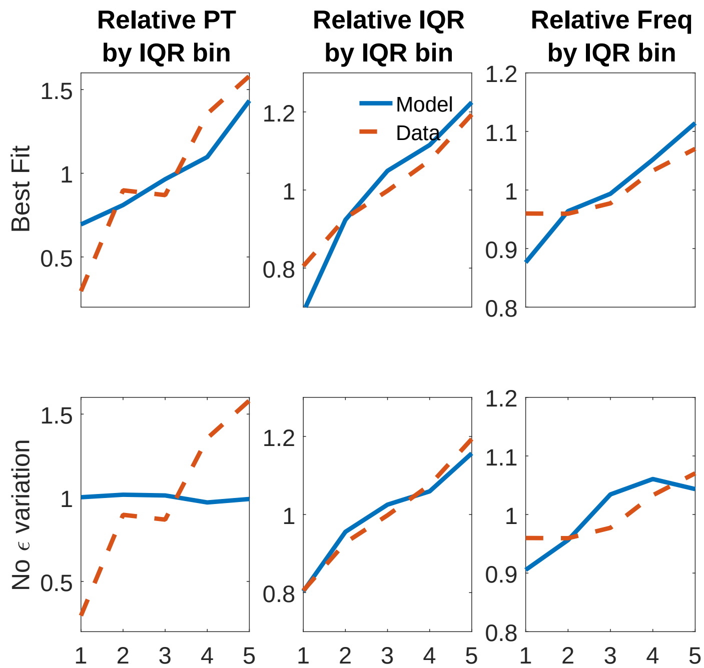

This figure sorts months into 5 bins by IQR. The left panels show how pass-through varies across these 5 bins, the middle panels show how IQR varies across the 5 bins and the right panels show how frequency varies. Model moments are shown in solid blue and data moments are shown in dashed red. The first row shows results for the best fit estimates of $\rho, \sigma_\varepsilon, \sigma_\sigma, \sigma_\kappa$ and the second row shows the best fit for the restricted model with $\sigma_\varepsilon = 0$. Weights on moment deviations are computed using an estimate of the efficient weight matrix.

conclusions in Gopinath and Itskhoki (2010), which arise from matching a different set of facts.

5.4 Interpreting Magnitudes: Implications for Other Observables

Beginning from the steady-state value of $\Gamma = 0.59$, our estimates imply that a one standard deviation decline in responsiveness lowers Γ to 0.43. Interpreting the plausibility of this variation directly is somewhat challenging for two reasons: 1) In reality changes in Γ are likely driven by a variety of mechanisms acting simultaneously so that our estimates of super-elasticity changes are likely standing in as a reduced form for a variety of mechanisms discussed in Appendix B. 2) More importantly, even if one views super-elasticity shocks as the sole source of responsiveness variation in reality, we have no empirical measures of how this elasticity-of-elasticity of demand varies across time to compare our model to. Just measuring the elasticity of demand is difficult much less how it moves with a firm's relative price. We thus instead argue for the plausibility of our estimates by showing that they imply reasonable time-series variation in many economic variables which we *can* measure in the data.

In particular, we show that our model is capable of explaining time-series patterns for the IQR of price changes, overall inflation, import inflation, output growth, frequency and pass-through while also generating plausible markup behavior over our 1993-2015 sample period.

Before describing these results, we must first introduce one additional shock which is necessary for this exercise to be well-defined. In our results thus far we have abstracted from aggregate nominal shocks by assuming total nominal ouptut PC is constant. While this assumption is not important for any of our prior conclusions, it must be relaxed in order to simultaneously match aggregate inflation and output in the data since if nominal output is fixed then inflation and real output growth are perfectly negatively correlated.[56] Following Nakamura and Steinsson (2010), we assume that these shocks are iid normal and calibrate their standard deviation to .005 to match the standard deviation of nominal output growth net of real output growth over our 1993-2015 sample period.

Figure 7: Time-Series Fit of Model

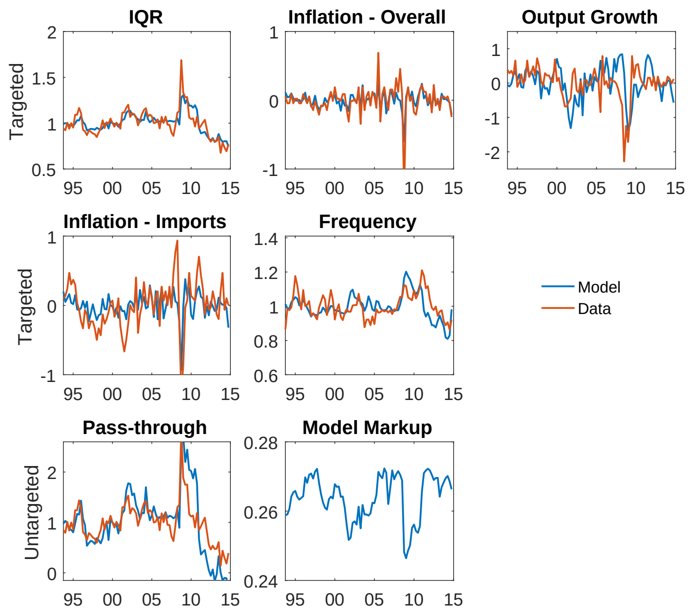

Beginning from the ergodic distribution, this figure shows results when we pick exchange rates in the model to match the major currencies trade-weighted exchange rate from 1993-2015 and pick the value of the nominal and responsiveness shock to fit the five targeted series.

Beginning from the ergodic steady-state of the model, we feed the observed sequence of exchange rates from 1993-2015 into the model and then pick the value of the two remaining shocks (ε and nominal output) period-by-period to best fit the IQR of price changes, overall inflation, import inflation, output growth and frequency.[57] Figure 7 shows that we are able to quite well-match these five series with only these two shocks.[58] Importantly, the model does a good job matching time-series patterns of pass-through shown in Figure 4, despite not directly targeting this series. Finally, we also show the time-series for the average realized markup in the model. There is a vast debate empirical debate on both the size and cyclicality of markup variation, so there is no obvious empirical counterpart to this series. Nevertheless, the overall variation we find is relatively small, meaning that the super-elasticity shocks necessary to match the other time-series do not imply markup variation which seems at odds with the data.

In sum, our parsimonious model provides a good fit to the data along multiple dimensions. We

[56] While it is infeasible to reestimate the model with this additional aggregate shock, we have recomputed moments for our best fit parameters (as well as for several other parameter values) when including this additional shock, and it has a negligible effect. This is because while this shock is important for explaining aggregate inflation and output comovement, it is much, much smaller than exchange rate and idiosyncratic shocks and so has almost no effect on firm pricing behavior.

[57] We use the major currencies trade-weighted exchange rate from the IMF. In order to estimate effects beginning from the ergodic distribution, we simulate an initial 10 year burn-in period and average our results over 20 replications to eliminate any effects of initial conditions.

[58] If we focus only on nominal variables and ignore output, we can match the other four series almost as well with only the single responsiveness shock.

view this as significant evidence that our estimated shocks to responsiveness are both plausible and an important driver of many aggregate variables.

6 Conclusion

An active theoretical literature debates whether time-variation in the dispersion of economic variables is driven by changes in the volatility of exogenous shocks or in the endogenous response to shocks of constant size. In this paper, we provide evidence from import prices that variation in price change dispersion is driven by changing responsiveness rather than volatility. Using confidential item-level micro price data from BLS import price indices, we document a robust positive relationship between price change dispersion and exchange rate pass-through, both across time and across items. We then estimate a structural price-setting model using indirect inference to match these facts and show it strongly supports variation in responsiveness while rejecting variation in volatility. This is because greater idiosyncratic volatility leads to price changes which are more orthogonal to exchange rate movements, reducing measured pass-through. In contrast, increasing endogenous responsiveness leads firms to respond more strongly to both idiosyncratic and exchange rate shocks and so increases dispersion and pass-through.

The result that volatility shocks induce a negative relationship between pass-through and dispersion is quite general since it arises from a simple statistical selection effect and so should hold in any any model in which larger shocks increase the probability of price adjustment.[59] Since volatility shocks induce a negative relationship between dispersion and pass-through, this means they are not just a worse fit to the data on this dimension than models with responsiveness shocks; they are a worse fit than models with no volatility or responsiveness shocks at all. This means that our paper simultaneously provides evidence against models with volatility shocks and evidence in favor of models with responsiveness changes, but these conclusions are to some extent independent. While we make a case that a particular model of responsiveness can match the data well, accepting that volatility shocks are at odds with the data does not actually require one to accept our model of responsiveness.

Conversely, our assertion that the price data favors time-variation in responsiveness does not depend crucially on our simplifying assumption that this variation arises from changes in the super-elasticity of demand. We adopt this parsimonious specification from Gopinath and Itskhoki (2010), but in Appendix B, we show explicitly that a variety of models which have been used to endogenously generate countercyclical dispersion all share a common reduce form representation. In particular, the forces in each of these models give rise to reduced form variation in responsiveness, Γ. While these models were designed in part to endogenously generate time-variation in dispersion, since they do so by changing Γ, they also lead to a positive relationship between dispersion and pass-through. In this sense, one should not interpret our modeling assumption as endorsing or rejecting any particular mechanism which generates endogenous responsiveness. It should instead be interpreted as providing broad support for models which generate time-variation in responsiveness.

Finally, it is important to reiterate the caveat in the introduction that our analysis is focused on import prices. However, there is a limited but growing body of research arriving at similar conclusions in other contexts. We focus on imports since we measure shocks to costs using exchange rate movements.

[59] In a Calvo model where frequency is exogenously fixed or in a flexible price model, price adjustment does not rise with volatility. In those models the correlation between dispersion and pass-through is zero, which is still counterfactual.

However, our methodological insight, that the joint behavior of pass-through and dispersion can be used to differentiate changes in volatility from changes in responsiveness, should apply to shocks other than exchange rates and outcomes other than prices. This means that with different microdata, a similar exercise could be performed studying the pass-through of any well-identified aggregate or sectoral shock into any outcome variable of interest and how that relates to the dispersion of that variable. We think it is an interesting avenue for future research to extend our analysis to variety of shocks such as credit, energy price, or monetary shocks identified using high frequency financial data and to explore pass-through into a variety of endogenous outcomes.

References

Bachmann, R. and C. Bayer (2014). Investment Dispersion and the Business Cycle. *American Economic Review*.

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

Baker, S. and N. Bloom (2013). Does uncertainty reduce growth? using disasters as natural experiments. *NBER Working Paper 19475*.

Baley, I. and J. Blanco (2016). Menu costs, uncertainty cycles, and the propagation of nominal shocks. *Mimeo*.

Berger, D., I. Dew-Becker, and S. Giglio (2016). Uncertainty shocks as second-moment news shocks.

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*.

Burstein, A. and G. Gopinath (2014). International prices and exchange rates.

Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2016). Changing business dynamism and productivity: Shocks vs. responsiveness.

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

Fleer, R., B. Rudolf, and M. Zurlinden (2015). Price change dispersion and time-varying pass-through into consumer prices.

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

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*.

Ilut, C., M. Kehrig, and M. Schneider (2014). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news.

Kaplan, G. and G. Menzio (2016). Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy*.

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).

Ludvigson, S. C., S. Ma, and S. Ng (2016). Uncertainty and business cycles: Exogenous impulse or endogenous response? Working Paper 21803, National Bureau of Economic Research.

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*.

Munro, D. (2016). Consumer behavior and firm volatility.

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 (2010, August). Monetary Non-Neutrality in a Multi-Sector Menu Cost Model. *Quarterly Journal of Economics 154*(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*.

Stroebel, J. and J. Vavra (2016). 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: Relationship Between pass-through and Dispersion

	(1) Overall	(2) IQR	(3) IQR+Freq	(4) IQR+All Ctrls	(5) XSD	(6) XSD+Freq	(7) XSD+All Ctrls
Δe	0.154	0.143	0.143	0.174	0.141	0.141	0.176
	(0.012)	(0.011)	(0.011)	(0.015)	(0.012)	(0.012)	(0.015)
IQR$\times\Delta$e		0.070	0.070	0.050			
		(0.009)	(0.009)	(0.009)			
IQR		-0.002	-0.002	-0.002			
		(0.001)	(0.001)	(0.001)			
XSD$\times\Delta$e					0.058	0.058	0.038
					(0.009)	(0.009)	(0.009)
XSD					-0.002	-0.002	-0.002
					(0.001)	(0.001)	(0.001)
Freq$\times\Delta$e			0.010	0.019		0.012	
			(0.009)	(0.009)		(0.009)	
Freq			0.004	0.004		0.004	
			(0.001)	(0.001)		(0.001)	
All Ctrls	No	No	No	Yes	No	No	Yes
Num obs	129260	129260	129260	129260	129260	129260	129260
R^2	0.035	0.038	0.039	0.040	0.037	0.038	0.039

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs$\times\Delta$e, gdp growth, gdp growth$\times\Delta$e, SDe, SDe$\times\Delta$e, month dummies, month dummies$\times\Delta$e, t, t$\times\Delta$e, Δ cpi, Δus gdp, Δuscpi. All regressions have country$\times$PSL fixed effects and standard errors are clustered by country$\times$PSL. Primary Strata Lower (PSL) 2 to 4-digit harmonized codes defined by BLS. Dispersion and freq results standardized so that coefficients give a one-standard deviation effect. Sample period: Oct 1993-Jan 2015.

Table 2: Relationship Between pass-through and Dispersion, Excluding Great Recession

	(1) Overall	(2) IQR	(3) IQR+Freq	(4) IQR+All Ctrls	(5) XSD	(6) XSD+Freq	(7) XSD+All Ctrls
Δe	0.129	0.131	0.132	0.159	0.128	0.129	0.160
	(0.012)	(0.012)	(0.012)	(0.014)	(0.012)	(0.012)	(0.014)
IQR$\times\Delta$e		0.037	0.038	0.027			
		(0.011)	(0.011)	(0.012)			
IQR		0.003	0.004	0.004			
		(0.001)	(0.001)	(0.001)			
XSD$\times\Delta$e					0.015	0.016	0.009
					(0.009)	(0.009)	(0.009)
XSD					0.003	0.003	0.003
					(0.001)	(0.001)	(0.001)
Freq$\times\Delta$e			0.014	0.024		0.013	0.024
			(0.009)	(0.010)		(0.009)	(0.010)
Freq			0.004	0.004		0.004	0.004
			(0.001)	(0.001)		(0.001)	(0.001)
All Ctrls	No	No	No	Yes	No	No	Yes
Num obs	119816	119816	119816	119816	119816	119816	119816
R^2	0.034	0.035	0.036	0.037	0.035	0.036	0.037

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs$\times\Delta$e, gdp growth, gdp growth$\times\Delta$e, SDe, SDe$\times$ Δe, month dummies, month dummies$\times$ Δe, t, t$\times\Delta$e, Δ cpi, Δ us gdp, Δ uscpi. Regressions have country$\times$PSL fixed effects and standard errors are clustered by country$\times$PSL. Dispersion and frequency results are standardized so that coefficients give a one-standard deviation effect. Sample period: Oct 1993-Jan 2015, excluding price changes from Dec 2007-Jun 2009.

Table 3: Sectoral vs. Aggregate Dispersion Effects

	(1) IQR_{sector}	(2) IQR_{sector}+ $IQR_{overall}$	(3) IQR_{abs_dev}	(4) $IQR_{\%\_dev}$	(5) IQR_{sector}+ Month dummy $\times$ Δe
Δe	0.184	0.174	0.193	0.187	-0.050
	(0.015)	(0.015)	(0.014)	(0.015)	(0.104)
IQR_{sector} $\times$ Δe	0.038	0.025			0.024
	(0.011)	(0.011)			(0.011)
IQR_{sector}	-0.003	-0.002			-0.001
	(0.001)	(0.001)			(0.001)
$IQR_{overall}$ $\times$ Δe		0.040			
		(0.009)			
$IQR_{overall}$		-0.001			
		(0.001)			
IQR_{abs_dev} $\times$ Δe			0.026		
			(0.011)		
IQR_{abs_dev}			-0.001		
			(0.001)		
$IQR_{\%\_dev}$ $\times$ Δe				0.030	
				(0.012)	
$IQR_{\%\_dev}$				-0.001	
				(0.001)	
Month Dummy	No	No	No	No	Yes
Month Dummy$\times\Delta$e	No	No	No	No	Yes
All Ctrls	Yes	Yes	Yes	Yes	Yes
Num obs	129232	129232	129232	129232	129232
R^2	0.040	0.040	0.039	0.039	0.067

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs$\times$ Δe, gdp growth, gdp growth $\times$ Δe, SDe, SDe $\times$ Δe, month dummies, month dummies $\times$ Δe, t, t x Δe, Δ cpi, Δ us gdp, Δ uscpi. $IQRsector, t$ is the month-level interquartile range of an items 1-digit sector in month t. $IQRoverall, t$ is month-level interquartile range across all items in month t. $IQRdeviation, t = IQRsector, t - mean(IQRsector, t)$ is the absolute deviation of the IQR in an item's sector from the average IQR across all sectors in month t. $IQRrelative, t = IQRsector, t/mean(IQRsector, t)$ is the percentage deviation of the IQR in an item's sector from the average IQR across all sectors in month t. See text for additional description. All regressions have country x PSL fixed effects and robust standard errors are clustered at the country x PSL level. Dispersion and frequency results are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table 4: Robustness to Misspecification

	(1) Full Interaction	(2) IQR+ IQR^2	(3) Δe+ Δe^2	(4) Δe+ l. Δe	(5) Δe+ l. Δe+ l2. Δe	(6) Δe>0	(7) Δe<0	(8) 5+chgs	(9) <=5 chgs	(10) Median Regs
Δe	0.158	0.164	0.188	0.239	0.265	0.168	0.178	0.199	0.098	0.163
	(0.014)	(0.015)	(0.015)	(0.035)	(0.039)	(0.020)	(0.018)	(0.017)	(0.014)	(0.010)
IQR $\times\Delta$e	0.026	0.032	0.066	0.076	0.073	0.054	0.037	0.060	0.024	0.029
	(0.009)	(0.011)	(0.012)	(0.022)	(0.026)	(0.014)	(0.014)	(0.011)	(0.013)	(0.006)
IQR	0.001	0.001	-0.002	-0.002	-0.002	-0.004	-0.000	-0.003	0.003	0.003
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$IQR^2\times\Delta$e		-0.004								
		(0.005)								
IQR^2	-0.004	-0.005								
	(0.000)	(0.000)								
IQR$\times\Delta e^2$			0.057							
			(0.022)							
Δe^2	-0.012		0.059							
	(0.022)		(0.023)							
IQR$\times$l.Δe				0.070	0.074					
				(0.019)	(0.030)					
l.Δe				0.065	0.086					
				(0.014)	(0.027)					
IQR$\times$l2.Δe					0.032					
					(0.025)					
l2.Δe					0.063					
					(0.018)					
All Ctrls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Num obs	129260	129260	129260	57589	42127	62551	62297	114141	28014	129260
R^2	0.044	0.044	0.041	0.025	0.028	0.049	0.057	0.029	0.095	

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs$\times$ Δe, gdp growth, gdp growth $\times$ Δe, SDe, SDe $\times$ Δe, month dummies, month dummies $\times$ Δe, t, t x Δe, Δ cpi, Δ us gdp, Δ uscpi. l.Δe and l2.Δe are the cumulative exchange rate movement in the lagged and twice lagged price spell, respectively. See text for additional description. All regressions have country $\times$ PSL fixed effects and robust standard errors are clustered at the country $\times$ PSL level. Dispersion and frequency results are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table 5: Cross-Item Results

	(1) XSD_{item}	(2) XSD_{item}+ $Freq_{item}$	(3) XSD_{item}+ $Freq_{item}$+IQR	(4) XSD_{item}+$Freq_{item}$ +IQR+ all controls
Δe	0.151	0.162	0.152	0.197
	(0.012)	(0.014)	(0.012)	(0.016)
$XSD_{item}\times\Delta$e	0.033	0.030	0.026	0.028
	(0.013)	(0.013)	(0.012)	(0.011)
XSD_{item}	0.001	0.001	0.002	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
$Freq_{item}\times\Delta$e		0.024	0.025	0.041
		(0.011)	(0.010)	(0.009)
$Freq_{item}$		-0.001	-0.002	0.004
		(0.001)	(0.001)	(0.001)
IQR$\times\Delta$e			0.069	0.047
			(0.009)	(0.009)
IQR			-0.002	-0.002
			(0.001)	(0.001)
All Ctrls	No	No	No	Yes
Num obs	129260	129260	129260	129260
R^2	0.036	0.036	0.039	0.041

"All controls" are frequency of adjustment (freq), freq$\times$ Δe,frequency of product substitutions (subs), freq and subs $\times$ Δe, gdp growth, gdp growth$\times\Delta$e, SDe, SDe$\times\Delta$e, month dummies, month dummies $\times$ Δe, t, t$\times\Delta$e, Δ cpi, Δ us gdp, Δ uscpi. See text for additional description. Regressions have country$\times$PSL fixed effects and robust standard errors clustered at the country$\times$PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table 6: Parameter Values

Parameter	Symbol	Menu Cost Model	Source
Discount factor	β	$0.96^{1/12}$	Annualized interest rate of 4%
Fraction of imports	$\omega/(1+\omega)$	14.5%	BEA input-output table
Cost sensitivity to ER shock			
Foreign firms	α	0.165	Estimation (see text)
U.S. firms	α^{US}	0	
Menu cost	κ	5.0%	Estimation (see text)
Markup elasticity	ε	2.35	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	8.6%	Estimation (see text)
Persistence of shock	ρ_A	0.85	Gopinath and Itshkoki (2010)

Table 7: Estimated Parameters and Fit

Parameter	Estimate	95% Confidence Interval	
σ_ε	0.365	(0.347,0.383)	
σ_σ	0.000	(0.00,.0053)	
σ_κ	0.014	(0.00,.0337)	
ρ	0.0845	(0.838,0.852)	

Models	Wald-Statistic/Likelihood Ratio	95% Critical Value	99% Critical Value
Unrestricted Model	41.64	19.68	24.72
$\sigma_\varepsilon = 0$	113.2851	3.84	6.64

Asymptotic s.e.'s for parameters in parantheses. Unrestricted model Wald-Statistic: $g\left(\widehat{\theta}\right)' W\left(\widehat{\theta}\right)' g\left(\widehat{\theta}\right) \sim \chi^2(11)$

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

.

Online Appendix Materials - Not For Print

A Empirical Appendix: Online Only

A.1 Additional Empirical Results

In this section, we provide a number of robustness checks and extensions of our primary analysis.

Figure A.1 replicates Figure 2 using alternative window lengths and shows that we continue to find a strong positive relationship between our non-parametric pass-through estimates and dispersion which arises both including and excluding the Great Recession.

Figure A.1: Dispersion vs. Pass-through: Different Windows

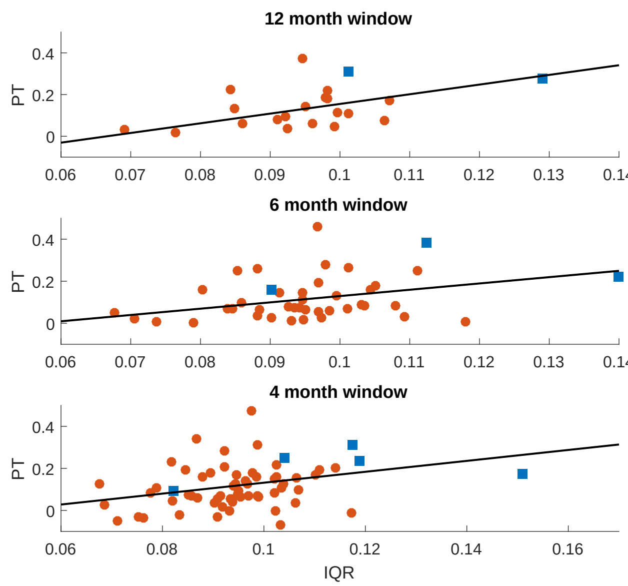

This figure shows the IQR of all non-zero price changes against our preferred measure of exchange rate pass-through, described below. Both statistics are computed separately in a series of disjoint windows which span our sample period. Our primary specification in the text uses 8 month windows, but this figure shows results are similar for 4, 6 and 12 month windows. Windows which have a majority of months during the Great Recession, as defined by NBER, are shown in blue. The black regression line includes all observations while the red-dotted line excludes Great Recession observations.

Figure A.2 repeats the binned time-series regression in Figure 3 using a much larger number of bins. This allows for a more non-parametric relationship between pass-through and dispersion and again shows that the linearity assumed in most of our empirical regressions is a reasonable approximation of the data. Unsurprisingly, there is somewhat more noise when performing this exercise, but the basic picture is unchanged.

Figure A.2: Non-Parameteric IQR-pass-through Relationship

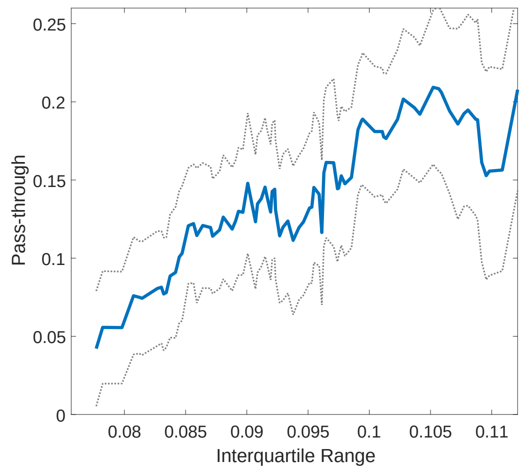

This figure shows separate estimates of regression (1) in each of 80-intervals by months' IQR. The first point includes observations from months with IQR in percentiles 1-20, the second observation months in percentiles 2-21, up to the last observation which includes observations from months in percentiles 80-100. All regressions have country $\times$ PSL fixed effects and robust standard errors are clustered at the country $\times$ PSL level. We also include controls for foreign CPI growth, US gdp growth and US CPI growth. 95% confidence intervals are shown with dotted lines, and the average IQR in each window is shown on the x-axis.

Figure A.3 repeats the binned time-series regression in Figure 3 instead using cross-item dispersion. In particular, we sort individual items by their item-level standard deviation into 5 quintiles and then run regression 1 separately in each bin. This shows that there is a positive relationship between item-level dispersion and pass-through using a specification that does not impose linearity like in Table 5.

Figure A.3: Item-Level Dispersion-pass-through Relationship

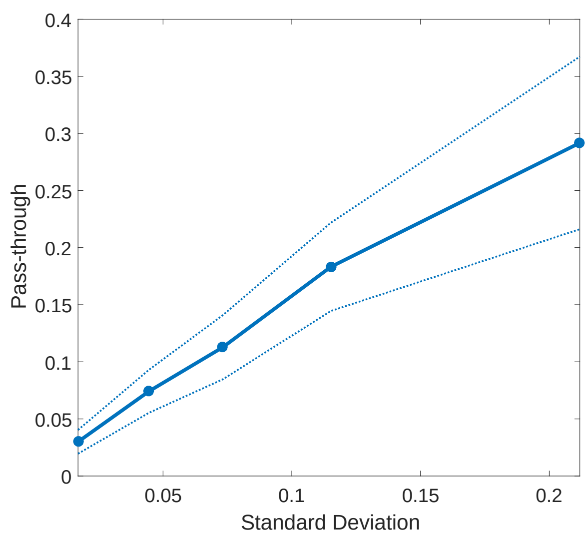

This figure shows separate estimates of regression (1) in each of 5-quintiles by the the item-level standard deviation of price changes. All regressions have country $\times$ PSL fixed effects and robust standard errors are clustered at the country $\times$ PSL level. We also include controls for foreign CPI growth, US gdp growth and US CPI growth. 95% confidence intervals are shown with dotted lines, and the average item-level standard deviation value in each quintile is shown on the x-axis.

In Table 1, Columns (4) and (7), we showed that despite the fact that dispersion is countercyclical, our patterns indeed reflect a pass-through-dispersion relationship and are not just proxying for a pass-through-business cycle relationship. In that table, we measured the business cycle using real GDP growth, but one might be concerned that real GDP growth is only a partial proxy for the business cycle. Table

A1 shows that our conclusions are robust to instead measuring the business cycle using NBER Recession indicators or using HP filtered log GDP instead of gdp growth. These results show that pass-through is indeed countercyclical (at least when measuring cyclicality using real GDP growth or business cycle dates), but that this does not drive our dispersion effects. The effects of dispersion on pass-through are very similar after controlling for business cycle effects.

One might also be concerned that our results could be driven by compositional effects as the mix of product-origin countries and bilateral exchange rates varies across time. Table A2 shows that this is not the case by redoing our results restricted to particular countries/country groups.[60] These compositional concerns are more of a concern for our cross-item effects than our time-series results since an item's country of origin is necessarily fixed across time. Thus, we also repeat our cross-item results for individual countries in Table A3.

In order to use a comprehensive sample, our baseline results include a broad set of items, described in Section 2.1. However, many of these products have less product differentiation or pricing power and so are likely less well described by our theoretical price-setting model. In Table A4 we also show that our results continue to go through when using a narrower set of manufactured products that map more naturally to our model.

Finally, as an additional check of misspecification as well as the importance of our sample selection, Table A5 shows our results using an alternative pass-through specification which does not specifically condition on adjustment. More specifically, we simply regress Δp on Δe plus various additional controls and interactions over various time-intervals, without conditioning on adjustment. For example, in column 1 we simply regress the one month change in price on the one month change in exchange rates, and items in this interval may have either zero or 1 price change. In column 4 we regress one-year changes on one-year changes and item in the regressions may thus have between 0 and 11 price changes in this interval. This specification is more akin to the long-run pass-through measures in Gopinath and Itskhoki (2010). It is less useful for identification purposes but is useful for checking the robustness of our sample selection and for diagnosing misspecification. This is because it can be computed for items with a single price change, in contrast to our primary pass-through measure which can only be computed for items with at least two price changes.[61] Thus, sample sizes are expanded in this specification and we can include items with fewer price changes.

A.2 Additional Sample Summary Statistics

This section provides additional detail on the construction of our benchmark empirical sample and various related summary statistics. From our raw data which includes 2,527,619 price observations from October, 1993 January, 2015, we begin by dropping the 203,562 price observations which are imputed and so flagged as "unusable" observations by the BLS. Row 1 of Table A6 shows the total number of price observations and items as well as various summary statistics of the raw data after dropping these unusable prices. The typical product is in the data set for a little over 3 years and changes prices roughly

[60] There are not enough imports from individual countries aside from Mexico and Canada to get precise individual country estimates.

[61] We require at least one price change so that we can correctly measure Δe. For items with no price changes, exchange rate movements are left-censored and cannot be accurately measured. Nevertheless, despite the concerns with this measure, repeating results simply using the cumulated exchange rate change since an item enters the BLS sample allows us to further expand our sample to include all items and delivers similar results.

9 times. The last 3 columns show the 25th, median and 75th percentile of non-zero price changes. From this raw data, we then exclude commodities, intrafirm transactions and non-dollar prices in our baseline sample. We exclude non-dollar prices because these items mechanically have pass-through of 1 when not adjusting prices and so cannot be used to measure responsiveness. This means, they contain no useful information for our identification purposes. Similarly, commodities exhibit extremely high competition and are undifferentiated. This means they also exhibit nearly 100% pass-through at all times and so cannot be used to measure variation in pass-through across time. Finally, we exclude intrafirm transactions and keep only arms-length price transactions. This is because intrafirm transactions are not necessarily allocative since these transfer prices are often set for tax purposes or other internal purposes and do not necessarily have any relationship to market values so they have little value for our analysis. In total, excluding these prices, which are not informative for our analysis, reduces our sample size substantially. However, of the 1,135,439 observations dropped, the vast majority, 923,978, are intrafirm transactions. This means that although our sample size drops substantially, this is largely just from dropping prices which are essentially mismeasured relative to their allocative value. Overall, this sample selection criteria is identical to the initial sample restriction in Gopinath and Itskhoki (2010), and so makes our results more comparable to the existing literature.

Furthermore, it is important to note that our goal is not to inform aggregate statistics with our analysis. So it is not important that our sample be representative of the overall composition of import price indices. Our goal is to instead use a subset of our data to provide sharp identification, and for these purposes it makes sense to focus on the subset of data most suited for this purpose, even if it does not necessarily aggregate up to national statistics as closely as broader data sets might.

The more relevant comparison is between this initial sample cleaning and our final analysis sample, which includes only observations with at least two price changes. Comparing row 2 and 3 shows that products in our analysis sample have slightly longer average lives in the data set. This is not surprising since items which are only in the data set briefly are less likely to have measured price changes. Even less surprising, the average number of price changes per item is higher in our analysis sample, but this will mechanically be the case since this is how we are selecting our sample. However, the distribution of price changes conditional on adjustment is essentially the same. Overall these comparisons reassure us that we are not performing our analysis on a particularly unusual subset of data. Again, it is worth noting the relationship between our final sample and that in Gopinath and Itskhoki (2010). Our sample is identical to theirs except that we require items to have 2+ price changes while they require items to have only 1+ price changes because MRPT can only be measured for items with two completed price spells while their LRPT measure can be computed for items with only a single price change. However, Appendix Table A5 shows results for alternative specifications that allow us to include items with 1+ price changes. These specifications are less useful for identification purposes but are useful for checking the robustness of our sample selection and for diagnosing misspecification, and overall we find similar patterns.

B Modeling Appendix: Online Only

B.1 Interpretation of Responsiveness Fluctuations

We refer to responsiveness, Γ, as anything that affects the elasticity of a firm's desired price to a cost shock. What economic forces generate time-series variation in this responsiveness parameter? In this section, we show that many of the proposed mechanisms put forward independently to explain countercyclical dispersion, such as ambiguity aversion, customer search, employer learning and experimentation, also map into this responsiveness parameter in a way that has not previously been noted. Conversely, we show that the other dominant mechanism (other than Kimball demand) used by the international finance literature to generate incomplete pass-through – variation in market power, implies a positive relationship between pass-through and dispersion. As a result, all of these mechanisms have similar implications for the relationship between pass-through and dispersion that is at the heart of our paper.

B.1.1 Mechanisms Which Have Been Used to Explain Time-Varying Dispersion

Ambiguity Aversion

Ilut et al. (2014) show that concave hiring rules (which they microfound using an information processing framework where agents are ambiguity averse but which could result from asymmetric adjustment costs) endogenously generate higher cross-sectional (employment) dispersion and shock pass-through during recessions.

It is easy to illustrate the basic mechanism and to see how it naturally maps into responsiveness. Assume firms receive a signal s about future productivity and that the signal has an aggregate and idiosyncratic component, $s = a + \epsilon$, where the idiosyncratic shock ϵ is mean zero and i.i.d. across firms and across time. Further assume, as Ilut et al. (2014) do, that all firms follow the same decision rule, $n = f(s)$, where n is net employment growth and $f(s)$ is strictly increasing and concave. This implies that firms exhibit asymmetric adjustment to shocks: firms respond more to a signal of a given magnitude during recessions than during booms because during recessions firms are in the more concave region of their policy function. As parts 2 and 3 of Proposition 1 in Ilut et al. (2014) prove, this implies that the dispersion of employment changes is higher in recessions.

Concave policy rules also imply that aggregate employment growth is more responsive to aggregate shocks (e.g. higher cost pass-through) in recessions than booms. Formally, for any two realizations of the aggregate shocks with $a' > a$,

$$\frac{d}{da}E[n|a] > \frac{d}{da'}E[n|a'],$$

which follows directly from the strict concavity of $f(s)$. (The formal proof is given in part 1 of Proposition 1 in Ilut et al. (2014)). Thus, a positive correlation between higher dispersion and higher pass-through is a direct implication of concave policy rules. Moreover, there is a mapping between this mechanism and our responsiveness measure. During booms, firms are in the flat region of their concave policy function where they have a low responsiveness to shocks of a given size. However, in recessions firms are in the steep part of their policy functions and endogenously respond more to shock of the same size.[62] Thus, any mechanism that generates concave policy rules as a function of the firms shocks is going

[62] Ilut et al. (2014) show empirically that for the U.S. manufacturing data that both employment dispersion and pass-through are higher in recessions. For pass-through, they estimate hiring rules both non-parametrically and parametrically and find higher pass-through to shocks of the same size in recessions for both rules. In particular, for the non-parametrically estimated hiring rule (see their Figure 6), the average response in boom to a 2 SD shock was +0.16% while the average response in recession to a 2 SD shock was -0.55%. These standard deviation values are calculated over the entire sample, so this shows that the response to a shock of the same size is larger in recessions. For the parametric hiring rule (see Column (I) in their Table 8), the average response in boom was +0.31% while the average response in recession was -1.05%.

to generate countercyclical dispersion and a positive correlation between pass-through and dispersion.

Learning

Baley and Blanco (2016) present a price-setting model with menu costs and imperfect information about idiosyncratic and aggregate productivity. They use this model analyze how price setting behavior is shaped by changes in information by analyzing the response to random increases in "uncertainty", in which firms become less informed about their underlying costs (but with no actual change in current idiosyncratic or aggregate productivity and with no changes in their volatility). That is, they study the response to a pure shock to information in which firms become less informed about their current level of productivity.

The basic logic of their model is simple to understand. Upon the arrival of a new uncertainty regime, a firm's uncertainty increases and then quickly decreases as the firm learns about the shocks they are facing. These informational shocks in turn lead to an increase in price dispersion, as proved in Baley and Blanco (2016) Proposition 6.

Is cost pass-through higher after information shocks? In order to gain intuition, it useful to examine how the level of firm's uncertainty about costs, $\varOmega_t$, affects firms incentive to learn about its markup, μ_t around a short interval of time Δ:

$$\Delta\mu_{t+\Delta} = \left(\frac{\gamma}{\varOmega_t + \gamma}\right)\mu_t + \left(\frac{\varOmega_t}{\varOmega_t + \gamma}\right)(s_{t-}s_{t-\Delta})$$

Firms update their guess of the new markup (which affects the optimal price it would like to set) as a convex combination of a weight on its previous markup and and a weight on the new information from its signals, s_t. Here γ captures the size of the information friction. It is obvious that when information is low and firms are more uncertain about their costs, they (optimally) put more weight on new information. This increases the speed of learning about the new monetary shocks hitting the economy and increases the level of pass-through. Baley and Blanco (2016) show that this implies that pass-through is higher for monetary shocks.

The intuition is simple. The response to monetary shocks is increasing in firms' information about the size of shocks. Since we already established that a decline in information quickens the speed of learning, in the sense that agents put relatively more weight on new signals, this means that firms put more (Bayesian) weight on the new, monetary policy shock and pass-through rises. Thus, their model implies a positive relationship between price change dispersion and pass-through of cost shocks. Baley and Blanco (2016) in fact devote an entire section of their paper to showing that this mechanism is economically important in their calibrated model (see Table 4 in Section 6) and induces variation in dispersion and pass-through that lines up with the empirical facts we document in Section 2.3.

The firm learning mechanism maps precisely into our responsiveness measure: variation in responsiveness corresponds to (endogenous) variation in the speed of firm learning in response to information shocks. When firms have less information, they respond by learning more quickly about the aggregate shocks they face, increasing the responsiveness of their prices to aggregate shocks and increasing price dispersion as they respond more aggressively to idiosyncratic shocks of constant size.

Consumer Search

A growing body of research highlights the importance of changing consumer shopping behavior for business cycle outcomes. For example, Kaplan and Menzio (2016) generate business cycle fluctuations from changes in “market competitiveness”. Unemployed workers spend less and search more for low prices than employed workers, so increases in unemployment increase competition. This increased competition increases incentives for firms to further reduce employment. This feedback between employment and competition can lead to self-fulfilling fluctuations and so endogenously give rise to recessions.

This mechanism is supported by a growing empirical literature. Aguiar, Hurst, and Karabarbounis (2013) document that households search more during recessions. Stroebel and Vavra (2016) show that firms adjust markups in response to changing customer price sensitivity over house price booms and busts. Munro (2016) uses UPC level panel data to show that, consistent with a changing demand elasticity story, dispersion of stores' growth rates increases during recessions and this increase is larger in markets where the increase in consumer shopping effort is highest.

Time-variation in the elasticity of demand naturally maps into our responsiveness framework. Recall, that the steady state level of responsiveness in our model is given by $\Gamma = \frac{\varepsilon}{\sigma-1}$. Thus as long as there is any adjustment of markups in response to shocks ($\varepsilon > 0$), then if certain periods of time such as recessions are characterized by increased competition (because consumers search more), with larger σ and lower markups, they will also be times of greater responsiveness and thus price change dispersion and cost pass-through.[63]

Indeed Munro (2016) explicitly explores the link between changes in the elasticity of demand (coming from variations in consumer search behavior over the cycle) and countercyclical dispersion. The logic is simple. If consumers spend more time shopping for lower prices during recessions in order to smooth consumption, then firms face more elastic demand during recessions. This means that firm sales are more responsive to a given size cost shock leading to higher dispersion of firm sales and employment in recessions. Munro (2016) formalizes this mechanism in a simple business cycle model where search frictions in product markets provide a role for consumer search effort to affect the elasticity of demand that firms face and shows that it generates quantitatively important fluctuations in dispersion even with no changes in the volatility of shocks.

Experimentation

Bachmann and Moscarini (2012) was one of the first papers to explore whether the increase in both macro and micro dispersion was a result of larger shocks or whether causation ran in the opposite direction. In particular, they explore whether time-varying price experimentation in response to negative aggregate shocks can explain countercyclical price dispersion dispersion in both the time-series and the cross-section of individual outcomes.

Bachmann and Moscarini (2012) start by adding imperfect information about demand to an otherwise standard monopolistically competitive model. The basic idea is that firms are heterogeneous in their elasticity of demand but face idiosyncratic demand shocks and so only gradually learn from sales about this elasticity. During booms, price dispersion is low as firms understand the demand curve they face

[63] The presence of markup adjustment can be induced by a wide-variety of strategic-complementarities and is a pervasive assumption. In the pass-through literature this assumption is used explain incomplete pass-through and in the monetary literature it is used to explain large and persistence responses to monetary shocks.

and the cost of deviating from the average price is large in terms of lost profits. However, in recessions, when the chance of bankruptcy is high, they show that the chance that firms will choose to experiment increases because the opportunity cost of price mistakes is lower and the chance of going out of business is higher. Thus, the model delivers countercyclical price dispersion without time-varying volatility shocks.

Their model also implies that pass-through is higher when experimentation is higher. To see this, consider a recession induced by a negative TFP shock. For the firm, this decrease in TFP is a negative cost shock that increases the probability of firm exit and incentive to experiment. Bachmann and Moscarini (2012) show that in this situation the firm will choose to experiment by raising its price, and the size of the price increase is decreasing in firms expectations of future demand. The logic is simple. If the firm does not change its price, it is more likely to go out of business soon, because it likely can no longer cover its costs (this is all probabilistic, based on its beliefs about demand). In principle, it could reduce the price, hoping that true demand is so elastic that revenues will boom, however, if such a high elasticity was plausible then it would have already lowered its price during the boom when the firm was confident demand was high and it could earn large profits.[64] So the only possible move is to raise the price. This generates twin benefits as it increases the chance of survival and also provides information about the demand curve. While firms can experiment at any time, it is not profitable to do so during booms when costs are low and revenues are high and becomes profitable when costs rise in recessions. Hit by these negative cost shocks, firms then choose to experiment in the direction that at least offsets costs. In addition, more pessimistic firms raise their prices by a larger amount than firms with strong beliefs about demand (see Figure 3 in Bachmann and Moscarini (2012)). This means that pessimistic firms have higher pass-through on average than optimistic firms.

Finally, recessions lead to an increase in the mass of pessimistic firms near exit. Since pessimistic firms experiment more and have higher pass-through, this implies that both pass-through and price dispersion rise. Thus variation in the incentive to experimentation acts just like time-varying responsiveness in our baseline framework: both mechanisms generate higher price dispersion and higher pass-through during recessions.

B.1.2 Mechanisms Which Have Been Used to Explain Incomplete Pass-through

In a recent survey of the pass-through literature (Burstein and Gopinath (2014)), they show that a number of mechanisms aside from Kimball demand map into our responsiveness parameter, Γ. Variation in markups arising from variation in firms' market power is the most common alternative to Kimball demand in the pass-through literature. Canonical references are Krugman (1986), Helpman and Krugman (1987), Dornbusch (1987) and more recently Atkeson and Burstein (2008). Since the body of our paper shows extensive results for Kimball demand, we focus here on this market power alternative and show that it also implies a positive correlation between pass-through and price dispersion.

Variation in Market Power

In this setting a discrete number of products and strategic complementarity gives rise to variable markups and markup elasticity, Γ, in the same form as our baseline model. The difference is that Γ is

[64]The logic is based on the envelope theorem. The first-order expected revenue gain from reducing the price cannot be large enough to more than offset the cost increase, because otherwise the previous price could not have been optimal.

determined by different parameters: variation in market power and elasticities of demand and whether there is Betrand or Cournot competition rather than from kinked demand. Otherwise the underlying structure of the problem is the same. See Section 4.2 of Burstein and Gopinath (2014) for full details. We now show that this model naturally generates time varying responsiveness: when competitive pressures are high, responsiveness is higher.

Despite a similar overall structure, since there are a finite number of firms and strategic complementarity, we must check whether the indirect effect of the exchange rate change coming through changes in other firms' prices overturns the basic results in Section 4. As in our baseline model, price changes depend on changes in the exchange rate, idiosyncratic shocks and changes in the industry level price index:

$$\Delta p_{ik} = \frac{\alpha \Delta e_i + \mathit{\Gamma}_i \Delta p_k + \epsilon_i}{1 + \Gamma_i}$$

Assume some common exchange rate variation across firms. Define firm $j's$ common exposure to firm $i's$ exchange rate variation in sector k as $\Delta e_j = \theta \Delta e_i + (1-\theta)\Delta v_j$ with $\Delta v_j \perp \Delta e_i$ for all j. If $\theta = 1$ then firm j is exposed to the same exchange rate variation as firm i and if $\theta = 0$, there is no common exchange rate variation. The most interesting case is if $0 < \theta < 1$ where there is some difference in exposure to the exchange rate between firm i and firm j. This could happen if competing firms within the same industry source inputs from different countries. In this case we can easily show (after some patient algebra) that pass-through is decreasing in Γ_i just as in our baseline case as long as $0 < \theta < 1$ (and $0 < w_{ik} < 1$ but this by construction).[65] In particular,

$$\frac{\partial \left(\frac{\Delta p_{ik}}{\Delta e_i} \right)}{\partial \mathit{\Gamma}_i} = \frac{\alpha \left((1-\theta)(1+\mathit{\Gamma}_i)\mathit{\Gamma}_i \frac{\partial w_{ik}}{\partial \mathit{\Gamma}_i} + (\theta + (1-\theta) w_{ik} - 1) \right)}{(1+\Gamma_i)^2} < 0$$

Thus as long as there are least two firms in a industry and the exchange rates relevant for each firm are not perfectly correlated, pass-through is decreasing in $\mathit{\Gamma}_i$. This is the empirically relevant case since firms import from a variety of different countries with different exchange rate exposure. How does the variance of price changes change with $\mathit{\Gamma}_i$?

$$\Delta p_{ik} = \frac{(\alpha + \alpha\theta \mathit{\Gamma}_i)\,\Delta e_i}{1+\Gamma_i} + \frac{\alpha(1-\theta)\mathit{\Gamma}_i \sum_j w_{jk} \Delta v_j}{1+\Gamma_i} + \frac{\mathit{\Gamma}_i \sum_j w_{jk}\epsilon_j}{1+\Gamma_i} + \frac{\epsilon_i}{1+\Gamma_i}$$

Assume that the variance of exchange rates are the same across firms:[66] $var(\Delta e_i) = var(\Delta v_j) = \sigma_e^2$ and the variance of the idiosyncratic shocks be equal to σ_ϵ^2. Then one can show that $\frac{\partial var(\Delta p_{ik})}{\partial \mathit{\Gamma}_i} < 0$ as long

[65] Here $w_{ik} \equiv \frac{\left(\frac{s_{ik}}{1+\Gamma_i}\right)}{\sum_i \left(\frac{s_{ik}}{1+\Gamma_i}\right)}$, where s_{ik} denotes the market share of firm i in sector k. By construction $\sum_i w_{ik} = 1$. One can show that these assumptions imply that $\frac{\partial w_{ik}}{\partial \mathit{\Gamma}_i} = \frac{\frac{-s_{ik}}{(1+\Gamma_i)^2}\left[\Sigma\left(\frac{s_{ik}}{1+\Gamma_i}\right) - \left(\frac{s_{ik}}{1+\Gamma_i}\right)\right]}{\left(\Sigma\left(\frac{s_{ik}}{1+\Gamma_i}\right)\right)^2} < 0$, making the first term in the above expression negative as well when $0 < \theta < 1$.

[66] This assumption is made for analytical convenience, however, a wide variety of simulations indicate that large deviations from this assumption do not change the relevant parameter bounds. The reason is that empirically the variance of idiosyncratic shocks is much larger than the variance of exchange rates and thus the terms with idiosyncratic shocks in them are what matter. Increasing significantly the size of the variance of the idiosyncratic exchange rate shocks, $var(\Delta v_j)$, or allowing the variance of these shocks to be correlated with w_i has minimal impact.

as $\mathit{\Gamma}_i < min\left\{\frac{1}{1-2\theta}, \frac{2-w_{ik}}{2(\sum_j w_{jk}^2)-w_i}\right\}$.[67] Thus, under reasonable parameter restrictions this model implies a positive relationship between pass-through and dispersion.

B.2 More General Flexible Price Results

In this section, we show that the intuition from our simple framework in Section 4.1, 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:

$$\Delta p_i = -\Gamma_i(\Delta p_i - \Delta p) + \alpha \Delta e_i + \Delta \eta_i$$

which can be rearranged to give:

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

In Section 4.1 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 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:

$$\frac{\Delta p_i}{\Delta e_i} = \frac{\alpha_i}{1+\Gamma_i} + \frac{\Gamma_i}{1+\Gamma_i}\frac{\Delta p}{\Delta e_i} \tag{10}$$

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

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

$$\begin{aligned} \frac{\partial \frac{\Delta p_i}{\Delta e_i}}{\partial \Gamma_i} &= -\frac{\alpha_i}{(1+\Gamma_i)^2} + \frac{1}{(1+\Gamma_i)^2}\frac{\Delta p}{\Delta e_i} \\ &= \frac{\frac{\Delta p}{\Delta e_i} - \alpha_i}{(1+\Gamma_i)^2} < 0 \text{ if } \alpha_i > \frac{\Delta p}{\Delta e_i} \end{aligned} \tag{11}$$

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

[67]For the LHS on the min, given the restrictions on θ, the strictest restriction here is $\mathit{\Gamma}_i < 1$. For the RHS, again, this is $\mathit{\Gamma}_i < 1$ as long as it is well defined (This object is not well defined (it is less than zero) when firm i is very large relative to the rest of the sector and all other firms are very small). In all other cases, this restriction is satisfied and is typically much larger than 1. For example, if all firms were the same size and there were N firms in the industry, the restriction would be $\mathit{\Gamma}_i < 2N - 1$.

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 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 (11). However, this is now an additional effect: a higher Γ_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 (11). 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-pass-through rate is affected by the level of strategic complementarities, Γ_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:

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

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

$$var(\Delta p_i) = \left(\frac{\alpha_i}{1+\Gamma_i}\right)^2 var(\Delta e_i) + \left(\frac{\Gamma_i}{1+\Gamma_i}\right)^2 var(\Delta p) + \left(\frac{1}{1+\Gamma_i}\right)^2 var(\Delta \eta_i) + \frac{\alpha_i \Gamma_i}{(1+\Gamma_i)^2} cov(\Delta e_i, \Delta p) \tag{12}$$

Using this expression, we get that

$$\frac{\partial var(\Delta p_i)}{\partial \Gamma_i} = -\frac{2\alpha_i^2}{(1+\Gamma_i)^3} var(\Delta e_i) + \frac{2\Gamma_i}{(1+\Gamma_i)^3} var(\Delta p) - \frac{2}{(1+\Gamma_i)^3} var(\eta_i) + \frac{\alpha_i(1-\Gamma_i)}{(1+\Gamma_i)^3} cov(\Delta e_i, \Delta p). \tag{13}$$

We now show that under a mild and empirically realistic restriction, the variance of price changes is declining in Γ_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 (12) to get that

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

or

$$var(\eta_i) > \left[(1+\Gamma_i)^2 - \Gamma_i^2\right] var(\Delta p) + \left[(1+\Gamma_i)^2 - \alpha_i^2\right] var(\Delta e_i) - \alpha_i \Gamma_i cov(\Delta e_i, \Delta p) \tag{14}$$

Using (13) we have

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

Substituting the inequality (14) for $var(\eta_i)$ gives

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

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 ($\alpha_i \uparrow$, $\Gamma_i \downarrow$) also increase the variance of price changes.

B.3 Steady-State Calibration

This subsection shows how super-elasticity ε), shock volatility (σ) and import shares (α) are identified in steady-state. As described in Section 5.1.4, we jointly target average pass-through, the R^2 from our MRPT regression and the mean standard deviation of item level price changes. Figure B.4 shows that varying each parameter produces a different patterns of movement between these moments. In this exercise, we hold all parameters at their best-fit calibration and then very one parameter at a time and show its implications for MRPT, R^2 and the standard deviation of price changes. Similar patterns arise if we fix parameters at other values instead, so these relationship are quite robust.

B.4 Cross-Item Indirect Inference

In this section, we repeat our indirect inference exercise but now allowing for permanent firm heterogeneity instead of time-series aggregate shocks. In particular, we allow firms to differ by κ, ε and σ_A. We assume

Figure B.4: Identification of Baseline Parameters

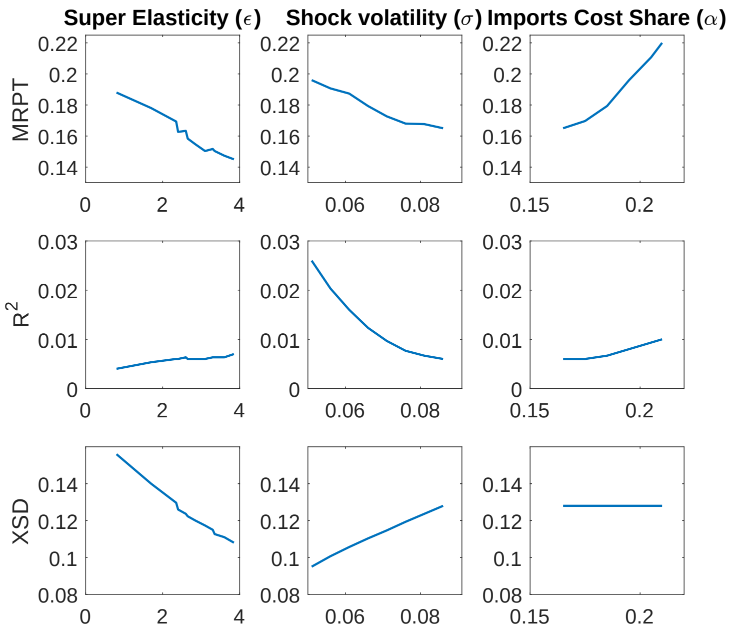

This figure shows how our three target moments (labeled on the left-hand side) vary with parameters (labeled as the titles of each column).

that each parameter takes on one of two values uniformly distributed around the steady-state value.[68] For example, we assume that for a particular firm, κ is either equal to $\kappa_h = .043 + \kappa_\Delta$ or $\kappa_l = .043 - \kappa_\Delta$ 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.[69] 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.

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.

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.[70] Just as in the time-series, this indirect inference estimation

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

[69] 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.

[70] 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.

strongly rejects restricted specifications with no ε variation as well as specifications with any significant heterogeneity in σ. Figure B.5 displays these results visually, showing the best-fit for all fifteen moments as well as the fit of restricted models which shut down various sources of heterogeneity.

Figure B.5: Cross-Item Indirect Inference

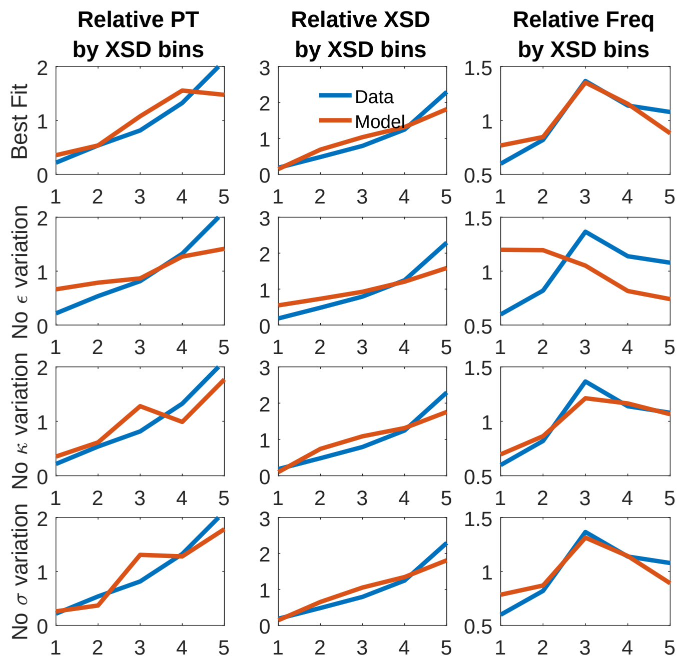

This figure shows the model fit to all fifteen moments as well as the fit of restricted models which shut down various sources of heterogeneity.

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.

B.5 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, for the same reason that changes in α have minimal affect on 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.

Table A1: Alternative Business Cycle Controls

	(1) IQR+Recession Dummy	(2) IQR+GDP growth	(3) IQR+HP filtered GDP	(4) XSD+Recession Dummy	(5) XSD+GDP growth	(6) XSD+ HP filtered GDP
Δe	0.128	0.150	0.143	0.122	0.152	0.140
	(0.012)	(0.011)	(0.011)	(0.012)	(0.012)	(0.012)
IQR$\times\Delta$e	0.049	0.057	0.072			
	(0.010)	(0.009)	(0.010)			
IQR	-0.001	-0.002	-0.002			
	(0.001)	(0.001)	(0.001)			
XSD$\times\Delta$e				0.033	0.043	0.055
				(0.009)	(0.008)	(0.010)
XSD				-0.001	-0.001	-0.001
				(0.001)	(0.001)	(0.001)
Recession Dummy$\times\Delta$e	0.119			0.164		
	(0.034)			(0.033)		
Recession Dummy	-0.008			-0.009		
	(0.002)			(0.002)		
GDP Growth$\times\Delta$e		-0.028			-0.042	
		(0.010)			(0.009)	
GDP Growth		0.000			0.001	
		(0.001)			(0.001)	
HP GDP$\times\Delta$e			0.002			-0.013
			(0.011)			(0.011)
HP GDP			0.002			0.002
			(0.001)			(0.001)
Num obs	129260	129260	129260	129260	129260	129260
R^2	0.039	0.038	0.038	0.038	0.038	0.037

All regressions control for Δ cpi, Δ us gdp, Δ uscpi and allow for exchange rate pass-through to vary with business cycle controls. Monthly recession dummies picked to match NBER dates, GDP growth is real chained quarterly GDP growth and HP filtered GDP is log real GDP level Hodrick-Prescott filtered with a smoothing parameter of 1600. Regressions have country$\times$PSL fixed effects and robust standard errors clustered at the country$\times$PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table A2: Time-Series Results by Country

	(1) OECD	(2) Asia	(3) Eurozone	(4) Canada	(5) Mexico
Δe	0.206	0.147	0.254	0.222	0.075
	(0.015)	(0.019)	(0.029)	(0.043)	(0.055)
IQR$\times\Delta$e	0.058	0.027	0.040	0.141	0.127
	(0.012)	(0.012)	(0.026)	(0.034)	(0.035)
IQR	-0.001	0.000	0.001	-0.004	0.002
	(0.001)	(0.001)	(0.002)	(0.026)	(0.001)
All Ctls	Yes	Yes	Yes	Yes	Yes
Num obs	68478	43590	14591	26309	8269
R^2	0.047	0.052	0.079	0.030	0.016

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs $\times$ Δe, gdp growth, gdp growth$\times$ Δe, SDe, SDe$\times\Delta$e, month dummies, month dummies $\times\Delta$e, t, t$\times\Delta$e, Δ cpi, Δ us gdp, Δ uscpi. See text for additional description. Regressions have country$\times$PSL fixed effects and robust standard errors clustered at the country$\times$PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table A3: Cross-Item Results by Country

	(1) OECD	(2) Asia	(3) Eurozone	(4) Canada	(5) Mexico
Δe	0.257	0.123	0.299	0.279	0.103
	(0.020)	(0.022)	(0.039)	(0.061)	(0.037)
$\text{XSD}_{\text{item}}\times\Delta$e	0.072	0.048	0.099	0.124	0.031
	(0.019)	(0.020)	(0.034)	(0.065)	(0.045)
XSD_{item}	0.002	-0.002	0.004	0.003	0.004
	(0.001)	(0.002)	(0.003)	(0.001)	(0.003)
$\text{Freq}_{\text{item}}\times\Delta$e	0.085	0.012	0.067	0.178	0.076
	(0.016)	(0.014)	(0.030)	(0.055)	(0.035)
$\text{Freq}_{\text{item}}$	-0.001	-0.001	-0.004	-0.001	0.010
	(0.001)	(0.001)	(0.002)	(0.002)	(0.006)
Num obs	68478	43590	14591	26309	8269
R^2	0.048	0.049	0.084	0.031	0.010

All regressions control for Δ cpi, Δ us gdp, Δ uscpi and have country$\times$PSL fixed effects with robust standard errors clustered at the country$\times$PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table A4: Results for Manufactured Goods

	(1) Overall	(2) IQR	(3) IQR+Freq	(4) IQR+All Ctrls	(5) XSD	(6) XSD+Freq	(7) XSD+All Ctrls
Δe	0.156	0.148	0.147	0.176	0.153	0.152	0.180
	(0.012)	(0.011)	(0.011)	(0.015)	(0.012)	(0.012)	(0.015)
IQR×Δe		0.062	0.061	0.042			
		(0.010)	(0.010)	(0.010)			
IQR		-0.002	-0.002	-0.002			
		(0.001)	(0.001)	(0.001)			
XSD×Δe					0.051	0.050	0.030
					(0.009)	(0.009)	(0.009)
XSD					-0.002	-0.002	-0.002
					(0.001)	(0.001)	(0.001)
Freq×Δe			0.011	0.017		0.013	0.021
			(0.009)	(0.011)		(0.009)	(0.009)
Freq			0.003	0.005		0.004	0.004
			(0.001)	(0.001)		(0.001)	(0.001)
All Ctrls	No	No	No	Yes	No	No	Yes
Num obs	129260	129260	129260	129260	129260	129260	129260
R^2	0.035	0.038	0.039	0.040	0.037	0.038	0.039

"All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs × Δe, gdp growth, gdp growth× Δe, SDe, SDe×Δe, month dummies, month dummies ×Δe, t, t×Δe, Δ cpi, Δ us gdp, Δ uscpi. See text for additional description. Regressions have country×PSL fixed effects and robust standard errors clustered at the country×PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table A5: pass-through at Fixed Horizons

	(1) 1 month	(2) 3 month	(3) 6 month	(4) 12 month
Δe	0.037	0.078	0.118	0.125
	(0.006)	(0.011)	(0.017)	(0.024)
IQR×Δe	0.017	0.024	0.032	0.023
	(0.005)	(0.008)	(0.010)	(0.011)
IQR	-0.000	0.001	0.004	0.011
	(0.000)	(0.000)	(0.001)	(0.002)
All Ctrls	Yes	Yes	Yes	Yes
Num obs	354851	335848	304041	249103
R^2	0.009	0.036	0.082	0.136

These show the relationship between dispersion and pass-through without conditioning on price adjustment, at various horizons. This specification allows us to expand our sample to items with 1+ price changes instead of the 2+ in our baseline sample. See Appendix for additional description. "All controls" are frequency of adjustment (freq), frequency of product substitutions (subs), freq and subs × Δe, gdp growth, gdp growth× Δe, SDe, SDe×Δe, month dummies, month dummies ×Δe, t, t×Δe, Δ cpi, Δ us gdp, Δ uscpi. Regressions have country×PSL fixed effects and robust standard errors clustered at the country×PSL level. Dispersion and frequency are standardized so that coefficients represent a one-standard deviation effect. Sample period is October 1993-January 2015.

Table A6: Sample Summary Statistics

	Price Observations	Items	Mean Life	Mean # Changes per item	# Items w/ < 2 changes	Δp 25th percentile	Δp median	Δp75th percentile
All non-imputed	2,324,069	107,549	41.1	8.9	36385	-.03	.002	.04
Exclude comm., intrafirm, nondollar	1,188,630	58,567	34.6	5.1	22826	-.04	.005	.054
Exclude items w/ < 2 price changes	772,341	35,741	38.5	7.1	0	-.041	0.004	0.055

This table shows summary statistics for our baseline sample. Price observations is the total number of month-item price observations, items is the total number of items in the sample, mean life is the average number of months between an item's first and last observation in the data set, mean changes per item calculates the total number of changes for each item and then averages across items, items w/ ¡ 2 changes is just a count of the total number of items with 0 or 1 price change, and the price change percentiles show the 25th, 50th, and 75th percentile of non-zero price changes in each sample. Note that since items sometimes have missing price observations within their sample llife, the total number of price observations in column 1 is less than the number of items times the mean item life.